Steve Suehring, Tim Converse, and Joyce Park

PHP6 and MySQL®

Explore PHP syntax, datatypes, and functions

Create database-driven, dynamic Web sites

Master server-side Web programming

Bible

The book you need to succeed!

PHP 6 and MySQL® 6 Bible

PHP 6 and MySQL® 6 Bible

Steve Suehring

Tim Converse

Joyce Park

About the Authors

Steve Suehring is a technology consultant with a diverse business and computing background. Steve's extensive experience enables him to work cross-functionally within organizations to help create computing architectures that fit the business need. Steve has written several books and magazine articles and contributed to many others. Steve has spoken internationally at user groups and conventions. When he has the chance, Steve plays just about any sport or any musical instrument, some with better success than others.

Tim Converse has written software to recommend neckties, answer questions about space stations, pick value stocks, and make simulated breakfast. He has an M.S. in Computer Science from the University of Chicago, where he taught several programming classes. He is now an engineering manager in the Web search group at Yahoo!.

Joyce Park has an M.A. in history from the University of Chicago, and has worked for several Silicon Valley startups including Epinions, KnowNow, and Friendster. She is a co-lead of the Mod-pubsub Open Source project.

Credits

Acknowledgments

People sometimes ask me how many books I've written. I never have the answer. You see, I've contributed to well over a dozen (maybe two dozen or more) books in one form or another, be it a chapter or two here, a section there, a rewrite of an existing title with much new material, a revision of another edition where the existing material is already pretty good (as was the case for this book), or an original, authored work. The short answer is: I don't know. It's really somewhat difficult to claim that I, alone, wrote a book. At best I put some words down into a word processor and several other people look them over, edit them, change them for both technical and grammatical usage, and the end result is my name on the cover or somewhere in the book, or sometimes not at all.

This brings me to the difficulty at hand. I've written a sufficient number books that writing acknowledgments is becoming a bit mundane. Sure, I'll thank my wife, Rebecca, and son, Jakob, for their patience while I wrote this. I'll thank my family for their continued support. I'll thank the Tueschers, Heins, Leus, and Guthries. I'll thank Jason Keup and Aaron Saray, too. I'll thank my agent Neil Salkind at Studio B., Jim Oliva and John Eckendorf, and the 90fm staff along with Nightmare Squad.

Of course, I'll thank Tim and Rob @ Partners, and Jay, Deb, and Brian, and Andy Hale and Eliot Irons and the SecAdmin team. Kyle Mac always gets mad if I don't include him. There are lot of people at Knob Hill who deserve thanking, and the like. And I'll always thank Mark Little and meek, Pat Dunn, AJ Prowant, and Andy Berkvam. But it's the people that I don't thank that always find me, asking why their name isn't in this book. With that in mind, I'll stop here and let them find me and hope that I write another book where I'll remember to include them. Just a hint: Everyone who was thanked here has paid me.

Contents at a Glance

Contents

Contents

Contents

Contents

Contents

Contents

Contents

Contents

Part IV: Other Databases 549

Chapter 33: Connecting PHP and PostgreSQL 551

Chapter 34: Using PEAR DB with PHP 567

Chapter 35: An Overview of Oracle. 575

Contents

Contents

Contents

Introduction

What Is PHP?

PHP is an open source, server-side, HTML-embedded web-scripting language that is compatible with all the major web servers (most notably Apache). PHP enables you to embed code fragments in normal HTML pages — code that is interpreted as your pages are served up to users. PHP also serves as a "glue" language, making it easy to connect your web pages to server-side databases.

Why PHP?

We devote nearly all of Chapter 1 to this question. The short answer is that it's free, it's open source, it's full featured, it's cross-platform, it's stable, it's fast, it's clearly designed, it's easy to learn, and it plays well with others.

What's New in This Edition?

This book is a new edition of the popular PHP Bible and PHP5 and MySQL Bible series. The book updates the elements from previous versions, where applicable, for PHP 6 and MySQL 6.

New PHP 6 features

Although much of PHP 5's functionality survives unchanged in PHP 6, there have been some changes. Among the ones we cover are:

- Unicode support, making internationalization easier
- Security enhancements such as removing safe_mode and register globals
- Enhancements to the object-oriented interfaces

Who wrote the book?

The first two editions were by Converse and Park, with a guest chapter by Dustin Mitchell and tech editing by Richard Lynch. For the third edition, Clark Morgan took on much of the revision work, with help from Converse and Park as well as from David Wall and Chris Cornell, who also contributed chapters and did technical editing. For this edition, Steve Suehring did revision work with Aaron Saray providing technical editing.

Whom This Book Is For

This book is for anyone who wants to build web sites that exhibit more complex behavior than is possible with static HTML pages. Within that population, we had the following three particular audiences in mind:

- Web site designers who know HTML and want to move into creating dynamic web sites
- Experienced programmers (in C, Java, Perl, and so on) without web experience who want to quickly get up to speed in server-side web programming
- Web programmers who have used other server-side technologies (Active Server Pages, Java Server Pages, or ColdFusion, for example) and want to upgrade or simply add another tool to their kit

We assume that the reader is familiar with HTML and has a basic knowledge of the workings of the web, but we do not assume much programming experience beyond that. To help save time for more experienced programmers, we include a number of notes and asides that compare PHP with other languages and indicate which chapters and sections may be safely skipped. Finally, see our appendixes, which offer specific advice for C programmers, ASP coders, and pure-HTML designers.

This Book Is Not the Manual

The PHP Documentation Group has assembled a great online manual, located at `www.php.net` and served up (of course) by PHP. This book is not that manual or even a substitute for it. We see the book as complementary to the manual and expect that you will want to go back and forth between them to some extent.

In general, you'll find the online manual to be very comprehensive, covering all aspects and functions of the language, but inevitably without a great amount of depth in any one topic. By contrast, we have the leisure of zeroing in on aspects that are most used or least understood and give background, explanations, and lengthy examples.

How the Book Is Organized

This book is divided into five parts, as the following sections describe.

Part I: PHP: The Basics

This part is intended to bring the reader up to speed on the most essential aspects of PHP, with complexities and abstruse features deferred to later parts.

- Chapters 1 through 3 provide an introduction to PHP and tell you what you need to know to get started.

- Chapters 4 through 9 are a guide to the most central facets of PHP (with the exception of database interaction): the syntax, the data types, and the most basic built-in functions.
- Chapter 10 is a guide to the most common pitfalls of PHP programming.

Part II: PHP and MySQL

Part II is devoted both to MySQL and to PHP's interaction with MySQL.

- Chapters 11 and 12 provide a general orientation to web programming with SQL databases, including installation of MySQL.
- Chapter 13 covers Structured Query Language (SQL), and Chapter 14 covers database administration basics.
- Chapter 15 is devoted to PHP functions for MySQL.
- Chapters 16 and 17 are detailed, code-rich case studies of PHP/MySQL interactions.
- Chapters 18 and 19 provide tips and gotchas specific to PHP/MySQL work.

Part III: Advanced Techniques

In this part we cover more advanced features of PHP, usually as self-contained chapters, including object-oriented programming, session handling, exception handling, using cookies, and regular expressions. Chapter 31 is a tour of debugging techniques, and Chapter 32 discusses programming style.

Part IV: Connections

In this part we cover advanced techniques and features that involve PHP talking to other services, technologies, or large bodies of code.

- Chapters 33 through 36 cover PHP's interaction with other database technologies (PostgreSQL, Oracle, PDO, and SQLite).
- Chapters 37 through 42 cover self-contained topics: PHP and e-mail programs, combining PHP with JavaScript, integrating PHP and Java, PHP and XML, PHP-based Web services, and creating graphics with the `gd` image library.

Part V: Case Studies

Here we present three extended case studies that wrap together techniques from various early chapters.

- Chapter 43 takes you through the design and implementation of a weblog.
- Chapter 44 discusses a soup-to-nuts implementation of a novel trivia quiz game.
- Chapter 45 uses the `gd` image library to visualize data from a MySQL database.

Appendices

At the end, we offer three "quick-start" appendixes, for use by people new to PHP but very familiar with either C (Appendix A), Perl (Appendix B), or pure HTML (Appendix C). If you are in any of these three situations, start with the appropriate appendix for an orientation to important differences and a guide to the book. Appendix (D) is a guide to important resources, web sites, and mailing lists for the PHP community. The final appendix (E) is information on the PEAR repository, which is no longer scheduled to be included in PHP 6. However, this information (from a previous edition of the book) may be helpful to someone maintaining a PHP site on an earlier version of PHP or one that uses PEAR.

Conventions Used in This Book

We use a monospaced font to indicate literal PHP code. Pieces of code embedded in lines of text look like `this`, while full code listing lines look as follows:

```
print("this");
```

If the appearance of a PHP-created web page is crucial, we include a screenshot. If it is not, we show textual output of PHP in monospaced font. If we want to distinguish the PHP output as seen in your browser from the actual output of PHP (which your browser renders), we call the former *browser output*.

If included in a code context, *italics* indicate portions that should be filled in appropriately, as opposed to being taken literally. In normal text, an *italicized* term means a possibly unfamiliar word or phrase.

What the Icons Mean

Icons similar to the following example are sprinkled liberally throughout the book. Their purpose is to visually set off certain important kinds of information.

TIP   Tip icons indicate PHP tricks or techniques that may not be obvious and that enable you to accomplish something more easily or efficiently.

NOTE   Note icons usually provide additional information or clarification but can be safely ignored if you are not already interested. Notes in this book are often audience-specific, targeted to people who already know a particular programming language or technology.

CAUTION   Caution icons indicate something that does not work as advertised, something that is easily misunderstood or misused, or anything else that can get programmers into trouble.

CROSS-REF   We use this icon whenever related information is in a different chapter or section.

Part I

Introducing PHP

Chapter 1

Why PHP and MySQL?

This first chapter is an introduction to PHP, MySQL, and the interaction of the two. In it, we'll try to address some of the most common questions about these tools, such as "What are they?" and "How do they compare to similar technologies?" Most of the chapter is taken up with an enumeration of the many, many reasons to choose PHP, MySQL, or the two in tandem. If you're a techie looking for some ammunition to lob at your PHB ("Pointy-Haired Boss," for those who don't know the Dilbert cartoons) or a manager asking yourself what is this P-whatever thing your geeks keep whining to get, this chapter will provide some preliminary answers.

What Is PHP?

PHP is the web development language written by and for web developers. PHP stands for *PHP: Hypertext Preprocessor*. The product was originally named *Personal Home Page Tools*, and many people still think that's what the acronym stands for, but as it expanded in scope, a new and more appropriate (albeit GNU-ishly recursive) name was selected by community vote. PHP is currently in its sixth major rewrite, called PHP6 or just plain PHP.

PHP is a server-side scripting language, usually used to create web applications in combination with a web server, such as Apache. PHP can also be used to create command-line scripts akin to Perl or shell scripts, but such use is much less common than PHP's use as a web language.

Strictly speaking, PHP has nothing to do with layout, events, on-the-fly Document Object Model (DOM) manipulation, or really anything about the look and feel of a web page. In fact, most of what PHP does is invisible to the end user. Someone looking at a PHP page will not necessarily be able to tell that it was not written purely in Hypertext Markup Language (HTML), because the result of PHP is usually HTML.

What Is MySQL?

MySQL (pronounced My Ess Q El) is an open source, SQL relational database management system (RDBMS) that is free for many uses (more detail on that later). Early in its history, MySQL occasionally faced opposition because of its lack of support for some core SQL constructs such as subselects and foreign keys. Ultimately, however, MySQL found a broad, enthusiastic user base for its liberal licensing terms, perky performance, and ease of use. Its acceptance was aided in part by the wide variety of other technologies such as PHP, Perl, Python, and the like that have encouraged its use through stable, well-documented modules and extensions.

Databases are generally useful, perhaps the most consistently useful family of software products (the "killer product") in modern computing. Like many competing products, both free and commercial, MySQL isn't a database until you give it some structure and form. You might think of this as the difference between a database and an RDBMS (that is, RDBMS plus user requirements equal a database).

There's lots more to say about MySQL, but then again, there's lots more space in which to say it.

Deciding on a Web Application Platform

There are many platforms upon which web applications can be built. This section compares PHP to a few other platforms and highlights some of PHP's and MySQL's strengths.

Cost

PHP is one of the "P's" in the popular LAMP stack. The LAMP stack refers to the popular combination of Linux, Apache, MySQL, and PHP/Perl/Python that runs many web sites and powers many web applications. Many of the components of the LAMP stack are free, and PHP is no exception. PHP is free, as in there is no cost to develop in and run programs made with PHP. Though MySQL's license and costs have changed, you can obtain the Community Server edition for free. MySQL offers several levels of support contracts for their database server. More information can be obtained at `www.mysql.com`. Both PHP and MySQL run on a variety of platforms, including many variants of Linux, Microsoft Windows, and others. Running on an operating system such as Linux gives the opportunity for a completely free web application platform, with no up-front costs.

Of course, when talking about software development and application platforms, the up-front cost of software licensing is only a portion of the total cost of ownership (TCO). Years of real-world experience with Linux, Apache, MySQL, and PHP in production environments has proved that the total cost of maintaining these platforms is lower, many times much lower, than maintaining an infrastructure with proprietary, non-free software.

Ease of Use

When compared to many other programming languages, PHP makes it easy to develop powerful web applications quickly (this is a blessing and a curse). Many of the most useful specific functions (such as those for opening a connection to an Oracle database or fetching e-mail from an Internet Message Access Protocol [IMAP] server) are predefined for you. A lot of complete scripts are waiting out there for you to look at as you're learning PHP.

Most advanced PHP users (including most of the development team members) are diehard hand-coders. They tend to share certain gut-level, subcultural assumptions — for instance, that hand-written code is beautiful and clean and maximally browser-compatible and therefore the only way to go — that they do not hesitate to express in vigorous terms. The PHP community offers help and trades tips mostly by e-mail, and if you want to participate, you have to be able to parse plain-text source code with facility. Some WYSIWYG users occasionally ask list members to diagnose their problems by looking at their web pages instead of their source code, but this rarely ends well.

That said, let us reiterate that PHP really is easy to learn and write, especially for those with a little bit of experience in a C-syntaxed programming language. It's just a little more involved than HTML. This small learning curve means that relatively inexperienced programmers can sometimes make mistakes that turn into large security issues. This is the curse of PHP. While this book has no specific chapter dedicated to security, I feel that security needs to be applied at every layer, during every phase of programming. Therefore dedicating a single chapter would not do justice to the importance of web application security.

If you have no relational database experience, or are coming from an environment such as Microsoft Access, MySQL's command-line interface and lack of implicit structure may at first seem a little daunting. MySQL has a few GUI (graphical user interface) tools to help work with databases. None of the GUI tools is a substitute for learning a little theory and employing good design practices, but that is a subject for another chapter.

HTML-embeddedness

PHP can be embedded within HTML. In other words, PHP pages are ordinary HTML pages that escape into PHP mode only when necessary. Here is an example:

```
<HEAD>
<TITLE>Example.com greeting</TITLE>
</HEAD>
<BODY>
<P>Hello,
<?php
// We have now escaped into PHP mode.
// Instead of static variables, the next three lines
// could easily be database calls or even cookies;
// or they could have been passed from a form.
$firstname = 'Joyce';
$lastname = 'Park';
```

```
$title = 'Ms.';
echo "$title $lastname";
// OK, we are going back to HTML now.
?>
.  We know who you are!  Your first name is <?php echo
$firstname; ?>.</P>

<P>You are visiting our site at <?php echo date('Y-m-d H:i:s');
?></P>

<P>Here is a link to your account management page:  <A
HREF="http://www.example.com/accounts/<?php echo
"$firstname$lastname"; ?>/"><?php echo $firstname; ?>'s account
management page</A></P>
</BODY>
</HTML>
```

When a client requests this page, the web server *preprocesses* it. This means it goes through the page from top to bottom, looking for sections of PHP, which it will try to resolve. For one thing, the parser will suck up all assigned variables (marked by dollar signs) and try to plug them into later PHP commands (in this case, the echo function). If everything goes smoothly, the preprocessor will eventually return a normal HTML page to the client's browser, as shown in Figure 1-1.

FIGURE 1-1

A result of preprocessed PHP

If you peek at the source code from the client browser (select Source or Page Source from the View menu, it will look like this:

```
<HEAD>
<TITLE>Example.com greeting</TITLE>
```

```
</HEAD>
<BODY>
<P>Hello,
Ms. Park
.  We know who you are!  Your first name is Joyce.</P>

<P>You are visiting our site at 2002-04-21 19:34:24</P>

<P>Here is a link to your account management page:  <A HREF="http://
www.example.com/accounts/JoycePark/">Joyce's account management page</
A></P>
</BODY>
</HTML>
```

This code is exactly the same as if you were to write the HTML by hand. So simple!

The HTML-embeddedness of PHP has many helpful consequences:

- PHP can quickly be added to code produced by WYSIWYG editors.
- PHP lends itself to a division of labor between designers and programmers.
- Every line of HTML does not need to be rewritten in a programming language.
- PHP can reduce labor costs and increase efficiency because of its shallow learning curve and ease of use.

Cross-platform compatibility

PHP and MySQL run native on every popular flavor of Linux/Unix (including Mac OS X) and Microsoft Windows. A huge percentage of the world's Hypertext Transfer Protocol (HTTP) servers run on one of these two classes of operating systems.

PHP is compatible with the leading web servers: Apache HTTP Server for Linux/Unix and Windows and Microsoft Internet Information Server. It also works with several lesser-known servers. Specific web server compatibility with MySQL is not required, since PHP will handle all the dirty work for you.

Stability

The word *stable* means two different things in this context:

- The server doesn't need to be rebooted or restarted often.
- The software doesn't change radically and incompatibly from release to release.

To our advantage, both of these connotations apply to both MySQL and PHP.

Apache Server is generally considered the most stable of major web servers, with a reputation for enviable uptime percentages. Most often, a server reboot isn't required for each setting change. PHP inherits this reliability; plus, its own implementation is solid yet lightweight.

7

PHP and MySQL are also both stable in the sense of feature stability. Their respective development teams have thus far enjoyed a clear vision of their project and refused to be distracted by every new fad and ill-thought-out user demand that comes along. Much of the effort goes into incremental performance improvements, communicating with more major databases, or adding better OOP support. In the case of MySQL, the addition of reasonable and expected new features has hit a rapid clip. For both PHP and MySQL, such improvements have rarely come at the expense of compatibility.

Many extensions

PHP makes it easy to communicate with other programs and protocols. The PHP development team seems committed to providing maximum flexibility to the largest number of users.

Database connectivity is especially strong, with native-driver support for about 15 of the most popular databases plus Open DataBase Connectivity (ODBC). In addition, PHP supports a large number of major protocols such as POP3, IMAP, and LDAP. Earlier versions of PHP added support for Java and distributed object architectures (Component Object Model [COM] and Common Object Request Broker Architecture [CORBA]), making n-tier development a possibility for the first time, fully incorporated GD graphics library and revamped Extensible Markup Language (XML) support with DOM and simpleXML.

Fast feature development

Users of proprietary web development technologies can sometimes be frustrated by the glacial speed at which new features are added to the official product standard to support emerging technologies. With PHP, this is not a problem. All it takes is one developer, a C compiler, and a dream to add important new functionality. This is not to say that the PHP team will accept every random contribution into the official distribution without community buy-in, but independent developers can and do distribute their own extensions that may later be folded into the main PHP package in more or less unitary form. For instance, Dan Libby's elegant xmlrpc-epi extension was adopted as part of the PHP distribution in version 4.1, a few months after it was first released as an independent package.

PHP development is also constant and ongoing. Although there are clearly major inflection points, such as the transition between PHP4 and PHP5, these tend to be most important deep in the guts of the parser — people were actually working on major extensions throughout the transition period without critical problems. Furthermore, the PHP group subscribes to the open source philosophy of "release early, release often," which gives developers many opportunities to follow along with changes and report bugs.

Not proprietary

The history of the personal computer industry to date has largely been a chronicle of proprietary standards: attempts to establish them, clashes between them, their benefits and drawbacks for the consumer, and how they are eventually replaced with new standards.

In the past few years the Internet has demonstrated the great convenience of voluntary, standards-based, platform-independent compatibility. E-mail, for example, works so well because it enjoys a clear, firm standard to which every program on every platform must conform. New developments that break with the standard (for example, HTML-based e-mail stationery) are generally regarded as deviations, and their users find themselves having to bear the burdens of early adoption.

Furthermore, customers (especially the big-fish businesses with large systems) are fed up with spending vast sums to conform to a proprietary standard only to have the market uptake not turn out as promised. Much of the current momentum toward XML and web services is driven by years of customer disappointment with Java RMI (Remote Method Invocation), CORBA, COM, and even older proprietary methods and data formats.

Right now, software developers are in a period of experimentation and flux concerning proprietary versus open standards. Companies want to be sure that they can maintain profitability while adopting open standards. There have been some major legal conflicts related to proprietary standards, which are still being resolved. These could eventually result in mandated changes to the codebase itself or even affect the futures of the companies involved. In the face of all this uncertainty, a growing number of businesses are attracted to solutions that they know will not have these problems in the foreseeable future.

PHP is in a position of maximum flexibility because it is, so to speak, *antiproprietary*. It is not tied to any one server operating system, unlike Active Server Pages. It is not tied to any proprietary cross-platform standard or middleware, as is Java Server Pages or ColdFusion. It is not tied to any one browser or implementation of a programming language or database. PHP isn't even doctrinaire about working only with other open source software. This independent but cooperative pragmatism should help PHP ride out the stormy seas that seem to lie ahead.

Strong user communities

PHP is developed and supported in a collaborative fashion by a worldwide community of users. Some animals (such as the core developers) are more equal than others, but that's hard to argue with, because they put in the most work, had the best ideas, and have managed to maintain civil relationships with the greatest number of other users.

The main advantage for most new users is technical support without charge, without boundaries, and without the runaround. People on the mailing list are available 24/7/52 to answer your questions, help debug your code, and listen to your gripes. The support is human and real. PHP community members might tell you to read the manual, take your question over to the appropriate database mailing list, or just stop your whining — but they'll never tell you to wipe your C drive and then charge you for the privilege. Often, they'll look at your code and tell you what you're doing wrong or even help you design an application from the ground up.

As you become more comfortable with PHP, you may wish to contribute. Bug tracking, offering advice to others on the mailing lists, posting scripts to public repositories, editing documentation, and, of course, writing C code are all ways you can give back to the community.

MySQL, while open source licensed for non-redistributive uses, is somewhat less community driven in terms of its development. Nevertheless, it benefits from a growing community of users who are actively listened to by the development team. Rarely has a software project responded so vigorously to community demand, and the community of users can be extremely responsive to other users who need help. It's a point of pride with a lot of SQL gurus that they can write the complicated queries that get you the results you are looking for but had struggled with for days. In many cases, they'll help you for nothing more than the enduring, if small, fame that comes with the archived presence of their name on Google Groups. Try comparing *that* with $100 per incident support.

Summary

PHP and MySQL, individually or together, aren't the panacea for every web development problem, but they present a lot of advantages. PHP is built by web developers for web developers and supported by a large and enthusiastic community. MySQL is a powerful standards-compliant RDBMS that comes in at an extremely competitive price point, even more so if you qualify for free use. Both technologies are clear-cut cases of the community banding together to address its own needs.

Chapter 2

Server-Side Scripting Overview

This chapter is about server-side scripting and its relationship to both static HTML and common client-side technologies. By the end, you can expect to gain a clear understanding of what kinds of things PHP can and cannot do for you, along with a general understanding of how it interacts with client-side code (JavaScript, Java applets, Flash, style sheets, and the like).

Static HTML

The most basic type of web page is a completely static, text-based one, written entirely in HTML. Take the simple HTML-only page that Figure 2-1 shows as an example.

The following example displays the source code for the web page shown in Figure 2-1:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN"
"http://www.w3.org/TR/html4/strict.dtd">
<html>
<head>
<title>Selected Constellations</title>
</head>
<body>
<h1>Constellations</h1>
<ul>
<li><a href="Aquila.html">Aquila</a></li>
<li><a href="Bootes.html">Bootes</a></li>
<li><a href="Cassiopeia.html">Cassiopeia</a></li>
```

```
<li><a href="Cygnus.html">Cygnus</a></li>
<li><a href="Deneb.html">Deneb</a></li>
<li><a href="Draco.html">Draco</a></li>
<li><a href="Gemini.html">Gemini</a></li>
<li><a href="Leo.html">Leo</a></li>
<li><a href="Libra.html">Libra</a></li>
<li><a href="Lynx.html">Lynx</a></li>
<li><a href="Orion.html">Orion</a></li>
<li><a href="Pegasus.html">Pegasus</a></li>
<li><a href="Perseus.html">Perseus</a></li>
<li><a href="Pisces.html">Pisces</a></li>
<li><a href="Taurus.html">Taurus</a></li>
<li><a href="Ursa_Major.html">Ursa Major</a></li>
<li><a href="Ursa_Minor.html">Ursa Minor</a></li>
<li><a href="Vega.html">Vega</a></li>
</ul>
</body>
</html>
```

FIGURE 2-1

A static HTML example

Client-Side Technologies

The most common additions to plain HTML are on the client side. These add-ons include formatting extensions, such as Cascading Style Sheets (CSS) and Dynamic HTML; client-side scripting languages, such as JavaScript; VBScript; Java applets; and Flash. Support for all these technologies is (or is not, as the case may be) built into the web browser. They perform the tasks described in Table 2-1, with some overlap.

TABLE 2-1

Client-Side HTML Extensions

Client-Side Technology	Main Use	Example Effects
Cascading Style Sheets, Dynamic HTML	Formatting pages: controlling size, color, placement, layout, timing of elements	Overlapping, different colored/sized fonts Layers, exact positioning
Client-side scripting (JavaScript, VBScript)	Event handling: controlling consequences of defined events	Link that changes color on mouseover Mortgage calculator
Java applets	Delivering small standalone applications	Moving logo Crossword puzzle
Flash animations	Animation	Short cartoon film

The page shown in Figure 2-2 is based on the same content as that in Figure 2-1. As you can see from the following source code, however, this example adds a bit of styling with basic inline CSS.

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN" "http://www.w3.org/
TR/html4/strict.dtd">
<html>
<head>
<STYLE TYPE="text/css">
BODY, P {color: black; font-family: verdana; font-size: 10 pt}
H1 {margin-top: 10; color: black; font-family: arial; font-size: 12 pt}
H2 {margin-bottom: -10; color: black; font-family: verdana; font-size:
18 pt}
A:link, A:visited    {color: #000080; text-decoration: none}
</STYLE>
<title>Selected Constellations</title>
</head>
<body>
<h1>Constellations</h1>
```

13

```
<ul>
<li><a href="Aquila.html">Aquila</a></li>
<li><a href="Bootes.html">Bootes</a></li>
<li><a href="Cassiopeia.html">Cassiopeia</a></li>
<li><a href="Cygnus.html">Cygnus</a></li>
<li><a href="Deneb.html">Deneb</a></li>
<li><a href="Draco.html">Draco</a></li>
<li><a href="Gemini.html">Gemini</a></li>
<li><a href="Leo.html">Leo</a></li>
<li><a href="Libra.html">Libra</a></li>
<li><a href="Lynx.html">Lynx</a></li>
<li><a href="Orion.html">Orion</a></li>
<li><a href="Pegasus.html">Pegasus</a></li>
<li><a href="Perseus.html">Perseus</a></li>
<li><a href="Pisces.html">Pisces</a></li>
<li><a href="Taurus.html">Taurus</a></li>
<li><a href="Ursa_Major.html">Ursa Major</a></li>
<li><a href="Ursa_Minor.html">Ursa Minor</a></li>
<li><a href="Vega.html">Vega</a></li>
</ul>
</body>
</html>
```

FIGURE 2-2

An example of HTML plus CSS.
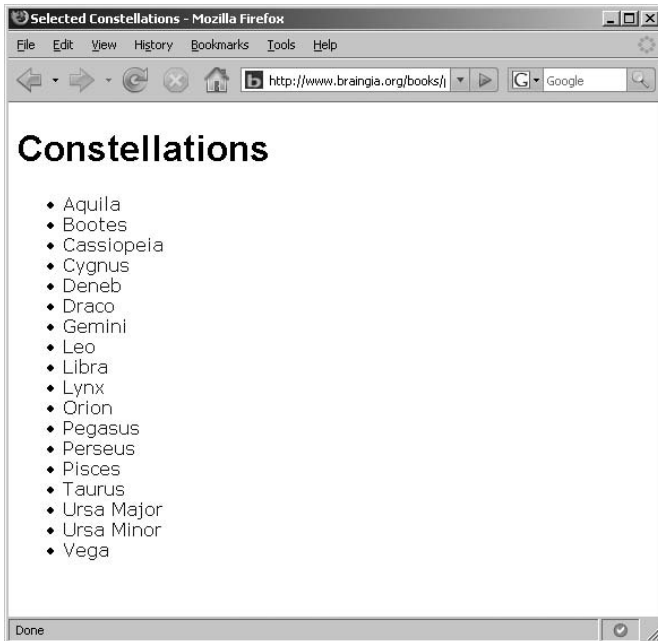
Unfortunately, the best thing about client-side technologies is also the worst thing about them: They depend entirely on the browser. Wide variations exist in the capabilities of each browser and even among versions of the same brand of browser. Individuals can also choose to configure their own browsers in awkward ways: Some people disable JavaScript for security reasons, for example, which makes it impossible for them to view sites that use JavaScript incorrectly or with little care.

The savvy web developer should also consider the implications of device-based browsing, universal accessibility, and a global audience. The stubborn unwillingness of the public to upgrade is the bane of client-side developers, causing them to frequently suffer screaming nightmares and/or existential meltdowns in the dark, vulnerable hours before dawn. The bottom-line irony is that, even after almost 15 years of explosive web progress, the only thing that a developer can absolutely, positively know that the client is going to see is plain text-based HTML (or, rather, the subset of HTML that's widely supported and has stood the tests of time and usefulness).

Server-Side Scripting

Client-side scripting is the glamorous, eye-catching part of web development. In contrast, server-side scripting is invisible to the user. Pity the poor server-side scripters, toiling away in utter obscurity, trapped in the no-man's land between the web server and the database while their arty brethren brazenly flash their wares before the public gaze.

Server-side web scripting is mostly about connecting web sites to backend servers, processing data and controlling the behavior of higher layers such as HTML and CSS. This enables the following types of two-way communication:

- ■ **Server to client**: Web pages can be assembled from backend-server output.
- ■ **Client to server**: Customer-entered information can be acted upon.

Common examples of client-to-server interaction are online forms with some drop-down lists (usually the ones that require you to click a button) that the script assembles dynamically on the server.

Server-side scripting products consist of two main parts: the scripting language and the scripting engine (which may or may not be built into the web server). The engine parses and interprets pages written in the language.

The following code shows a simple example of server-side scripting — a page assembled on the fly from a database. We include database calls (which we don't get around to explaining until Part II of this book) and leave out some of the included files, because we intend this example to show the final product of PHP rather than serve as a piece of working code.

The following PHP code shows the source on the server:

```php
<?php

require_once('db-config.inc.');
```

```
$dbh = mysql_connect(DB_HOST,DB_USER,DB_PASSWORD) or die("Unable to
connect to database.");
mysql_select_db('webdb') or die("Cannot access database.");
$query = "SELECT pagetitle FROM sitepages
            WHERE site = 'braingia.org'
            AND page_id = '1'";
$qresult = mysql_query($query) or die("Unable to query database.");
$title = mysql_fetch_array($qresult);

<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN" "http://www.w3.org/
TR/html4/strict.dtd">
<html>
<head>
<STYLE TYPE="text/css">
BODY, P {color: black; font-family: verdana; font-size: 10 pt}
H1 {margin-top: 10; color: black; font-family: arial; font-size: 12 pt}
H2 {margin-bottom: -10; color: black; font-family: verdana; font-size:
18 pt}
A:link, A:visited    {color: #000080; text-decoration: none}
</STYLE>
<title><?php echo $title[0] ?></title>
</head>
<body>
<h1>$title[0]</h1>
<ul>
<?php
$linksQuery = "SELECT description,href FROM sitepagedata
                WHERE site = 'braingia.org'
                AND pagetitle = '{$title}'";
$linksResult = mysql_query($linksQuery) or die("Unable to query
database.");
while ($row = mysql_fetch_array($linksResult)) {
    print "<li><a href=\"{$row[1]}\">$row[0]</a></li>\n";
}
?>
</ul>
</body>
</html>
```

This particular page isn't significantly more impressive to look at than the version shown in Figure 2-2.

Compare the version with the PHP code to the HTML versions shown earlier in the chapter. The source code that uses PHP is shorter because it retrieves the information from a database. Nevertheless, this server-side code is never viewable by end users. The version that they see is exactly the same as the HTML shown earlier. The only evidence that it's a PHP file is the filename extension, .php. All the heavy lifting happens before the code gets shoved down the pipe to the client. After emerging from the web server, the code appears on the other end as normal HTML

and JavaScript, which also means that you can't tell which server-side scripting language was used unless something in the header or URL gives it away (which usually is the case, as the page you are requesting often ends with `.jsp` or `.php`). These scripts, incidentally, were written in PHP using the MySQL database as backend; you can learn all about these techniques in Part II of this book.

Server-Side or Client-Side?

There are client-side methods and server-side methods to accomplish many tasks. When sending e-mail, for example, the client-side way is to open up the mail client software with a preaddressed blank e-mail message after the user clicks a MAILTO link. The server-side method is to make the user fill out a form, and the contents are formatted as an e-mail that is sent via a Simple Mail Transfer Protocol (SMTP) server (which very well could be the same machine that the server-side script is executing on). You can also choose between client methods and server methods of browser-sniffing, form validation, drop-down lists, and arithmetic calculation. Sometimes you see subtle but meaningful differences in functionality (server-side drop-downs can be assembled dynamically; client-side cannot) but not always.

How to choose? Know your audience. Server-side methods are generally a bit slower at runtime because of the extra transits they must make, but they don't assume anything about your visitor's browser capabilities and take less developer time to maintain.

What Is Server-Side Scripting Good For?

Server-side scripting languages such as PHP perfectly serve most of the truly useful aspects of the web, such as the items in this list:

- Content sites (both production and display)
- Community features (forums, bulletin boards, and so on)
- E-mail (web mail, mail forwarding, and sending mail from a web application)
- Customer-support and technical-support systems
- Advertising networks
- Web-delivered business applications
- Directories and membership rolls
- Surveys, polls, and tests
- Filling out and submitting forms online
- Personalization technologies
- Groupware

- Catalog, brochure, and informational sites
- Games (for example, chess) with lots of logic but simple/static graphics
- Any other application that needs to connect a backend server (database, Lightweight Directory Access Protocol [LDAP], and so on) to a web server

PHP can handle all these essential tasks — and then some.

But enough rhetoric! Now that you have a grasp of the differences between client-side and server-side technologies, you can get on to the practical stuff. In Chapter 3, we show you how to get, install, and configure PHP for yourself (or find someone to do it for you).

Summary

To understand what PHP (or any server-side scripting technology) can do for you, having a firm grasp on the division of labor between client and server is crucial. In this chapter, we worked through examples of plain, static HTML; HTML with client-side additions such as JavaScript and Cascading Style Sheets; and PHP-generated web pages as viewed from both the server and the client.

Client-side scripting can be visually attractive and quickly responsive to user inputs, but anything beyond the most basic HTML is subject to browser variation. Static client-side scripts also require more developer time to maintain and update, because pages cannot be dynamically generated from a constantly changing datastore. Server-side programming and scripting languages, such as PHP, can connect databases and other servers to web pages.

Chapter 3

Getting Started with PHP

In this chapter, we'll give detailed directions for installing PHP and finish with a few tips on finding the right development tool. By the end of the chapter, you should be ready to write your first script.

Installing PHP

This section looks at the installation of PHP onto a computer. If you're going to be using a hosting provider that provides PHP or if you have a friendly sysadmin who has installed PHP for you, then this section will be of limited usefulness. PHP runs on various platforms, including Linux, various Unix flavors, Microsoft Windows, and Mac OS X. Linux is the most popular platform for PHP, and when combined with the Apache web server, and MySQL forms the acronym LAMP (although the "P" can also be Perl or Python).

If you plan to install PHP on Windows, you'll also need:

- A working PHP-supported web server. Under previous versions of PHP, IIS/PWS was the easiest choice because a module version of PHP was available for it; but PHP now has added a much wider selection of modules for Windows. These days, Apache works very well with Windows, so we'll be focusing on PHP with Apache on Windows.

- The PHP Windows binary distribution (download it at `www.php .net/downloads.php`)

- A utility to unzip files (search `http://download.cnet.com` for PC file compression utilities), if your version of Windows doesn't include such a utility.

19

If you plan to install PHP on Linux, you may be able to take advantage of your distribution's PHP package. Most Linux distributions, including Red Hat, Debian, SuSE, and Ubuntu, include PHP as an available package, and, where possible, you should use the distribution's official PHP package.

There are certain instances where you need to compile PHP from source, in order to take advantage of a bleeding-edge feature, for example, but these are the rare exceptions. It is much easier and much more stable to use the distribution's PHP package.

Additionally, you need a web server that supports PHP. Most of the time this will be the Apache web server, but others work well with PHP. For this book, we'll be concentrating on Apache as the web server of choice. Therefore, you'll need to install Apache from your distribution, as well.

Installation procedures

Because of PHP's strong commitment to cross-platform operability, there are far too many specific installation methods to fully list here. We have tried to cover what we believe to be the most popular platforms for PHP, but trying to write the installation instructions for every possible operating system and web server would have resulted in a prohibitively long chapter.

Furthermore, while PHP installation procedures under Unix have been stable for years, Windows installs have gone through quite a bit of flux since PHP4 was first released. Part of this is the result of actions on the part of the PHP team; part of this is because of changes in the Windows product line. PHP also runs on Macintosh OS X, and that installation has only fairly recently stabilized.

In response to such rapid change, we can only caution you that for the freshest information on installation you should visit the PHP web site (`www.php.net/docs.php`) on each download. Even if you've installed PHP a gazillion times before, there might be something new and different on the gazillion-and-first occasion.

For those who have already successfully built an earlier version of PHP, the procedure is exactly the same — only it takes a lot longer than before.

CAUTION **Your Red Hat, Mandrake, or SuSE Linux installation may have come with RPM versions of Apache and PHP, or your Debian Linux may have come with a deb package. You** *must* **remove these packages before compiling your new PHP! In addition, you may have RPM or apt versions of third-party servers, such as MySQL or PostgreSQL, which are generally installed differently from their source counterparts. If you encounter problems, look in the documentation for installation locations, or uninstall the packages and reinstall from scratch. Nevertheless, I strongly recommend using the distribution's version of the package unless you have specific reasons for doing otherwise. If you choose to compile your own versions of PHP and Apache from source then you must maintain them by hand. This means that each and every time a security update is released for either, or for a library touching either, PHP or Apache, you need to recompile the server in order to remain up to date. Otherwise, just use the distribution's package. They'll maintain the security updates, leaving you to concentrate on things like programming PHP!**

The following procedures give an overview of PHP installation on CentOS and Debian. As of this writing, the only version of PHP officially available with these distributions is PHP5. We expect these instructions to be valid when PHP6 becomes available with the distributions.

Installing PHP on CentOS

The YellowDog Update Manager (yum) is available with CentOS and is somewhat like the dpkg and apt toolset from Debian. Therefore, installation of PHP and Apache on CentOS is rather trivial. From the command-line as root, type:

```
yum install php
```

Doing so will cause the yum system to examine the system, gather any prerequisites, and inform you of the installation's progress. Our example system is a fresh CentOS 5.1 install with a minimal package set. Therefore, yum needs to install several prerequisites, and a summary is shown.

After downloading the prerequisites (if necessary), yum will go about its business and install PHP. Part of the install includes Apache, known as "httpd" in CentOS terminology. Apache 2 is installed as part of the installation of PHP.

Apache isn't started by default. To start it, run:

```
/etc/init.d/httpd start
```

While Apache is installed, it is firewalled by default in CentOS, meaning that you can't get to the web server through its default protocol and port, tcp/80. To alleviate this problem, edit /etc/sysconfig/iptables and add this line, second from the bottom:

```
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 80 -j
ACCEPT
```

The final file looks like this:

```
# Firewall configuration written by system-config-securitylevel
# Manual customization of this file is not recommended.
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
:RH-Firewall-1-INPUT - [0:0]
-A INPUT -j RH-Firewall-1-INPUT
-A FORWARD -j RH-Firewall-1-INPUT
-A RH-Firewall-1-INPUT -i lo -j ACCEPT
-A RH-Firewall-1-INPUT -p icmp --icmp-type any -j ACCEPT
-A RH-Firewall-1-INPUT -p 50 -j ACCEPT
-A RH-Firewall-1-INPUT -p 51 -j ACCEPT
-A RH-Firewall-1-INPUT -p udp --dport 5353 -d 224.0.0.251 -j ACCEPT
-A RH-Firewall-1-INPUT -p udp -m udp --dport 631 -j ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m tcp --dport 631 -j ACCEPT
```

```
-A RH-Firewall-1-INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 22 -j
ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 80 -j
ACCEPT
-A RH-Firewall-1-INPUT -j REJECT --reject-with icmp-host-prohibited
COMMIT
```

Restart the iptables firewall by running:

```
/etc/init.d/iptables restart
```

With that, you'll be able to access your web server with PHP enabled by visiting `http://your
.ip.address/` in the browser. For example, my CentOS computer is 192.168.1.155 and so pointing
to that in the web browser looks like this:

```
http://192.168.1.155
```

You may also want to install MySQL through the yum installer and the PHP/MySQL libraries:

```
yum install mysql php-mysql mysql-server mysql-devel
```

Installing PHP on Debian

Installation of PHP (or really anything) on Debian is probably the easiest and most manageable of all
Linux distributions with which I've worked (and that's more than a few). Installation of the Debian
PHP package is done through the apt-get utility:

```
apt-get install libapache2-mod-php5
```

NOTE   **This example shows the installation of the PHP5 module on Debian because the PHP6
module was not yet available at the time of this writing.**

This will install not only the PHP module for Apache 2 but also Apache 2 itself, if the web server
software hasn't already been installed.

Once installed, the web server is ready to use. You'll find the default location for PHP files at `/var/
www/ apache2-default/`, though that location may change in future releases of Debian.

Installing PHP from source

In the following directions, you will type the code fragments into each shell prompt, substituting the
version of software shown in the examples for the version that you're compiling.

You'll need a C compiler, with GCC being a good choice. On Debian you can install gcc by typing
`apt-get install gcc`, whereas on CentOS you can install GCC by typing `yum install gcc`.

You'll also need ICU (International Components for Unicode) for Unicode support. On CentOS, this
is installed with `yum install icu libicu-devel`.

Finally, you'll also need development libraries for libxml, which can be installed on CentOS through the libxml2-devel package, `yum install libxml2-devel`.

If you'll be using MySQL you can install it and the libraries from the command line with the yum installer:

```
yum install mysql mysql-server mysql-devel
```

TIP   Remember to log in as the root user first if you are installing in a root-owned directory. Remember to stop and uninstall your previous Apache server if you had one.

To start your build, just follow these steps:

1. If you haven't already done so, unzip and untar your Apache source distribution. Unless you have a reason to do otherwise, `/usr/local` is the standard place to do so.

```
tar -zxvf httpd-2.2.x.tar.gz
```

2. Build the Apache server: If you are installing somewhere other than `/usr/local`, this is the time to say so with the `--prefix` flag as follows. If you are installing in `/usr/local`, don't worry that the `apache` directory mentioned in a moment doesn't exist — it will by the end of the build process. The `--enable-so` flag will allow Apache to load PHP support (and many other things) as a module called a *Shared Object*. This is how you'll build your PHP module later on. After the configuration finishes, the next two commands will build the binaries and then drop everything in the appropriate place according to the target of the `--prefix` flag.

```
cd apache_2.2.x
./configure --prefix=/usr/local/apache --enable-so
make
make install
```

3. Unzip and untar your PHP source distribution. Unless you have a reason to do otherwise, `/usr/local` is the standard place to do so.

```
tar -zxvf php-6.x.tar.gz
cd php-6.x
```

4. Configure your PHP build. (Configuring PHP is a topic so large and important that it would not fit into this chapter, so please flip over to Chapter 29 for more information.) The most common options are the ones to build as an Apache module, which you almost certainly want, and to do so with specific database support. The example build here is an Apache module with MySQL support, built using apxs.

```
./configure
--with-apxs2=/usr/local/apache/bin/apxs
--with-mysql
```

5. Make and install the PHP module.

```
make
make install
```

6. Install the `php.ini` file. Edit this file to get configuration directives; see the options listed in Chapter 29. At this point, we highly recommend that new users set error reporting to `E_ALL` on their development machines.

```
cd ../../php-6.x
cp php.ini-dist /usr/local/lib/php.ini
```

7. Tell your Apache server what extension(s) you want to identify PHP files (`.php` is the standard, but you can use `.html`, `.phtml`, or whatever you want). Go to your HTTP configuration files (`/usr/local/apache/conf` or whatever your path is), and open `httpd.conf` with a text editor. Add at least one PHP extension directive, as shown in the first line of code that follows. In the second line, we've also added a second handler to have all HTML files parsed as PHP (which does impose a small performance hit and should not be done if your architecture uses the `.html` file extension strictly for HTML-only files). This would also be a good time for you to ensure that Apache knows what domain alias or IP address to listen for. (If you have no idea what this means, search `httpd.conf` for the word `ServerName`, add the word `localhost` right after it, and use that as your domain name until you get a better one.)

```
AddType application/x-httpd-php .php
AddType application/x-httpd-php .html
```

8. Restart your server. Every time you change your HTTP configuration or `php.ini` files, you must stop and start your server again. An HUP signal will not suffice.

```
cd ../bin
./apachectl start
```

9. Set the document root directory permissions to world-executable. The actual PHP files in the directory need only be world-readable (644). If necessary, replace `/home/httpd` with your document root in the code that follows.

```
chmod 755 /home/httpd/html/php
```

10. Open a text editor. Type: `<?php phpinfo(); ?>`. Save this file in your web server's document root as `info.php`. Start any web browser and browse the file — you must always use an HTTP request (`http://www.example.com/info.php` or `http://localhost/info.php` or `http://127.0.0.1/info.php`) rather than a filename (`/home/httpd/info.php`) for the file to be parsed correctly. You should see a long table of information about your new PHP6 installation. Congratulations!

CROSS-REF   Many Apache production servers do not use a `php.ini` file; it can be undesirable to have two different configuration files in two different locations. You can replicate many of the configuration directives of `php.ini` in your Apache `httpd.conf` file. At a minimum, you probably want to set the include path and error-reporting levels, because the default settings for these are often unsatisfactory. See Chapter 29 for more details.

Microsoft Windows and Apache

As with the LAMP (Linux/Apache/MySQL/Perl/PHP/Python) stack, the last several years has seen a rise in the WAMP stack (Windows/Apache/MySQL/Perl/PHP/Python). If Microsoft Windows is your OS of choice, then you'll have no problem running any of these popular packages, just like your Linux brethren. Apache, PHP, and MySQL all offer installers and source code for Windows. This section examines installation on Microsoft Windows Server 2008, Windows Server 2003, and Windows Vista.

> **NOTE**   Microsoft Windows XP is still quite popular on the desktop, and installation of these components on Windows XP is roughly the same as the installation on Windows Server 2003.

To install Apache with PHP on Microsoft Windows Vista and Windows Server 2003 and 2008:

1. Download Apache server from `http://httpd.apache.org/download.cgi`. You want the current stable release version with the `no_src.msi` extension (You can try the `.exe` version if there is one, but it doesn't work on all systems and isn't any easier). Once downloaded, double-click the installer file to install. The installer will run through a wizard. For our intents and purposes in this book, you can accept the defaults. As you gain experience with the Apache server, you may find that you want to adjust and tweak the configuration, but for now, the defaults are fine.

   You may need to stop Internet Information Server (IIS) in Windows prior to starting Apache, since both will attempt to listen on TCP port 80. You may also need to allow Apache through the firewall in Windows. In Vista, this is accomplished through the Security Center Control Panel in Windows Vista. Specifically, by using the "Allow a program through Windows Firewall" option, clicking on Add Port, and then configuring TCP port 80 within the Add a Port dialog. In Windows Server 2008, the Windows Firewall with Advanced Security applet is found in Administrative Tools. Within the Windows Firewall with Advanced Security applet, clicking on Inbound Rules on the left and then New Rule on the right will result in a New Inbound Rule Wizard. Follow the wizard to add a TCP port of 80 inbound.

2. Next, download PHP from `www.php.net/downloads.php`. If there's an installer available, get it. Otherwise get the zip file version. If you download the installer, then you can merely follow through the Installation Wizard. Otherwise, for the zip version of PHP, extract the PHP binary archive using your unzip utility placing it in `C:\PHP`.

3. Copy some `.dll` files from your PHP directory to your system directory (usually `C:\Windows\System32`). You need `php6ts.dll` for every case. You will also probably need to copy the file corresponding to your web server module — `C:\PHP\php6apache2_2.dll` — to your Apache modules directory. It's possible that you will also need other files from the `dlls` subfolder — but start with the two mentioned previously and add more if you need them. For instance, it's quite common to need to copy `libmysql.dll` from `C:\PHP` to `C:\Windows\System32` as well, so you might as well copy it there now. In Windows Vista, I've found that the easiest way to do this is to right-click on the command prompt, select Run as Administrator, and then copy the files using the `copy` command, as in `copy c:\php\php6ts.dll c:\windows\system32\`.

4. Rename either `php.ini-dist` or `php.ini-recommended` (preferably the latter) as `php.ini` within your `C:\PHP` directory. Open this file in a text editor (for example, Notepad). Edit this file to get configuration directives; see the options listed in Chapter 29. At this point, we highly recommend that new users set error reporting to `E_ALL` on their development machines. Note that it's not strictly necessary to edit the file at this time, but you should be familiar with its contents nonetheless.

5. Go to your HTTP configuration files (`C:\Program Files\Apache Software Foundation\Apache2.2\conf` or whatever your path is), and open `httpd.conf` with a text editor. Add the PHP module load directive as shown in the first line of the following code and add the handler for `.php` and `.phtml` files, too:

   ```
   LoadModule php6_module modules/php6apache2_2.dll
   AddType application/x-httpd-php .php .phtml
   ```

6. Stop and restart the WWW service. Go to the Start menu ➪ All Programs ➪ Apache HTTP Server 2.2 ➪ Control Apache HTTP Server ➪ Stop/Start; or Restart, or even run Apache from the MS-DOS prompt.

7. Open a text editor (for example, Notepad). Type: `<?php phpinfo(); ?>`. Save this file in your web server's document root (`C:\Program Files\Apache Software Foundation\Apache2.2\htdocs` by default) as `info.php`. Start any web browser and request the file: `http://localhost/info.php` or `http://127.0.0.1/info.php`). You should see a long table of information about your new PHP6 installation. Congratulations! If things didn't go as planned, check the error log for Apache, usually located at `C:\Program Files\Apache Software Foundation\Apache\logs\error.log`.

CROSS-REF   **If you follow these directions and don't get the results you expected, don't panic! Check out Chapter 10 for common gotchas and quirks. If that doesn't help, check out the comments on the relevant pages in the PHP online manual — users leave specific tips for specific setups they've had problems with.**

Other web servers

PHP has been successfully built and run with many other web servers, such as Netscape Enterprise Server, Xitami, Zeus, and thttpd. Module support for AOLServer, NSAPI, and fhttpd is available. See the relevant pages on the PHP online manual's installation section.

Development tools

When it comes to development tools, PHP used to fall between the cracks — between tools originally designed for other programming languages and those mainly used to create pretty HTML. It's certainly possible to write a complex 2000-line program that touches several other services and file-systems and outputs the string 1 to the browser on completion. On the other hand, there are many people whose main use of PHP is to slap common headers and footers on what amounts to a bunch of static HTML pages. With such a diversity of usages, it's perhaps not so amazing that the perfect PHP development environment — user-friendly enough for the designers, but light and powerful enough for the geeks — has been elusive.

Those coming to PHP from a strictly client-side perspective probably have the hardest adjustment to make. There's no such thing as a plush development environment with wizards and drag-and-drop icons and built-in graphics manipulation. If that sort of thing is important to you, you can use a WYSIWYG editor to format the page and then add PHP functionality later using a text editor. The downside of this strategy is, of course, that machine-written code is often not very human-readable — but one must suffer to be pretty.

The last year and a half, however, has seen substantial change in the market. Plenty of editors for both Windows and Linux now offer at least syntax highlighting for PHP. Several of these can map drive locations to server names, so you can debug in place.

CAUTION   Be particularly careful with using Microsoft FrontPage or Adobe Dreamweaver as a PHP editor, as they both leave something to be desired for PHP development. .

Old-school programmers will have less of a learning curve, since they can treat PHP like any other server-side programming language that may or may not happen to output HTML to a browser. Most PHP users in this category seem to prefer simple text editors. Generally, these products will afford you a modest amount of help, such as syntax highlighting, brace matching, or tag closing — most of which is about helping you avoid stupid mistakes rather than actually writing the script for you.

My favorite is good old Vi, or Vi-Enhanced, Vim, although many people have problems using Vi. An excellent GUI tool is Eclipse. I've been using Eclipse for quite some time and feel comfortable recommending it for development in PHP, JavaScript, HTML, and just about any other language. Get Eclipse from `www.eclipse.org`.

What's to Come?

The remainder of this chapter looks at some basics of PHP, focusing on getting you up to speed for the rest of the book!

Your HTML Is Already PHP-Compliant!

PHP is already perfectly at home with HTML — in fact, it is generally embedded within HTML. As you'll see in later chapters, PHP rides piggyback on some of the cleverer parts of the HTML standard, such as forms and cookies, to do all kinds of useful things.

Anything compatible with HTML on the client side is also compatible with PHP. PHP could not care less about chunks of JavaScript, calls to music and animation, applets, or anything else on the client side. PHP will simply ignore those parts, and the web server will happily pass them on to the client.

It should be clear that you can use any method of developing web pages and simply add PHP to that method. If you're comfortable having teams work on each page using huge multimedia graphics suites, you can keep doing that. The general point is that you don't need to change tools or workflow order, just do what you've been doing and add the server-side functionality at the end.

Escaping from HTML

By now you're probably wondering: How does the PHP parser recognize PHP code inside your HTML document? The answer is that you tell the program when to spring into action by using special PHP tags at the beginning and end of each PHP section. This process is called *escaping from HTML* or *escaping into PHP*.

CAUTION   **Not to confuse you, but** *escape* **in this sense should not be confused with another common use of the term** *escape* **in PHP: putting a backslash in front of certain special characters (such as tab and newline) within double-quoted strings. Escaping strings is explained in Chapter 7.**

Everything within these tags is understood by the PHP parser to be PHP code. Everything outside of these tags does not concern the server and will simply be passed along and left for the client to sort out whether it's HTML or JavaScript or something else.

There are several styles of PHP, but it's best to stick with the tried-and-true tags that will always work no matter which version of PHP you're using:

Canonical PHP tags

The most universally effective PHP tag style is:

```
<?php  ?>
```

If you use this style, you can be positive that your tags will always be correctly interpreted. Unless you have a very, very strong reason to prefer another style, use this one. Some or all of the other styles of PHP tag may be phased out in the future — only this one is certain to be safe.

Hello World

Now you're ready to write your first PHP program. Open a new file in your preferred editor. Type:

```
<HTML>
<HEAD>
<TITLE>My first PHP program</TITLE>
</HEAD>

<BODY>
<?php
print("Hello, World<BR />\n");
phpinfo();
?>
</BODY>
</HTML>
```

In most browsers, nothing but the PHP section is strictly necessary; however, it's a good idea to get in the habit of always using a well-formed HTML structure in which to embed your PHP.

If you don't see something pretty close to the output shown in Figure 3-1, you have a problem — most likely some kind of installation or configuration glitch. Review Chapter 2 and make doubly sure that your installation succeeded.

FIGURE 3-1

Your first PHP script

Refer back to Chapter 2 for installation instructions and forward to Chapter 29 for configuration options. Chapter 10 diagnoses some common early problems and gives debugging hints.

Jumping in and out of PHP mode

At any given moment in a PHP script, you are either in PHP mode or you're out of it in HTML. There's no middle ground. Anything within the PHP tags is PHP; everything outside is plain HTML, as far as the server is concerned.

You can escape into PHP mode with giddy abandon, as often and as briefly or lengthily as necessary. For example:

```
<?php $id = 1; ?>
<FORM METHOD="POST" ACTION="registration.php"">
<P>First name:
<INPUT TYPE="TEXT" NAME="firstname" SIZE="20">
<P>Last name:
<INPUT TYPE="TEXT" NAME="lastname" SIZE="20">
<P>Rank:
<INPUT TYPE="TEXT" NAME="rank" SIZE="10">
<INPUT TYPE="HIDDEN" NAME="serial number" VALUE="<?php
echo $id; ?>">
<INPUT TYPE="submit"SUBMIT" VALUE="INPUT"">
</FORM>
```

Notice that things that happened in the first PHP mode instance — in this case, a variable being assigned — are still valid in the second. In Chapter 4, you'll learn more about what happens to variables when you skip in and out of PHP mode. In Chapter 32, you'll also learn about different styles of using PHP mode.

Including files

Another way you can add PHP to your HTML is by putting it in a separate file and calling it by using PHP's `include` functions. There are four `include` functions:

- `include('/filepath/filename')`
- `require('/filepath/filename')`
- `include_once('/filepath/filename')`
- `require_once('/filepath/filename')`

In previous versions of PHP, there were significant differences in functionality and speed between the `include` functions and the `require` functions. This is no longer true; the two sets of functions differ only in the kind of error they throw on failure. `Include()` and `include_once()` will merely generate a warning on failure, while `require()` and `require_once()` will cause a fatal error and termination of the script.

As suggested by the names of the functions, `include_once()` and `require_once()` differ from simple `include()` and `require()` in that they will allow a file to be included only once per PHP script. This is extremely helpful when you are including files that contain PHP functions, because

redeclaring functions results in an automatic fatal error. In larger PHP systems, it's quite common to include files that include other files that include other files — it can be difficult to remember whether you've included a particular function before, but with `include_once()` or `require_once()` you don't have to.

How do you decide on a preferred `include` function? In essence, you must decide whether you want to force yourself to write good code on pain of fatal error or whether you want it to run regardless of certain common errors on your part. The strictest alternative is `require()`, which will bring everything grinding to a halt if your code isn't perfect; the least strict is `include_once()`, which will good-naturedly hide the consequences of some of your bad coding habits.

The most common use of PHP's `include` capability is to add common headers and footers to all the web pages on a site. For example, a simple header file (cleverly named `header.php`) might look like this:

```
<HTML>
<HEAD>
<TITLE>A site title</TITLE>
</HEAD>

<BODY>
```

Similarly, a footer file called `footer.php` might consist of:

```
<P>Copyright 1995 - 2002</P>
</BODY>
</HTML>
```

They are called from a PHP page this way:

```
<?php
require_once($_SERVER['DOCUMENT_ROOT'].'/header.php');
?>
<P>This is some body text for this particular page.</P>
<?php
require_once($_SERVER['DOCUMENT_ROOT'].'/footer.php');
?>
```

Obviously, this single move greatly enhances the maintainability and scalability of an entire site. Now, if you want a different look and feel or if you need to update the copyright notice, you can alter one file instead of identical lines in dozens of HTML pages.

> **TIP** When including files, remember to set the `include_path` directive correctly in your `php.ini` file. Remember that you can include files from above or entirely outside your web tree by proper use of this directive. See Chapter 29 for more information.

As you can see from the preceding example, PHP's `include` functions simply pass along the contents of the included file *as text*. Many people think that because an `include` function occurs inside PHP mode, the included file will also be in PHP mode. This is not true! Actually, the server escapes

back into HTML mode at the beginning of each included file and silently returns to PHP mode at the end, just in time to catch the semicolon.

As always, you need to say when you intend something to be PHP by using PHP opening and closing tags. Any part of an included file that needs to be executed as PHP should be enclosed in valid PHP tags. If the entire file is PHP (very common in files of functions), the entire file must be enclosed within PHP tags.

Take the following file, `database.php`:

```
$db = mysql_connect('localhost', 'db_user', 'db_password');
mysql_select_db('my_database');
```

CAUTION We can't emphasize this enough: If you're having problems including PHP files, particularly if you're seeing output you don't expect or not seeing output you do expect, be *ABSOLUTELY POSITIVE* **that you've put PHP tags at the beginning and end of the included file.**

If you were to foolishly include this file from a PHP script, your database variables would be visible to the world in plain text — because you neglected to use PHP tags, the parser assumes that this block of code is HTML. A correct version of the `database.php` file would look like this:

```
<?php
$db = mysql_connect('localhost', 'db_user', 'db_password');
mysql_select_db('my_database');
?>
```

CAUTION **For all PHP files included from other files, you must ensure that there are no empty new lines at the end of the file. Remember, anything outside a PHP block is considered HTML, even a blank line. Blank lines, or even blank spaces outside a closing PHP tag, will be interpreted as output. If you include the file in a situation where you cannot have output — say before using HTTP headers — your script will fail with a big error message about the output stream having already been started in your included file. See Chapter 10 for an example.**

Summary

This chapter gets you up to speed with PHP, beginning with installation instructions for the several common platforms. Finally, some coding was shown in this chapter through the venerable "Hello World" example, illustrating not only that your PHP installation is working, but also that you can code in PHP!

Chapter 4

Learning PHP Syntax and Variables

I n this chapter, we cover the basic syntax of PHP — the rules that all well-formed PHP code must follow. We explain how to use variables to store and retrieve information as your PHP code executes and the type of system that governs what kinds of values can be stored in the first place. Finally, we look at the simplest ways to display text that will show up in your user's browser window.

PHP Is Forgiving

The first and most important thing to say about the PHP language is that it tries to be as forgiving as possible. Programming languages vary quite a bit in terms of how stringently syntax is enforced. Pickiness can be a *good* thing because it helps make sure that the code you're writing is really what you mean. If you are writing a program to control a nuclear reactor and you forget to assign a variable, it is far better to have the program be rejected than to create behavior different from what you intended. PHP's design philosophy, however, is at the other end of the spectrum. Because PHP started life as a handy utility for making quick web pages, it emphasizes convenience for the programmer over correctness; rather than have a programmer do the extra work of redundantly specifying what is meant by a piece of code, PHP requires the minimum and then tries its best to figure out what was meant. Among other things, this means that certain syntactical features that show up in other languages, such as variable declarations and function prototypes, are simply not necessary.

With that said, though, PHP can't read your mind; it has a minimum set of syntactical rules that your code must follow. Whenever you see the words `parse error` in your browser window instead of the cool web page you thought you had just written, it means that you've broken these rules to the point that PHP has given up on your page.

HTML Is Not PHP

The second most important thing to understand about PHP syntax is that it applies only within PHP. Because PHP is embedded in HTML documents, every part of such a document is interpreted as either PHP or HTML, depending on whether that section of the document is enclosed in PHP tags.

PHP syntax is relevant only within PHP, so we assume for the rest of this chapter that PHP mode is in force — that is, most code fragments will be assumed to be embedded in an HTML page and surrounded with the appropriate tags.

PHP's Syntax Is C-Like

The third most important thing to know about PHP syntax is that, broadly speaking, it is like the C programming language. If you happen to be one of the lucky people who already know C, this is very helpful; if you are uncertain about how a statement should be written, try it first the way you would do it in C, and if that doesn't work, look it up in the manual. The rest of this section is for the other people, the ones who don't already know C. (C programmers might want to skim the headers of this section and also see Appendix A, which is specifically for C programmers.)

PHP is whitespace insensitive

Whitespace is the stuff you type that is typically invisible on the screen, including spaces, tabs, and carriage returns (end-of-line characters). PHP's whitespace insensitivity does not mean that spaces and such never matter. (In fact, they are crucial for separating the *words* in the PHP language.) Instead, it means that it almost never matters *how many* whitespace characters you have in a row — one whitespace character is the same as many such characters.

For example, each of the following PHP statements that assigns the sum of 2 + 2 to the variable `$four` is equivalent:

```
$four = 2 + 2;      // single spaces
$four <tab>=<tab>2<tab>+<tab>2  ;     // spaces and tabs
$four              =
2
+
2;  // multiple lines
```

The fact that end-of-line characters count as whitespace is handy, because it means you never have to strain to make sure that a statement fits on a single line.

PHP is sometimes case sensitive

Having read that PHP isn't picky, you may be surprised to learn that it is sometimes case sensitive (that is, it cares about the distinction between lowercase and capital letters). In particular, all variables are case sensitive. If you embed the following code in an HTML page:

```php
<?php
  $capital = 67;
  print("Variable capital is $capital<BR>");
  print("Variable CaPiTaL is $CaPiTaL<BR>");
?>
```

The output you will see is:

```
Variable capital is 67
Variable CaPiTaL is
```

The different capitalization schemes make for different variables. (Surprisingly, under the default settings for error reporting, code like this fragment will not produce a PHP error — see the section "Unassigned variables," later in this chapter.)

On the other hand, unlike in C, function names are *not* case sensitive, and neither are the basic language constructs (`if`, `then`, `else`, `while`, and the like).

Statements are expressions terminated by semicolons

A *statement* in PHP is any *expression* that is followed by a semicolon (`;`). If expressions correspond to phrases, statements correspond to entire sentences, and the semicolon is the full stop at the end. Any sequence of valid PHP statements that is enclosed by the PHP tags is a valid PHP program. Here is a typical statement in PHP, which in this case assigns a string of characters to a variable called `$greeting`:

```php
$greeting = "Welcome to PHP!";
```

The rest of this subsection is about how such statements are built from smaller components and how the PHP interpreter handles the evaluation of statements. (If you already feel comfortable with statements and expressions, feel free to skip ahead.)

Expressions are combinations of tokens

The smallest building blocks of PHP are the *indivisible tokens*, such as numbers (`3.14159`), strings (`"two"`), variables (`$two`), constants (`TRUE`), and the special words that make up the syntax of PHP

itself (`if`, `else`, and so forth). These are separated from each other by whitespace and by other special characters such as parentheses and braces.

The next most complex building block in PHP is the *expression*, which is any combination of tokens that has a value. A single number is an expression, as is a single variable. Simple expressions can also be combined to make more complicated expressions, usually either by putting an *operator* in between (for example, 2 + (2 + 2)) or by using them as input to a function call (for example, `pow(2 * 3, 3 * 2)`). Operators that take two inputs go in between their inputs, whereas functions take their inputs in parentheses immediately after their names, with the inputs (known as *arguments*) separated by commas.

Expressions are evaluated

Whenever the PHP interpreter encounters an expression in code, that expression is immediately *evaluated*. This means that PHP calculates values for the smallest elements of the expression and successively combines those values connected by operators or functions, until it has produced an entire value for the expression. For example, successive steps in an imaginary evaluation process might look like:

```
$result = 2 * 2 + 3 * 3 + 5;
        (= 4 + 3 * 3 + 5)  //imaginary evaluation steps
        (= 4 + 9 + 5)
        (= 13 + 5)
        (= 18)
```

with the result that the number 18 is stored in the variable `$result`.

Precedence, associativity, and evaluation order

There are two kinds of freedom PHP has in expression evaluation: how it groups or associates subexpressions and the order in which it evaluates them. For example, in the evaluation process just shown, multiplications were associated more tightly than additions, which affects the end result.

The particular ways that operators group expressions are called *precedence rules* — operators that have higher precedence win in grabbing the expressions around them. If you want, you can memorize the rules, such as the fact that `*` always has higher precedence than +. Or you can just use the following cardinal rule: When in doubt, use parentheses to group expressions.

For example:

```
$result1 = 2 + 3 * 4 + 5; // is equal to 19
$result2 = (2 + 3) * (4 + 5); // is equal to 45
```

Operator precedence rules remove much of the ambiguity about how subexpressions are associated. But what about when two operators have the same precedence? Consider this expression:

```
$how_much = 3.0 / 4.0 / 5.0;
```

Whether this is equal to 0.15 or 3.75 depends on which division operator gets to grab the number 4.0 first. There is an exhaustive list of rules of associativity in the online manual, but the rule to remember is that associativity is usually left-before-right — that is, the preceding expression would evaluate to 0.15, because the leftmost of the two division operators wins the dispute over precedence.

The final wrinkle is order of evaluation, which is not quite the same thing as associativity. For example, look at the arithmetic expression:

```
3 * 4 + 5 * 6
```

We know that the multiplications will happen before the additions, but that is not the same as knowing which multiplication PHP will perform first. In general, you need not worry about evaluation order, because in almost all cases it will not affect the result. You can construct weird examples where the result does depend on order of evaluation, usually by making assignments in subexpressions that are used in other parts of the expression. For example:

```
$huh = ($this = $that + 5) + ($that = $this + 3);  // BAD
```

But don't do this, okay? PHP may or may not have a predictable order of evaluation of expressions, but you shouldn't depend on it — so we're not going to tell you! (The one legitimate use of relying on left-to-right evaluation order is in short-circuiting Boolean expressions, which we cover in Chapter 5.)

Expressions and types

Usually, the programmer is careful to match the types of expressions with the operators and functions that combine them. Common expressions are *mathematical* (with mathematical operators combining numbers) or *Boolean* (combining true-or-false statements with *and*s and *or*s) or *string expressions* (with operators and functions constructing strings of characters). As with the rest of PHP, however, the treatment of types is surprisingly forgiving. Consider the following expression, which deliberately mixes the types of subexpressions in an inappropriate way:

```
2 + 2 * "nonsense" + TRUE
```

Rather than produce an error, this evaluates to the number 3. (You can take this as a puzzle for now, but we will explain how such a thing can happen in the "Types in PHP" section of this chapter.)

Assignment expressions

A very common kind of expression is the *assignment*, where a variable is set to equal the result of evaluating some expression. These have the form of a variable name (which always starts with a $), followed by a single equal sign, followed by the expression to be evaluated. For example:

```
$eight = 2 * (2 * 2);
```

assigns the variable $eight the value you would expect.

An important thing to remember is that even assignment expressions are expressions and so have values themselves! The value of an expression that assigns a variable is the same as the value assigned. This means that you can use assignment expressions in the middle of more complicated expressions. If you evaluate the statement:

```
$ten = ($two = 2) + ($eight = 2 * (2 * 2));
```

each variable would be assigned a numerical value equal to its name.

Reasons for expressions and statements

There are usually only two reasons to write an expression in PHP: for its *value* or for a *side effect*. The value of an expression is passed on to any more complicated expression that includes it; side effects are anything else that happens as a result of the evaluation. The most typical side effects involve assigning or changing a variable, printing something to the user's screen, or making some other persistent change to the program's environment (such as interacting with a database).

Although statements are expressions, they are not themselves included in more complicated expressions. This means that the only good reason for a statement is a side effect! It also means that it is possible to write legal (yet totally useless statements) such as the second of these:

```
print("Hello");  // side effect is printing to screen

2 * 3 + 4;  // useless - no side effect

$value_num = 3 * 4 + 5;  // side effect is assignment

store_in_database(49.5);  // side effect to DB
```

Braces make blocks

Although statements cannot be combined like expressions, you can always put a sequence of statements anywhere a statement can go by enclosing them in a set of curly braces.

For example, the `if` construct in PHP has a test (in parentheses) followed by the statement that should be executed if the test is true. If you want more than one statement to be executed when the test is true, you can use a brace-enclosed sequence instead. The following pieces of code (which simply print a reassuring statement that it is still true that 1 + 2 is equal to 3) are equivalent:

```
if (3 == 2 + 1)
  print("Good - I haven't totally lost my mind.<BR>");

if (3 == 2 + 1)
  {
    print("Good - I haven't totally ");
    print("lost my mind.<BR>");
  }
```

You can put any kind of statement in a brace-enclosed block, including, say, an `if` statement that itself has a brace-enclosed block. This means that `if` statements can have other `if` statements inside them. In fact, this kind of nesting can be done to an arbitrary number of levels.

Comments

A *comment* is the portion of a program that exists only for the human reader. The very first thing that a program executor does with program code is to strip out the comments, so they cannot have any effect on what the program does. Comments are invaluable in helping the next person who reads your code figure out what you were thinking when you wrote it, even when that person is yourself a week from now.

PHP drew its inspiration from several different programming languages, most notably C, Perl, and Unix shell scripts. As a result, PHP supports styles of comments from all those languages, and those styles can be intermixed freely in PHP code.

C-style multiline comments

The *multiline* style of commenting is the same as in C: A comment starts with the character pair `/*` and terminates with the character pair `*/`. For example:

```
/*  This is
    a comment in
    PHP */
```

The most important thing to remember about multiline comments is that they cannot be nested. You cannot put one comment inside another. If you try, the comment will be closed off by the first instance of the `*/` character pair, and the rest of what was intended to be an enclosing comment will instead be interpreted as code, probably failing horribly. For example:

```
/* This comment will /* fail horribly on the
    last word of this */ sentence
*/
```

This is an easy thing to do unintentionally, usually when you try to deactivate a block of commented code by "commenting it out."

Single-line comments: # and //

In addition to the `/* ... */` multiple-line comments, PHP supports two different ways of commenting to the end of a given line: one inherited from C++ and Java and the other from Perl and shell scripts. The shell-script-style comment starts with a pound sign, whereas the C++ style comment starts with two forward slashes. Both of them cause the rest of the current line to be treated as a comment, as in the following:

```
#  This is a comment, and
 #  this is the second line of the comment
```

39

```
 // This is a comment too.  Each style comments only
 // one line so the last word of this sentence will fail
horribly.
```

The very alert reader might argue that single-line comments are incompatible with what we said earlier about whitespace insensitivity. That would be correct — you will get a very different result if you take a single-line comment and replace one of the spaces with an end-of-line character. A more accurate way of putting it is that, after the comments have been stripped out of the code, PHP code is whitespace insensitive.

Variables

The main way to store information in the middle of a PHP program is by using a variable — a way to name and hang on to any value that you want to use later.

Here are the most important things to know about variables in PHP (more detailed explanations will follow):

■ All variables in PHP are denoted with a leading dollar sign ($).

■ The value of a variable is the value of its most recent assignment.

■ Variables are assigned with the = operator, with the variable on the left-hand side and the expression to be evaluated on the right.

■ Variables can, but do not need, to be declared before assignment.

■ Variables have no intrinsic type other than the type of their current value.

■ Variables used before they are assigned have default values.

PHP variables are Perl-like

All variables in PHP start with a leading $ sign just like scalar variables in the Perl scripting language, and in other ways they have similar behavior (need no type declarations, may be referred to before they are assigned, and so on). (Perl hackers may need to do no more than skim the headings of this section, which is really for the rest of us.)

After the initial $, variable names must be composed of letters (uppercase or lowercase), digits (0–9), and underscore characters (_). Furthermore, the first character after the $ may not be a number.

Declaring variables (or not)

This subheading is here simply because programmers from some other languages might be looking for it — in languages such as C, C++, and Java, the programmer must declare the name and type of any variable before making use of it. However in PHP, because types are associated with values rather than variables, no such declaration is necessary — the first step in using a variable is to assign it a value.

Assigning variables

Variable assignment is simple — just write the variable name, and add a single equal sign (=); then add the expression that you want to assign to that variable:

```
$pi = 3 + 0.14159; // approximately
```

Note that what is assigned is the result of evaluating the expression, not the expression itself. After the preceding statement is evaluated, there is no way to tell that the value of `$pi` was created by adding two numbers together.

It's conceivable that you will want to actually print the preceding math expression rather than evaluate it. You can force PHP to treat a mathematical variable assignment as a string by quoting the expression:

```
$pi = "3 + 0.14159";
```

Reassigning variables

There is no interesting distinction in PHP between assigning a variable for the first time and changing its value later. This is true even if the assigned values are of different types. For example, the following is perfectly legal:

```
$my_num_var = "This should be a number - hope it's reassigned";
$my_num_var = 5;
```

If the second statement immediately follows the first one, the first statement has essentially no effect.

Unassigned variables

Many programming languages will object if you try to use a variable before it is assigned; others will let you use it, but if you do you may find yourself reading the random contents of some area of memory. In PHP, the default error-reporting setting allows you to use unassigned variables without errors, and PHP ensures that they have reasonable default values.

CROSS-REF   **If you would like to be warned about variables that have not been assigned, you should change the error-reporting level to** `E_ALL` **(the highest level possible) from the default level of error reporting. You can do this either by including the statement** `error_reporting(E_ALL);` **at the top of a script or by changing your** `php.ini` **file to set the default level (see Chapters 29 and 30).**

Default values

Variables in PHP do not have intrinsic types — a variable does not know in advance whether it will be used to store a number or a string of characters. So how does it know what type of default value to have when it hasn't yet been assigned?

The answer is that, just as with assigned variables, the type of a variable is interpreted depending on the context in which it is used. In a situation where a number is expected, a number will be produced, and this works similarly with character strings. In any context that treats a variable as a number, an unassigned variable will be evaluated as 0; in any context that expects a string value, an unassigned variable will be the empty string (the string that is zero characters long).

Checking assignment with isset

Because variables do not have to be assigned before use, in some situations you can actually convey information by selectively setting or not setting a variable! PHP provides a function called isset that tests a variable to see whether it has been assigned a value.

As the following code illustrates, an unassigned variable is distinguishable even from a variable that has been given the default value:

```
$set_var = 0; //set_var has a value
              //never_set does not
print("set_var print value: $set_var<BR>");
print("never_set print value: $never_set<BR>");
if ($set_var == $never_set)
  print("set_var is equal to never_set!<BR>");
if (isset($set_var))
  print("set_var is set.<BR>");
else
  print("set_var is not set.<BR>");
if (isset($never_set))
  print("never_set is set.<BR>");
else
  print("never_set is not set.");
```

Oddly enough, this code will produce the following output:

```
set_var print value: 0
never_set print value:
set_var is equal to never_set!
set_var is set.
never_set is not set.
```

The variable $never_set has never been assigned, so it produces an empty string when a string is expected (as in the print statement) and a zero value when a number is expected (as in the comparison test that concludes that the two variables are the same). Still, isset can tell the difference between $set_var and $never_set.

Assigning a variable is not irrevocable — the function unset() will restore a variable to an unassigned state (for example, unset($set_var); will make $set_var into an unbound variable, regardless of its previous assignments).

Variable scope

Scope is the technical term for the rules about when a name (for, say, a variable or function) has the same meaning in two different places and in what situations two names spelled exactly the same way can actually refer to different things.

Any PHP variable not inside a function has *global* scope and extends throughout a given "thread" of execution. In other words, if you assign a variable near the top of a PHP file, the variable name has the same meaning for the rest of the file; and if it is not reassigned, it will have the same value as the rest of your code executes (except inside the body of functions and classes).

The assignment of a variable will not affect the value of variables with the same name in other PHP files or even in repeated uses of the same file. For example, let's say that you have two files, `startup.php` and `next_thing.php`, which are typically visited in that order by a user. Let's also say that near the top of `startup.php`, you have the line:

```
$username = "Jane Q. User";
```

which is executed only in certain situations. Now, you might hope that, after setting that variable in `startup.php`, it would also be preset automatically when the user visited `next_thing.php`, but no such luck. Each time a PHP page executes, it assigns and reassigns variables as it goes, and those variables disappear at the end of a page's production. Assignments of variables in one file do not affect variables of the same name in a different file or even in other requests for the same file.

Obviously, there are many situations in which you would like to hold onto information for longer than it takes to generate a particular web page. There are a variety of ways you can accomplish this, and the different techniques are a lot of what the rest of this book is about. For example, you can pass information from page to page using `GET` and `POST` variables (Chapter 6), store information persistently in a database (all of Part II of this book), associate it with a user's session using PHP's session mechanism (see Chapter 24), or store it on a user's hard disk via a cookie (see Chapter 24).

Functions and variable scope

Except inside the body of a function, variable scope in PHP is quite simple: Within any given execution of a PHP file, just assign a variable, and its value will be there for you later. We haven't yet covered how to define your own functions, but it's worth a look-ahead note: Variables assigned within a function are *local* to that function, and unless you make a special declaration in a function, that function won't have access to the global variables defined outside the function, even when they are defined in the same file. (We will discuss the scope of variables in functions in depth when we cover function definitions in Chapter 5.)

You can switch modes if you want

One scoping question that we had the first time we saw PHP code was: Does variable scope persist across tags? For example, we have a single file that looks like:

```
<HTML>
<HEAD>
```

```
<?php
  $username = "Jane Q. User";
?>
</HEAD>
<BODY>
<?php
  print("$username<BR>");
?>
</BODY>
</HTML>
```

Should we expect our assignment to $username to survive through the second of the two PHP-tagged areas? The answer is yes — variables persist throughout a thread of PHP execution (in other words, through the whole process of producing a web page in response to a user's request). This is a single manifestation of a general PHP rule, which is that the *only* effect of the tags is to let the PHP engine know whether you want your code to be interpreted as PHP or passed through untouched as HTML. You should feel free to use the tags to switch back and forth between modes whenever it is convenient.

Constants

In addition to variables, which may be reassigned, PHP offers constants, which have a single value throughout their lifetime. Constants do not have a $ before their names, and by convention the names of constants usually are in uppercase letters. Constants can contain only scalar values (numbers and string). Constants have global scope, so they are accessible everywhere in your scripts after they have been defined — even inside functions.

For example, the built-in PHP constant E_ALL represents a number that indicates to the error_reporting() function that all errors and warnings should be reported. A call to error_reporting() might look like this:

```
error_reporting(E_ALL);
```

This is identical to calling error_reporting() on the integer value of E_ALL, but is better because the actual value of E_ALL may change from one version of PHP to the next.

It's also possible to create your own constants using the define() function. The code:

```
define(MY_ANSWER, 42);
```

would cause MY_ANSWER to evaluate to 42 everywhere it appears in your code. There is no way to change this assignment after it has been made, and like variables, user-defined constants that are not part of PHP itself do not persist across pages unless they are explicitly passed to a new page. When created constants are used, they are generally most usefully defined in an external include file and might be used for such information as a sales-tax rate or perhaps an exchange rate.

Types in PHP: Don't Worry, Be Happy

All programming languages have some kind of type system, which specifies the different kinds of values that can appear in programs. These different types often correspond to different bit-level representations in computer memory, although in many cases programmers are insulated from having to think about (or being able to mess with) representations in terms of bits.

PHP's type system is simple, streamlined, and flexible, and it insulates the programmer from low-level details. PHP makes it easy not to worry too much about typing of variables and values, both because it does not require variables to be typed and because it handles a lot of type conversions for you.

No variable type declarations

As you saw in Chapter 3, the type of a variable does not need to be declared in advance. Instead, the programmer can jump right ahead to assignment and let PHP take care of figuring out the type of the expression assigned:

```
$first_number = 55.5;
$second_number = "Not a number at all";
```

Automatic type conversion

PHP does a good job of automatically converting types when necessary. Like most other modern programming languages, PHP will do the right thing when, for example, doing math with mixed numerical types. The result of the expression

```
$pi = 3 + 0.14159;
```

is a floating-point (double) number, with the integer 3 implicitly converted into floating point before the addition is performed.

Types assigned by context

PHP goes further than most languages in performing automatic type conversions. Consider:

```
$sub = substr(12345, 2, 2);
print("sub is $sub<BR>");
```

The `substr` function is designed to take a string of characters as its first input and return a substring of that string, with the start point and length determined by the next two inputs to the function. Instead of handing the function a character string, however, we gave it the integer `12345`. What happens? As it turns out, there is no error, and we get the browser output:

```
sub is 34
```

Because `substr` expects a character string rather than an integer, PHP converts the number `12345` to the character string `'12345'`, which `substr` then slices and dices.

Because of this automatic type conversion, it is very difficult to persuade PHP to give a type error — in fact, PHP programmers need to exercise a little care sometimes to make sure that type confusions do not lead to error-free but unintended results.

Type Summary

PHP has a total of eight types: integers, doubles, Booleans, strings, arrays, objects, NULL, and resources.

- *Integers* are whole numbers, without a decimal point, like 495.
- *Doubles* are floating-point numbers, like 3.14159 or 49.0.
- *Booleans* have only two possible values: `TRUE` and `FALSE`.
- *NULL* is a special type that only has one value: `NULL`.
- *Strings* are sequences of characters, like `'PHP 4.0 supports string operations.'`
- *Arrays* are named and indexed collections of other values.
- *Objects* are instances of programmer-defined classes, which can package up both other kinds of values and functions that are specific to the class.
- *Resources* are special variables that hold references to resources external to PHP (such as database connections).

Of these, the first five are *simple types*, and the next two (arrays and objects) are *compound* — the compound types can package up other arbitrary values of arbitrary type, whereas the simple types cannot. We treat only the simple types in this chapter, since arrays (see Chapter 8) and objects (see Chapter 20) need chapters all to themselves. Finally, the thorniest details of the type system, including discussion of the resource type, are deferred to Chapter 25.

The Simple Types

The most of the simple types in PHP (integers, doubles, Booleans, NULL, and strings) should be familiar to those with programming experience (although we will not assume that experience and will explain them in detail). The only thing likely to surprise C programmers is how few types there are in PHP.

Many programming languages have several different sizes of numerical types, with the larger ones allowing a greater range of values, but also taking up more room in memory. For example, the C language has a `short` type (for relatively small integers), a `long` type (for possibly larger integers), and an `int` type (which might be intermediate, but in practice is sometimes identical either to the `short` or `long` type). It also has floating-point types, which vary in their precision. This kind of typing choice made sense in an era when tradeoffs between memory use and functionality were often agonizing. The PHP designers made what we think is a good decision to simplify this by having only two numerical types, corresponding to the largest of the integral and floating-point types in C.

Integers

Integers are the simplest type — they correspond to simple whole numbers, both positive and negative. Integers can be assigned to variables, or they can be used in expressions, like this:

```
$int_var = 12345;
$another_int = -12345 + 12345; // will equal zero
```

Read formats

Integers can actually be read in three formats, which correspond to bases: *decimal* (base 10), *octal* (base 8), and *hexadecimal* (base 16). Decimal format is the default, octal integers are specified with a leading 0, and hexadecimals have a leading 0x. Any of the formats can be preceded by a - sign to make the integer negative. For example:

```
$integer_10 = 1000;
$integer_8 = -01000;
$integer_16 = 0x1000;
print("integer_10: $integer_10<BR>");
print("integer_8: $integer_8<BR>");
print("integer_16: $integer_16<BR>");
```

yields the browser output:

```
integer_10: 1000
integer_8: -512
integer_16: 4096
```

Note that the read format affects only how the integer is converted as it is read — the value stored in $integer_8 does not remember that it was originally written in base 8. Internally, of course, these numbers are represented in binary format; we see them in their base 10 conversion in the preceding output because that is the default for printing and incorporating int variables into strings.

Range

How big (or small) can integers get? Because PHP integers correspond to the C long type, which in turn depends on the word-size of your machine, this is difficult to answer definitively. For most common platforms, however, the largest integer is $2^{31} - 1$ (or 2,147,483,647), and the smallest (most negative) integer is $-(2^{31} - 1)$ (or –2,147,483,647).

The PHP constant PHP_INT_MAX will tell you the maximum integer for your implementation. If you really need integers even larger or smaller than the preceding, PHP does have some arbitrary-precision functions — see the BC section of the "Mathematics" chapter (see Chapter 27).

Doubles

Doubles are floating-point numbers, such as:

```
$first_double = 123.456;
```

```
$second_double = 0.456;
$even_double = 2.0;
```

Note that the fact that `$even_double` is a "round" number does not make it an integer. Integers and doubles are stored in different underlying formats, and the result of:

```
$five = $even_double + 3;
```

is a double, not an integer, even if it prints as 5. In almost all situations, however, you should feel free to mix doubles and integers in mathematical expressions, and let PHP sort out the typing.

By default, doubles print with the minimum number of decimal places needed — for example, the code:

```
$many = 2.2888800;
$many_2 = 2.2111200;
$few = $many + $many_2;
print("$many + $many_2 = $few<BR>");
```

produces the browser output:

```
2.28888 + 2.21112 = 4.5
```

CROSS-REF If you need finer control of printing, see the `printf` function in Chapter 7.

Read formats

The typical *read format* for doubles is `-X.Y`, where the `-` optionally specifies a negative number, and both X and Y are sequences of digits between 0 and 9. The X part may be omitted if the number is between –1.0 and 1.0, and the Y part can also be omitted. Leading or trailing zeros have no effect. All the following are legal doubles:

```
$small_positive = 0.12345;
$small_negative = -.12345;
$even_double = 2.00000;
$still_double = 2.;
```

In addition, doubles can be specified in scientific notation, by adding the letter e and a desired integral power of 10 to the end of the previous format — for example, `2.2e-3` would correspond to 2.2×10^{-3}. The floating-point part of the number need not be restricted to a range between 1.0 and 10.0. All the following are legal:

```
$small_positive = 5.5e-3;
print("small_positive is $small_positive<BR>");
$large_positive = 2.8e+16;
print("large_positive is $large_positive<BR>");
$small_negative = -2222e-10;
print("small_negative is $small_negative<BR>");
$large_negative = -0.00189e6;
print("large_negative is $large_negative<BR>");
```

The preceding code produces the following browser output:

```
small_positive is 0.0055
large_positive is 2.8E+16
small_negative is -2.222E-07
large_negative is -1890
```

Notice that, just as with octal and hexadecimal integers, the read format is irrelevant once PHP has finished reading in the numbers — the preceding variables retain no memory of whether or not they were originally specified in scientific notation. In printing the values, PHP is making its own decisions to print the more extreme values in scientific notation, but this has nothing to do with the original read format.

Booleans

Booleans are true-or-false values, which are used in control constructs like the testing portion of an `if` statement. As you will see in Chapter 5, Boolean truth values can be combined using logical operators to make more complicated Boolean expressions.

Boolean constants

PHP provides a couple of constants especially for use as Booleans: `TRUE` and `FALSE`, which can be used like this:

```
if (TRUE)
 print("This will always print<BR>");
else
 print("This will never print<BR>");
```

Interpreting other types as Booleans

Here are the rules for determine the "truth" of any value not already of the Boolean type:

- If the value is a number, it is false if the number is zero and true otherwise.

- If the value is a string, it is false if the string is empty (has zero characters) *or* is the string `"0"`, and is true otherwise.

- Values of type NULL are always false.

- If the value is a compound type (an array or an object), it is false if it contains no other values, and it is true otherwise. For an object, *containing a value* means having a member variable that has been assigned a value.

- Valid resources are true (although some functions that return resources when they are successful will return `FALSE` when unsuccessful).

CROSS-REF   **For a more complete account of converting values across types, see Chapter 25.**

Examples

Each of the following variables has the truth value embedded in its name when it is used in a Boolean context.

```
$true_num = 3 + 0.14159;
$true_str = "Tried and true";
$true_array[49] = "An array element"; // see next section
$false_array = array();
$false_null = NULL;
$false_num = 999 - 999;
$false_str = ""; // a string zero characters long
```

Don't use doubles as Booleans

Note that, although Rule 1 implies that the double 0.0 converts to a false Boolean value, it is dangerous to use floating-point expressions as Boolean expressions, because of possible rounding errors. For example:

```
$floatbool = sqrt(2.0) * sqrt(2.0) - 2.0;
if ($floatbool)
  print("Floating-point Booleans are dangerous!<BR>");
else
  print("It worked ... this time.<BR>");
print("The actual value is $floatbool<BR>");
```

The variable $floatbool is set to the result of subtracting two from the square of the square root of two — the result of this calculation should be equal to zero, which means that $floatbool is false. Instead, the browser output we get is:

```
Floating-point Booleans are dangerous!
The actual value is 4.4408920985006E-16
```

The value of $floatbool is very close to 0.0, but it is nonzero and, therefore, unexpectedly true. Integers are much safer in a Boolean role — as long as their arithmetic happens only with other integers and stays within integral sizes, they should not be subject to rounding errors.

NULL

The world of Booleans may seem small, since the Boolean type has only two possible values. The *NULL type*, however, takes this to the logical extreme: The type NULL has only one possible value, which is the value NULL. To give a variable the NULL value, simply assign it like this:

```
$my_var = NULL;
```

The special constant NULL is capitalized by convention, but actually it is case insensitive; you could just as well have typed:

```
$my_var = null;
```

So what is special about NULL? NULL represents the *lack* of a value. (You can think of it as the *non-value* or the *unvalue*.) A variable that has been assigned the value NULL is nearly indistinguishable from a variable that has not been set at all. In particular, a variable that has been assigned NULL has the following properties:

- It evaluates to FALSE in a Boolean context.
- It returns FALSE when tested with IsSet(). (No other type has this property.)
- PHP will not print warnings if you pass the variable to functions and back again, whereas passing a variable that has never been set will sometimes produce warnings.

The NULL value is best used for situations where you want a variable not to have a value, intentionally, and you want to make it clear to both a reader of your code and to PHP that this is what you want. The latter point is particularly relevant when passing variables to functions.

For example, the following pseudocode may print a warning (depending on your error-reporting settings) if the variable $authorization has never been assigned before you pass it to your test_authorization() function.

```
if (test_authorization($authorization)) {
  // code that grants a privilege of some sort
}
```

On the other hand, code like this:

```
$authorization = NULL;
// code that might or might not set $authorization
if (test_authorization($authorization)) {
  // code that grants a privilege of some sort
}
```

does not cause an unbound-variable warning, assuming that you have written test_authorization() to handle arguments that might be NULL. It also makes clear to a reader of the code that you intend for the variable to lack a value unless there's a case where it is assigned.

Strings

Strings are character sequences, as in the following:

```
$string_1 = "This is a string in double quotes.";
$string_2 = 'This is a somewhat longer, singly quoted string';
$string_39 = "This string has thirty-nine characters.";
$string_0 = ""; // a string with zero characters
```

Strings can be enclosed in either single or double quotation marks, with different behavior at read time. Singly quoted strings are treated almost literally, whereas doubly quoted strings replace variables with their values as well as specially interpreting certain character sequences.

Singly quoted strings

Except for a couple of specially interpreted character sequences, singly quoted strings read in and store their characters literally. The following code:

```
$literally = 'My $variable will not print!\\n';
print($literally);
```

produces the browser output:

```
My $variable will not print!\n
```

Singly quoted strings also respect the general rule that quotation marks of a different type will not break a quoted string. This is legal:

```
$singly_quoted = 'This quote mark: " is no big deal';
```

To embed a single quotation mark (such as an apostrophe) in a singly quoted string, escape it with a backslash, as in the following:

```
$singly_quoted = 'This quote mark\'s no big deal either';
```

Although in most contexts backslashes are interpreted literally in singly quoted strings, you may also use two backslashes (\\) as an escape sequence for a single (nonescaping) backslash. This is useful when you want a backslash as the final character in a string, as in:

```
$win_path = 'C:\\InetPub\\PHP\\';
print("A Windows-style pathname: $win_path<BR>");
```

which is displayed as:

```
A Windows-style pathname: C:\InetPub\PHP\
```

NOTE   **We could have used single backslashes to produce the first two backslashes in the output, but the escaping is necessary at the end of the string so that the closing quotation mark will *not* be escaped.**

These two escape sequences (\\ and \') are the *only* exceptions to the literal-mindedness of singly quoted strings.

Doubly quoted strings

Strings that are delimited by double quotes (as in `"this"`) are preprocessed in both the following two ways by PHP:

- Certain character sequences beginning with backslash (\) are replaced with special characters.
- Variable names (starting with $) are replaced with string representations of their values.

The escape-sequence replacements are:

- \n is replaced by the newline character
- \r is replaced by the carriage-return character
- \t is replaced by the tab character
- \$ is replaced by the dollar sign itself ($)
- \" is replaced by one double quotation mark (")
- \\ is replaced by a single backslash (\)

The first three of these replacements make it easy to visibly include certain whitespace characters in your strings. The \$ sequence lets you include the $ symbol when you want it, without it being interpreted as the start of a variable. The \" sequence is there so that you can include a double quotation mark symbol without terminating your doubly quoted string. Finally, because the \ character starts all these sequences, you need a way to include that character literally, without it starting an escape sequence — to do this, you preface it with itself.

Just as with singly quoted strings, quotes of the opposite type can be freely included without an escape character:

```
$has_apostrophe = "There's no problem here";
```

Single versus double quotation marks

PHP does some preprocessing of doubly quoted strings (strings with quotation marks like "this") before constructing the string value itself. For one thing, variables are replaced by their values (as in the preceding example). To see that this replacement is really about the quoted string rather than the print construct, consider the following code:

```
$animal = "antelope"; // first assignment
$saved_string = "The animal is $animal<BR>";
$animal = "zebra"; // reassignment
print("The animal is $animal<BR>"); //first display line
print($saved_string); //second display line
```

What output would you expect here? As it turns out, your browser would display:

```
The animal is zebra
The animal is antelope
```

And the browser displays the preceding output in exactly that order. This is because "antelope" is spliced into the string $saved_string, before the $animal variable is reassigned. In addition to splicing variable values into doubly quoted strings, PHP also replaces some special multiple-character *escape sequences* with their single-character values. The most commonly used is the end-of-line sequence ("\n") — in reading a string like:

```
"The first line \n\n\nThe fourth line"
```

Variable interpolation

Whenever an unescaped $ symbol appears in a doubly quoted string, PHP tries to interpret what follows as a variable name and splices the current value of that variable into the string. Exactly what kind of substitution occurs depends on how the variable is set:

■ If the variable is currently set to a string value, that string is interpolated (or spliced) into the doubly quoted string.

■ If the variable is currently set to a nonstring value, the value is converted to a string, and then that string value is interpolated.

■ If the variable is not currently set, PHP interpolates nothing (or, equivalently, PHP splices in the empty string).

For example:

```
$this = "this";
$that = "that";
$the_other = 2.2000000000;
print("$this,$not_set,$that+$the_other<BR>");
```

produces the PHP output

```
this,,that+2.2<BR>
```

which in turn, when seen in a browser, looks like:

```
this,,that+2.2
```

If you find any part of this example puzzling, it is worth working through exactly what PHP does to parse the string in the print statement. First, notice that the string has four $ signs, each of which is interpreted as starting a variable name. These variable names terminate at the first occurrence of a character that is not legal in a variable name. Legal characters are letters, numbers, and underscores; the *illegal* terminating characters in the preceding print string are (in order) a comma, another comma, the plus symbol (+), and a left angle bracket (<). The first two variables are bound to strings ('this' and 'that'), so those strings are spliced in literally. The next variable ($not_set) has never been assigned, so it is omitted entirely from the string under construction. Finally, the last variable ($the_other) is discovered to be bound to a double — that value is converted to a string ("2.2"), which is then spliced into our constructed string.

CROSS-REF   **For more about converting numbers to strings, see the "Assignment and Coercion" section in Chapter 25.**

As we said earlier in this chapter, all this interpretation of doubly quoted strings happens when the string is *read,* not when it is printed. If we saved the example string in a variable and printed it out later, it would reflect the variable values in the preceding code even if the variables had been changed in the meantime.

Newlines in strings

Although PHP offers an escape sequence (\n) for newline characters, it is good to know that you can literally include new lines in the middle of strings, which PHP also treats as a newline characters. This capability turns out to be convenient when creating HTML strings, because browsers will ignore the line breaks anyway, so you can format your strings with line breaks to make your PHP code lines short:

```
print("<HTML><HEAD></HEAD><BODY>My HTML page is too big
to fit on a single line, but that doesn't mean that I
need multiple print statements!</BODY></HTML>");
```

We produced this statement in our text editor by literally hitting the Enter key at the end of the first two lines — these newlines are preserved in the string, so the single `print` statement will produce three distinct lines of PHP output. (Your mileage may vary depending on your text editor — if your editor automatically wraps lines in displaying them, you may see three lines of code that are actually one long line.) Of course, the browser program will ignore these newlines and will make its own decisions about whether and where to break the lines in display, but you will see the linebreaks if you use View Source in your browser to see the HTML itself.

Limits

There are no artificial limits on string length — within the bounds of available memory, you ought to be able to make arbitrarily long strings.

Output

Most of the constructs in the PHP language execute *silently* — they don't print anything to output. The only way that your embedded PHP code will display anything in a user's browser program is either by means of statements that print something to output or by calling functions that, in turn, call `print` statements.

Echo and print

The two most basic constructs for printing to output are `echo` and `print`. Their language status is somewhat confusing, because they are basic constructs of the PHP language, rather than being functions. As a result, they can be used either with parentheses or without them. (Function calls always have the name of the function first, followed by a parenthesized list of the arguments to the function.)

Echo

The simplest use of echo is to print a string as argument, for example:

```
echo "This will print in the user's browser window.";
```

Or equivalently:

```
echo("This will print in the user's browser window.");
```

Both of these statements will cause the given sentence to be displayed, without displaying the quote signs. (*Note for C programmers:* Think of the HTTP connection to the user as the standard output stream for these functions.)

You can also give multiple arguments to the unparenthesized version of echo, separated by commas, as in:

```
echo "This will print in the ", "user's browser window.";
```

The parenthesized version, however, will not accept multiple arguments:

```
echo ("This will produce a ", "PARSE ERROR!");
```

Print

The command print is very similar to echo, with two important differences:

- Unlike echo, print can accept only one argument.
- Unlike echo, print returns a value, which represents whether or not the print statement succeeded.

The value returned by print is always 1.

Both echo and print are usually used with string arguments, but PHP's type flexibility means that you can throw pretty much any type of argument at them without causing an error. For example, the following two lines will print exactly the same thing:

```
print("3.14159");  // print a string
print(3.14159);    // print a number
```

Technically, what is happening in the second line is that, because print expects a string argument, the floating-point version of the number is converted to a string value before print gets hold of it. However, the effect is that both print and echo will reliably print out numbers as well as string arguments.

For the sake of simplicity and uniformity, we will typically use the parenthesized version of print in our examples, rather than using echo.

Variables and strings

C programmers are accustomed to using a function called `printf`, which allows you to splice values and expressions into a specially formatted printing string. PHP has analogous functions (which we will cover in Chapter 6), but as it turns out we can get much of the same functionality just by using `print` (or `echo`) with quoted strings. For example, the fragment:

```
$animal = "antelope";
$animal_heads = 1;
$animal_legs = 4;
print("The $animal has $animal_heads head(s).<BR>");
print("The $animal has $animal_legs leg(s).<BR>");
```

will produce the following output in the browser:

```
The antelope has 1 head(s).
The antelope has 4 leg(s).
```

The values for the variables we included in the string have been neatly spliced into the printed output. This makes it very easy to quickly produce web pages with content that varies depending on how variables have been set. It is not the result of any magical properties of `print`, however — the magic is really happening in the interpretation of the quoted string itself.

HTML and linebreaks

One mistake often made by new PHP programmers (especially those from a C background) is to try to break lines of text in their browsers by putting end-of-line characters (`"\n"`) in the strings they print. To understand why this doesn't work, you have to distinguish the *output* of PHP (which is usually HTML code, ready to be sent over the Internet to a browser program) from the way that output is rendered by the user's browser. Most browser programs will make their own choices about how to split up lines in HTML text, unless you force a line break with the `<BR>` tag. End-of-line characters in strings will put line breaks in the HTML source that PHP sends to your user's browser (which can still be useful for creating readable HTML source), but they will usually have no effect on the way that text looks in a web page.

Summary

PHP code follows a basic set of syntactical rules, mostly borrowed from programming languages such as C and Perl. The syntactical requirements of PHP are minimal, and in general PHP tries to display results when it can rather than generating an error.

PHP has eight types: integer, double, Boolean, NULL, string, array, object, and resource. Five of these are simple types: Integers are whole numbers, doubles are floating-point numbers, Booleans are true-or-false values, NULL has just one value (NULL), and strings are sequences of characters. *Arrays* are a compound type that holds other PHP values, indexed either by integers or by strings. *Objects* are instances of programmer-defined classes, which can contain both member variables and member functions, and which can inherit functions and data types from other classes. (We address arrays in Chapter 8 and objects in Chapter 20.) Finally, *resources* are special references to memory allocated from external programs, which memory PHP frees automatically when they are no longer needed (we cover resources in Chapter 25).

Only values are typed in PHP — variables have no inherent type other than the value of their most recent assignment. PHP automatically converts value types as demanded by the context in which the value is used. The programmer can also explicitly control types by means of both conversion functions and type casts.

PHP code is whitespace insensitive, and although variable names are case sensitive, basic language constructs and function names are not. Simple PHP expressions are combined into larger expressions by operators and function calls, and statements are expressions with a terminating semicolon. Variables are denoted by a leading $ character and are assigned using the = operator. They need no type declarations and have reasonable default values if used before they are assigned. Variable scope is global except inside the body of functions, where it is local to the function unless explicitly declared otherwise.

The simplest way to send output to the user is by using either echo or print, which output the string arguments. They are particularly useful in combination with doubly quoted strings, which automatically replace embedded variables with their values.

Chapter 5

Learning PHP Control Structures and Functions

It's difficult to write interesting programs if you can't make the course of program execution depend on anything. In a weak sense, the behavior of code that prints variables depends on the variable values, but that is as exciting as filling out a template. As programmers, we want programs that react to something (the world, the time of day, user input, or the contents of a database) by doing something different.

This kind of program reaction requires a *control structure,* which indicates how different situations should lead to the execution of different code. In Chapter 4, we informally used the `if` control structure without really explaining it; in this chapter, we lay out every kind of control structure offered by PHP and study their workings in detail.

> **NOTE** Experienced C programmers: Of all the features in PHP, control is probably the most reliably C-like — all the structures you are used to are here, and they work the same way.

The two broad types of control structures we will talk about are *branches* and *loops.* A branch is a fork in the road for a program's execution — depending on some test or other, the program goes either left or right, possibly following a different path for the rest of the program's execution. A loop is a special kind of branch, where one of the execution paths jumps back to the beginning of the branch, repeating the test and possibly the body of the loop.

Before we can make interesting use of control structures, however, we have to be able to construct interesting tests. We'll start from the very simplest of tests, working our way up from the constants `TRUE` and `FALSE` and then move on to using these tests in more complicated code.

Any real programming language has some kind of capability for *procedural abstraction* — a way to name pieces of code so that you can use them as building blocks in writing other pieces of code. Some scripting languages lack this capability, and we can tell you from our own sorrowful experience that complex server-side code can quickly become unmanageable without it.

PHP's mechanism for this kind of abstraction is the *function*. There are really two kinds of functions in PHP — those that have been built into the language by the PHP developers and those defined by individual PHP programmers.

In this chapter, we also look at how to use the large body of functions already provided in PHP and then, a bit later, how to define your own functions. Luckily, there is no real difference between using a built-in function and using your own functions. But first, let's discuss control.

Boolean Expressions

Every control structure in this chapter has two distinct parts: the *test* (which determines which part of the rest of the structure executes), and the *dependent code* itself (whether separate branches or the body of a loop). Tests work by evaluating a *Boolean expression*, an expression with a result treated as either true or false.

Boolean constants

The simplest kind of expression is a simple value, and the simplest Boolean values are the constants `TRUE` and `FALSE`. We can use these constants anywhere we would use a more complicated Boolean expression, and vice versa. For example, we can embed them in the test part of an `if-else` statement:

```
if (TRUE)
 print("This will always print<BR>");
else
 print("This will never print<BR>");
```

Or equivalently:

```
if (FALSE)
 print("This will never print<BR>");
else
 print("This will always print<BR>");
```

Logical operators

Logical operators combine other logical (aka Boolean) values to produce new Boolean values. The standard logical operations (and, or, not, and `exclusive-or`) are supported by PHP, which has alternate versions of the first two, as shown in Table 5-1.

TABLE 5-1

Logical Operators

Operator	Behavior
and	Is true if and only if both of its arguments are true.
or	Is true if either (or both) of its arguments are true.
!	Is true if its single argument (to the right) is false and false if its argument is true.
xor	Is true if either (but not both) of its arguments are true.
&&	Same as and but binds to its arguments more tightly. (See the discussion of precedence later in the chapter.)
||	Same as or but binds to its arguments more tightly.

The && and || operators will be familiar to C programmers. The ! operator is usually called not, since it negates the argument it operates on.

As an example of using logical operators, consider the following expression:

```
(($statement_1 && $statement_2) ||
 ($statement_1 && !$statement_2) ||
 (!$statement_1 && $statement_2) ||
 (!$statement_1 && !$statement_2))
```

This is a *tautology*, meaning that it is always true regardless of the values of the statement variables. There are four possible combinations of truth values for the two variables, each of which is represented by one of the && expressions. One of these four must be true, and because they are linked by the || operator, the entire expression must be true.

Here's another, slightly trickier tautology using xor:

```
(($statement_1 and $statement_2 and
  $statement_3) xor
 ((!($statement_1 and $statement_2)) or
  (!($statement_1 and $statement_3)) or
  (!($statement_2 and $statement_3))))
```

In English, this expression says, "Given three statements, one and only one of the following two things hold — either 1) all three statements are true, or 2) there are two statements that are not both true."

Precedence of logical operators

Just as with any operators, some logical operators have higher precedence than others, although precedence can always be overridden by grouping subexpressions using parentheses. The logical operators listed in declining order of precedence are: !, &&, ||, and, xor, or. Actually, and, xor, and or

have much lower precedence than the others, so that the assignment operator (=) binds more tightly than `and` but less tightly than `&&`.

A complete table of operator precedence and associativity can be found in the online manual at `www.php.net`.

Logical operators short-circuit

One very handy feature of Boolean operators is that they associate left to right, and they *short-circuit*, meaning that they do not even evaluate their second argument if their truth value is unambiguous from their first argument. For example, imagine that you wanted to determine a very approximate ratio of two numbers but also wanted to avoid a possible division-by-zero error. You can first test to make sure that the denominator is not zero by using the `!=` (not-equal-to) operator:

```
if ($denom != 0 && $numer / $denom > 2)
    print("More than twice as much!");
```

In the case where `$denom` is zero, the `&&` operator should return false regardless of whether the second expression is true or false. Because of short-circuiting, the second expression is not evaluated, so an error is avoided. In the case where `$denom` is not zero, the `&&` operator does not have enough information to reach a conclusion about its truth value, so the second expression is evaluated.

So far, all we've formally covered are the `TRUE` and `FALSE` constants and how to combine them to make other true-or-false values. Now we'll move on to operators that actually let you make meaningful Boolean tests.

Comparison operators

Table 5-2 shows the comparison operators, which can be used for either numbers or strings (although you should see the cautionary sidebar entitled "Comparing Things That Are Not Integers").

TABLE 5-2

Comparison Operators

Operator	Name	Behavior
==	Equal	True if its arguments are equal to each other, false otherwise
!=	Not equal	False if its arguments are equal to each other, true otherwise
<	Less than	True if the left-hand argument is less than its right-hand argument but false otherwise
>	Greater than	True if the left-hand argument is greater than its right-hand argument but false otherwise
<=	Less than or equal to	True if the left-hand argument is less than its right-hand argument or equal to it but false otherwise

Operator	Name	Behavior
>=	Greater than or equal to	True if the left-hand argument is greater than its right-hand argument or equal to it but false otherwise
===	Identical	True if its arguments are equal to each other and of the same type but false otherwise

As an example, here are some variable assignments, followed by a compound test that is always true:

```
$three = 3;
$four = 4;
$my_pi = 3.14159;
if (($three == $three) and
    ($four === $four) and
    ($three != $four) and
    ($three < $four) and
    ($three <= $four) and
    ($four >= $three) and
    ($three <= $three) and
    ($my_pi > $three) and
    ($my_pi <= $four))
 print("My faith in mathematics is restored!<BR>");
else
 print("Sure you typed that right?<BR>");
```

CAUTION   Watch out for a very common mistake: confusing the assignment operator (=) with the comparison operator (==). The statement `if ($three = $four)` will (probably unexpectedly) set the variable `$three` to be the same as `$four`; what's more, the test will be true if `$four` is a true value!

Operator precedence

Although overreliance on precedence rules can be confusing for the person who reads your code next, it's useful to note that comparison operators have higher precedence than Boolean operators. This means that a test like the following:

```
if ($small_num > 2 && $small_num < 5) ...
```

doesn't need any parentheses other than those shown.

String comparison

The comparison operators may be used to compare strings as well as numbers (see the cautionary sidebar). We would expect the following code to print its associated sentence (with apologies to Billy Bragg):

```
if (("Marx" < "Mary") and
    ("Mary" < "Marzipan"))
```

```
{
  print("Between Marx and Marzipan in the ");
  print("dictionary, there was Mary.<BR>");
}
```

The comparisons are case sensitive, and the only reason that this example will print anything is because our values are case-consistent. Because of the capitalization of Dennis, the following will not print anything:

```
if (("deep blue sea" < "Dennis") and
    ("Dennis" < "devil"))
{
  print("Between the deep blue sea and ");
  print("the devil, that was me.<BR>");
}
```

Comparing Things That Are Not Integers

Although comparison operators work with numbers or strings, a couple of gotchas lurk here.

First of all, although it is always safe to do less-than or greater-than comparisons on doubles (or even between doubles and integers), it can be dangerous to rely on equality comparisons on doubles, especially if they are the result of a numerical computation. The problem is that a rounding error may make two values that are theoretically equal differ slightly.

Second, although comparison operators work for strings as well as numbers, PHP's automatic type conversions can lead to counterintuitive results when the strings are interpretable as numbers. For example, the code:

```
$string_1 = "00008";
$string_2 = "007";
$string_3 = "00008-OK";
if ($string_2 < $string_1)
  print("$string_2 is less than $string_1<BR>");
if ($string_3 < $string_2)
  print("$string_3 is less than $string_2<BR>");
if ($string_1 < $string_3)

  print("$string_1 is less than $string_3<BR>");
```

gives this output (with comments added):

```
007 is less than 00008  // numerical comparison
00008-OK is less than 007 // string comparison
00008 is less than 00008-OK // string comp. - contradiction!
```

When it can, PHP will convert string arguments to numbers, and when both sides can be treated that way, the comparison ends up being numerical, not alphabetic. The PHP designers view this as a feature, not a bug. Our view is that if you are comparing strings that have any chance of being interpreted as numbers, you're better off using the strcmp() function.

The ternary operator

One especially useful construct is the *ternary conditional operator*, which plays a role somewhere between a Boolean operator and a true branching construct. Its job is to take three expressions and use the truth value of the first expression to decide which of the other two expressions to evaluate and return. The syntax looks like:

```
testExpression ? yesExpression : noExpression
```

The value of this expression is the result of `yes-expression` if `test-expression` is true; otherwise, it is the same as `no-expression`.

For example, the following expression assigns to `$max_num` either `$first_num` or `$second_num`, whichever is larger:

```
$max_num = $first_num > $second_num ? $first_num : $second_num;
```

As you will see, this is equivalent to:

```
if ($first_num > $second_num)
   $max_num = $first_num;
else
   $max_num = $second_num;
```

but is somewhat more concise.

Branching

The two main structures for branching are `if` and `switch`. If is a workhorse and is usually the first conditional structure anyone learns. Switch is a useful alternative for certain situations where you want multiple possible branches based on a single value and where a series of `if` statements would be cumbersome.

If-else

The syntax for `if` is:

```
if (test)
   statement-1
```

Or with an optional `else` branch:

```
if (test)
   statement-1
else
   statement-2
```

When an `if` statement is processed, the `test` expression is evaluated, and the result is interpreted as a Boolean value. If `test` is true, `statement-1` is executed. If `test` is not true, and there is an `else` clause, `statement-2` is executed. If test is false, and there is no `else` clause, execution simply proceeds with the next statement after the `if` construct.

Note that a *statement* in this syntax can be a single statement that ends with a semicolon, a brace-enclosed block of statements, or another conditional construct (which itself counts as a single statement). Conditionals can be nested inside each other to arbitrary depth. Also, the Boolean expression can be a genuine Boolean (TRUE, FALSE, or the result of a Boolean operator or function), or it can be a value of another type interpreted as a Boolean.

CROSS-REF   For the full story on how values of non-Boolean types are treated as Booleans, see Chapter 25. The short version is that the number 0, the string "0", and the empty string, "", are false, and almost every other value is true.

The following example, which prints a statement about the absolute difference between two numbers, shows both the nesting of conditionals and the interpretation of the test as a Boolean:

```
if ($first - $second)
  if ($first > $second)
    {
      $difference = $first - $second;
      print("The difference is $difference<BR>");
    }
  else
    {
      $difference = $second - $first;
      print("The difference is $difference<BR>");
    }
else
  print("There is no difference<BR>");
```

This code relies on the fact that the number 0 is interpreted as a false value — if the difference is zero, then the test fails, and the `no difference` message is printed. If there is a difference, a further test is performed. (This example is artificial, because a test like `$first != $second` would accomplish the same thing comprehensibly.)

Else attachment

At this point, former Pascal programmers may be warily wondering about `else` attachment — that is, how does an `else` clause know which `if` it belongs to? The rules are simple and are the same as in most languages other than Pascal. Each `else` is matched with the nearest unmatched `if` that can be found, while respecting the boundaries of braces. If you want to make sure that an `if` statement stays solo and does not get matched to an `else`, wrap it up in braces like this:

```
if ($num % 2 == 0) // $num is even?
  {
    if ($num > 2)
```

```
        print("num is not prime<BR>");
    }
  else
    print("num is odd<BR>");
```

This code will print num is not prime if $num happens to be an even number greater than 2, num is odd if $num is odd, and nothing if $num happens to be 2. If we had omitted the curly braces, the else would attach to the inner if, and so the code would buggily print num is odd if $num were equal to 2 and would print nothing if $num were actually odd.

> **NOTE**   In this chapter's examples, we often use the modulus operator (%), which is explained in Chapter 9. For the purposes of these examples, all you need to know is that if $x % $y is zero, $x is evenly divisible by $y.

Elseif

It's very common to want to do a cascading sequence of tests, as in the following nested if statements:

```
if ($day == 5)
  print("Five golden rings<BR>");
else
  if ($day == 4)
    print("Four calling birds<BR>");
  else
    if ($day == 3)
      print("Three French hens<BR>");
    else
      if ($day == 2)
        print("Two turtledoves<BR>");
      else
        if ($day == 1)
          print("A partridge in a pear tree<BR>");
```

> **NOTE**   We have indented this code to show the real syntactic structure of inclusions — although this is always a good idea, you will often see code that does not bother with this and where each else line starts in the first column.

This pattern is common enough that there is a special elseif construct to handle it. We can rewrite the preceding example as:

```
if ($day == 5)
  print("Five golden rings<BR>");
elseif ($day == 4)
  print("Four calling birds<BR>");
elseif ($day == 3)
  print("Three French hens<BR>");
elseif ($day == 2)
  print("Two turtledoves<BR>");
elseif ($day == 1)
  print("A partridge in a pear tree<BR>");
```

Branching and HTML Mode

As you may have learned from earlier chapters, you should feel free to use the PHP tags to switch back and forth between HTML mode and PHP mode, whenever it seems convenient. If you need to include a large chunk of HTML in your page that has no dynamic code or interpolated variables, it can be simpler and more efficient to escape back into HTML mode and include it literally than to send it using `print` or `echo`.

What may not be as obvious is that this strategy works even inside conditional structures. That is, you can use PHP to decide what HTML to send and then "send" that HTML by temporarily escaping back to HTML mode.

For example, the following cumbersome code uses `print` statements to construct a complete HTML page based on the supposed gender of the viewer. (We're assuming a nonexistent Boolean function called `female()` that tests for this.)

```
<HTML><HEAD>
<?php
if (cat())
   {
     print("<TITLE>The cat-only site</TITLE><BR>");
     print("</HEAD><BODY>");
     print("This site has been specially constructed ");
     print("for cats only.<BR> No dogs allowed here!");
   }
else
   {
     print("<TITLE>The dog-only site</TITLE><BR>");
     print("</HEAD><BODY>");
     print("This site has been specially constructed ");
     print("for dogs only.<BR> No cats allowed here!");
   }
?>


</BODY></HTML>
```

Instead of all these `print` statements, we can duck back into HTML mode within each of the two branches:

```
<HTML><HEAD>
<?php
if (cat())
   {
?>
<TITLE>The cat-only site</TITLE>
</HEAD><BODY>
This site has been specially constructed
for cats only.<BR> No dogs allowed here!
<?php
```

```
   }
else
   {
?>
<TITLE>The dog-only site</TITLE><BR>
</HEAD><BODY>
This site has been specially constructed
for dogs only.<BR> No cats allowed here!
<?php
   }
?>
</BODY></HTML>
```

This version is somewhat more difficult to read, but the only difference is that it replaces each set of `print` statements with a block of literal HTML that starts with a closing PHP tag (`?>`) and ends with a starting PHP tag (`<?php`).

In this book's examples, we mostly avoid this kind of conditional inclusion, simply because we feel that it may be harder for the novice PHP programmer to decipher. But that shouldn't stop you — literal inclusion has advantages, including fast execution. (In HTML mode, all the PHP engine must do is pass on characters and watch for the next PHP start tag, which is inevitably faster than parsing and executing `print` statements, especially if they include doubly quoted strings.)

A third alternative, when large blocks of HTML are conditionally included, is the heredoc, alluded to in Chapter 4 and explained fully in Chapter 7. The heredoc will allow you to include large blocks of HTML code inside a chunk of PHP without several consecutive print statements.

The `if`, `elseif` construct allows for a sequence of tests that executes only the first branch that has a successful test. In theory, this is syntactically different from the previous example (we have a single construct with five branches rather than a nesting of five two-branch constructs), but the behavior is identical. Use whichever syntax you find more appealing.

Switch

For a specific kind of multiway branching, the `switch` construct can be useful. Rather than branch on arbitrary logical expressions, `switch` takes different paths according to the value of a single expression. The syntax is as follows, with the optional parts enclosed in square brackets (`[]`):

```
switch(expression)
{
   case value-1:
      statement-1;
      statement-2;
      ...
      [break;]
   case value-2:
```

```
        statement-3;
        statement-4;
        ...
        [break;]
    ...
    [default:
        default-statement;]
}
```

The expression can be a variable or any other kind of expression, as long as it evaluates to a simple value (that is, an integer, a double, or a string). The construct executes by evaluating the expression and then testing the result for equality against each case value. As soon as a matching value is found, subsequent statements are executed in sequence until the special statement (`break;`) or until the end of the `switch` construct. (As we'll see in the "Looping" section of this chapter, `break` can also be used to break out of looping constructs.) A special `default` tag can be used at the end, which will match the expression if no other case has matched it so far.

For example, we can rewrite the `if-else` example as follows:

```
switch($day)
{
  case 5:
    print("Five golden rings<BR>");
    break;
  case 4:
    print("Four calling birds<BR>");
    break;
  case 3:
    print("Three French hens<BR>");
    break;
  case 2:
    print("Two turtledoves<BR>");
    break;
  default:
    print("A partridge in a pear tree<BR>");
}
```

This will print a single appropriate line for days 2–5; for any day other than those, it will print A partridge in a pear tree. Although `switch` will accept only a single argument, there's no reason why that argument can't be the value of expressions evaluated previously in your code.

> **CAUTION**   The single most confusing aspect of `switch` is that all cases after a matching case will execute, unless there are `break` statements to stop the execution. In the "partridge" example, the `break` statements ensure that we see only one line from the song at a time. If we remove the `break` statements, we will see a sequence of lines counting down to the final line, just as in the song.

Looping

Congratulations! You just passed the boundary from scripting into real programming. The branching structures we have looked at so far are useful, but there are limits to what can be computed with them alone. On the other hand, it's well established in theoretical computer science that any language with tests plus unbounded looping can do pretty much anything that any other language can do. You may not actually want to write a C compiler in PHP, for example, but it's nice to know that no inherent language limits are going to stop you.

Bounded loops versus unbounded loops

A *bounded loop* executes a fixed number of times — you can tell by looking at the code how many times the loop will iterate, and the language guarantees that it won't loop more times than that. An *unbounded loop* repeats until some condition becomes true (or false), and that condition is dependent on the action of the code within the loop. Bounded loops are predictable, whereas unbounded loops can be as tricky as you like.

Unlike some languages, PHP doesn't actually have any constructs specifically for bounded loops — `while`, `do-while`, and `for` are all unbounded constructs — but as you will see in this section, an unbounded loop can do anything a bounded loop can do.

CROSS-REF   **In addition to the looping constructs in this chapter, PHP provides functions for iterating over the contents of arrays, which are covered in Chapter 8.**

While

The simplest PHP looping construct is `while`, which has the following syntax:

```
while (condition)
   statement
```

The `while` loop evaluates the *condition* expression as a Boolean — if it is true, it executes *statement* and then starts again by evaluating *condition*. If the condition is false, the `while` loop terminates. Of course, just as with `if`, *statement* may be a single statement or it may be a brace-enclosed block. The body of a `while` loop may not execute even once, as in:

```
while (FALSE)
   print("This will never print.<BR>");
```

Or it may execute forever, as in this code snippet:

```
while (TRUE)
   print("All work and no play makes
         Jack a dull boy.<BR>");
```

Or it may execute a predictable number of times, as in:

```
$count = 1;
while ($count <= 10)
   {
     print("count is $count<BR>");
     $count = $count + 1;
   }
```

which will print exactly 10 lines. (For more interesting examples, see the "Looping examples" section, later in this chapter.)

Do-while

The do-while construct is similar to while, except that the test happens at the end of the loop. The syntax is:

```
do statement
   while (expression);
```

The statement is executed once, and then the expression is evaluated. If the expression is true, the statement is repeated until the expression becomes false. The only practical difference between while and do-while is that the latter will always execute its statement at least once. For example:

```
$count = 45;
do
  {
     print("count is $count<BR>");
     $count = $count + 1;
  }
  while ($count <= 10);
```

prints the single line:

```
count is 45
```

For

The most complicated looping construct is for, which has the following syntax:

```
for (initial-expression;
     termination-check;
     loop-end-expression)
   statement
```

In executing a for statement, first the *initial-expression* is evaluated just once, usually to initialize variables. Then *termination-check* is evaluated — if it is false, the for statement concludes, and if it is

true, the statement executes. Finally, the *loop-end-expression* is executed and the cycle begins again with *termination-check*. As always, by *statement* we mean a single (semicolon-terminated) statement, a brace-enclosed block, or a conditional construct.

If we rewrote the preceding `for` loop as a `while` loop, it would look like this:

```
initial-expression;
while (termination-check)
  {
    statement
    loop-end-expression;
  }
```

Actually, although the typical use of `for` has exactly one initial-expression, one termination-check, and one loop-end-expression, it is legal to omit any of them. The termination-check is taken to be always true if omitted, so:

```
for (;;)
    statement
```

is equivalent to:

```
while (TRUE)
    statement
```

It is also legal to include more than one of each kind of `for` clause, separated by commas. The termination-check will be considered to be true if any of its subclauses is true; it is like an `'or'` test. For example, the following statement:

```
for ($x = 1, $y = 1, $z = 1;   //initial expressions
     $y < 10, $z < 10;         // termination checks
     $x = $x + 1, $y = $y + 2, // loop-end expressions
     $z = $z + 3)
  print("$x, $y, $z<BR>");
```

would give the browser output:

```
1, 1, 1
2, 3, 4
3, 5, 7
```

Although the `for` syntax is the most complex of the looping constructs, it is often used for simple bounded loops, using the following idiom:

```
for ($count = 0; $count < $limit; $count = $count + 1)
    statement
```

Looping examples

Now let's look at some examples.

A bounded for loop

Listing 5-1 shows a typical use of bounded `for` loops. The page produced by Listing 5-1 is shown in Figure 5-1.

LISTING 5-1

A division table

```php
<?php
  $start_num = 1;
  $end_num = 10;
?>
<HTML>
<HEAD>
<TITLE>A division table</TITLE>
</HEAD>
<BODY>
<H2>A division table</H2>
<TABLE BORDER=1>
<?php
  print("<TR>");
  print("<TH> </TH>");
  for ($count_1 = $start_num;
       $count_1 <= $end_num;
       $count_1++)
    print("<TH>$count_1</TH>");
  print("</TR>");

  for ($count_1 = $start_num;
       $count_1 <= $end_num;
       $count_1++)
  {
    print("<TR><TH>$count_1</TH>");
    for ($count_2 = $start_num;
         $count_2 <= $end_num;
         $count_2++)
    {
      $result = $count_1 / $count_2;
      printf("<TD>%.3f</TD>",
             $result);  // see Chapter 7
    }
    print("</TR>\n");
  }
```

```
?>
</TABLE>
</BODY>
</HTML>
```

A division table

The main body of this code simply has one `for` loop nested inside another, with each loop executing 10 times, resulting in a 10 x 10 table. Each iteration of the outer loop prints a row, whereas each inner iteration prints a cell. The only novel feature is the way we chose to print the numbers — we used `printf` (covered in Chapter 7), which allows us to control the number of decimal places printed.

NOTE: The `$variable_name++` **feature used above is called an increment. It's a fairly standard shorthand for** `$variable_name + 1`**.**

An unbounded while loop

Now let's look at a loop not so obviously bounded. The sole purpose of the code in Listing 5-2 is to approximate the square root of 81 (using Newton's method). The approximation starts with a guess of 1 and then "zeros in" on the actual square root of 9 by improving the guesses. A trace of this approximation is shown in Figure 5-2.

LISTING 5-2

Approximating a square root

```
<HTML>
<HEAD>
<TITLE>Approximating a square root</TITLE>
</HEAD>
<BODY>
<H3>Approximating a square root</H3>

<?php
$target = 81;
$guess = 1.0;
$precision = 0.0000001;

$guess_squared = $guess * $guess;
while (($guess_squared - $target > $precision) or
       ($guess_squared - $target < - $precision))
{
  print("Current guess: $guess is the square
        root of $target<BR>");
  $guess = ($guess + ($target / $guess)) / 2;
  $guess_squared = $guess * $guess;
}
print("$guess squared = $guess_squared<BR>");
?>
</BODY>
</HTML>
```

Now, although it nicely illustrates a potentially unbounded loop, this approximation example is very artificial — first, because PHP already has a perfectly good square-root function (sqrt) and second, because the number 81 is hardcoded into the page. We can't use this page to find the square root of any other number.

Break and continue

The standard way to get out of a looping structure is for the main test condition to become false. The special commands break and continue offer an optional side exit from all the looping constructs, including while, do-while, and for:

■ The break command exits the innermost loop construct that contains it.

■ The continue command skips to the end of the current iteration of the innermost loop that contains it.

Approximating a square root
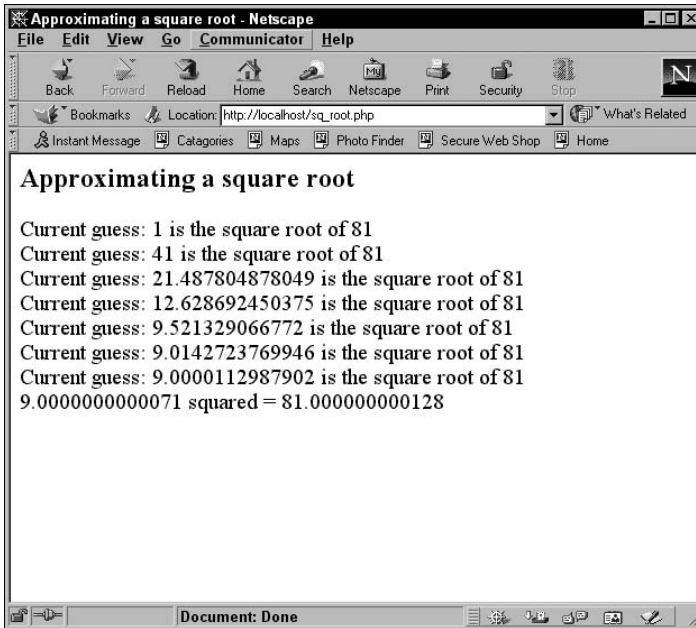
For example, the following code:

```
for ($x = 1; $x < 10; $x++)
{
  // if $x is odd, break out
  if ($x % 2 != 0)
    break;
  print("$x ");
}
```

prints nothing, because 1 is odd, which terminates the `for` loop immediately. On the other hand, the code:

```
for ($x = 1; $x < 10; $x++)
{
  // if $x is odd, skip this loop
  if ($x % 2 != 0)
    continue;
  print("$x ");
}
```

prints:

```
2 4 6 8
```

because the effect of the `continue` statement is to skip the printing of any odd numbers.

Using the `break` command, the programmer can choose to dispense with the main termination test altogether. Consider the following code, which prints a list of prime numbers (that is, numbers not divisible by something other than 1 or the number itself):

```
$limit = 500;
$to_test = 2;
while(TRUE)
{
  $testdiv = 2;
  if ($to_test > $limit)
    break;
  while (TRUE)
  {
    if ($testdiv > sqrt($to_test))
      {
        print "$to_test ";
        break;
      }
    // test if $to_test is divisible by $testdiv
    if ($to_test % $testdiv == 0)
      break;
    $testdiv = $testdiv + 1;
  }
  $to_test = $to_test + 1;
}
```

In the preceding code, we have two `while` loops — the outer loop works through all the numbers between 1 and 500, and the inner loop actually does the testing with each possible divisor. If the inner loop finds a divisor, the number is not prime, so it breaks out without printing anything. If, on the other hand, the testing gets as high as the square root of the number, we can safely assume that the number must be prime, and the inner loop is broken without printing. Finally, the outer loop is broken when we have reached the limit of numbers to test. The result in this case is a list of primes less than 500:

```
2 3 5 7 11 13 17 19 23 29 31 37 41 43 47 53 59 61 67 71 73 79 83
89 97 101 103 107 109 113 127 131 137 139 149 151 157 163 167
173 179 181 191 193 197 199 211 223 227 229 233 239 241 251 257
263 269 271 277 281 283 293 307 311 313 317 331 337 347 349 353
359 367 373 379 383 389 397 401 409 419 421 431 433 439 443 449
457 461 463 467 479 487 491 499
```

Notice that it is crucial to this code that `break` interrupt the inner `while` loop only.

CROSS-REF There is another iteration construct, called `foreach`, which is used only for iterating over arrays. We cover it in Chapter 8.

A note on infinite loops

If you've ever programmed in another language, you've probably had the experience of accidentally creating an infinite loop (a looping construct whose exit test never becomes true and so never returns). The first thing to do when you realize this has happened is to interrupt the program, which will otherwise continue "forever" and use up a lot of CPU time. But what does it mean to interrupt a PHP script? Is it sufficient to click the Stop button on your browser?

As it turns out, the answer is dependent on some PHP configuration settings — you can set the PHP engine to ignore interruptions from the browser (like the result of clicking Stop) and also to impose a time limit on script execution (so that "forever" will only be a short time). The default configuration for PHP is to ignore interruptions, but with a script time limit of 30 seconds — the time limitation means that you can afford to forget about infinite loops that you may have started.

CROSS-REF For more on the configuration of PHP, see Chapter 29.

Alternate Control Syntaxes

PHP offers another way to start and end the bodies of the `if`, `switch`, `for`, and `while` constructs. It amounts to replacing the initial brace of the enclosed block with a colon and the closing brace with a special ending statement for that construct (`endif`, `endswitch`, `endfor`, or `endwhile`). For example, the `if` syntax becomes:

```
if (expression):
  statement1
  statement2
  ..
endif;
```

Or:

```
if (expression):
  statement1
  statement2
  ..
elseif (expression2):
  statement3
  ..
else:
  statement4
  ..
endif;
```

Note that the `else` and `elseif` bodies also begin with colons. The corresponding `while` syntax is:

```
while (expression):
  statement
endwhile;
```

Which syntax you use is a matter of taste. The nonstandard syntax in PHP is largely used for historical reasons and for the comfort of people who are familiar with it from the early versions of PHP. We will consistently use the standard syntax in the rest of this book.

Terminating Execution

Sometimes you just have to give up, and PHP offers a construct that helps you do just that. The `exit()` construct takes either a string or a number as argument, prints out the argument, and then terminates execution of the script. Everything that PHP produces up to the point of invoking `exit()` is sent to the client browser as usual, and nothing in your script after that point will even be parsed — execution of the script stops immediately. If the argument given to exit is a number rather than a string, the number will be the return value for the script's execution. Because exit is a construct, not a function, it's also legal to give no argument and omit the parentheses.

The `die()` construct is an alias for `exit()` and so behaves exactly the same way. (We'll usually use the `die()` version because we find the name more evocative.) So what's the point of `exit()` and `die()`? One possible use is to cut off production of a web page when your script has determined that there is no more interesting information to send, without bothering to wrap up the different branches in a conditional construct. This usage can make long scripts somewhat difficult to read and debug, however.

A better use for `die()` is to make your crashes informative. It's good to get into the habit of testing for unexpected conditions that would crash your script if they were true, and throw in a `die()` statement with an informative message. If you're correct in your expectations, the `die()` will never be invoked; if you're wrong, you will have an error message of your own rather than a possibly obscure PHP error. For example, consider the following pseudocode, which assumes that we have functions to make a database connection and that we then use that database connection:

```
$connection = make_database_connection();
if (!$connection)
  die("No database connection!");
use_database_connection($connection);
```

This example assumes that our imaginary function `make_database_connection()`, like many PHP functions, returns a useful value if it succeeds, and a false value if it fails. An even more compact version of the preceding code takes advantage of the fact that `or` has lower precedence than the `=` assignment operator.

```
$connection = make_database_connection()
    or die("No database connection!");
use_database_connection($connection);
```

This works because the `or` operator short-circuits, and therefore the `die()` construct will only be evaluated if the expression `$connection = make_database_connection()` has a false value. Because the value of an assignment expression is the value assigned, this code ends up being equivalent to the earlier version. (Note that this would not work the same way if we used `||` instead of `or`, because `||` has higher precedence than assignment, and so `$connection` would end up being assigned to the true-or-false result of the `||` expression.)

> **NOTE** Before PHP5, the control structures we've presented so far were really the only alternatives; control would flow from the first statement in a file to the last (possibly bounced around by function calls), unless prematurely terminated with `die()`. With exception handling, PHP5 introduces an alternate way to deal with problematic conditions, and one that is much more flexible than `die()`. We treat exceptions briefly later in this chapter, and more thoroughly in Chapter 30.

In Table 5-3, we summarize all the control structures you've seen thus far.

TABLE 5-3

PHP Control Structures

Name	Syntax	Behavior

If (or if-else)	`if (test)statement-1`
	        `-or-`
	`if (test)`
	  `statement-1`
	`else`
	  `statement-2`
	        `-or-`
	`if (test)`
	  `statement-1`
	`elseif (test2)`
	  `statement-2`
	`else`
	  `statement-3`
		Evaluate *test* and if it is true, execute *statement-1*. If *test* is false and there is an `else` clause, execute *statement-2*. The `elseif` construct is a syntactic shortcut for `else` clauses, where the included statement is itself an `if` construct. Statements may be single statements terminated with a semicolon or brace-enclosed blocks.

Ternary operator	`expression-1 ?`
	`expression-2 :`
	`expression-3`
		Evaluate *expression-1* and interpret it as a Boolean. If it is true, evaluate *expression-2* and return it as the value of the entire expression. Otherwise, evaluate and return *expression-3*.

continued

TABLE 5-3	*(continued)*	

<div align="center">

PHP Control Structures

</div>

Name	Syntax	Behavior

If
(or if-else)
	`if (test)statement-1`
	`        -or-`
	`if (test)`
	`  statement-1`
	`else`
	`  statement-2`
	`        -or-`
	`if (test)`
	`  statement-1`
	`elseif (test2)`
	`  statement-2`
	`else`
	`  statement-3`
		Evaluate *test* and if it is true, execute *statement-1*. If *test* is false and there is an `else` clause, execute *statement-2*. The `elseif` construct is a syntactic shortcut for `else` clauses, where the included statement is itself an `if` construct.
		Statements may be single statements terminated with a semicolon or brace-enclosed blocks.

Ternary operator	`expression-1 ?`
	`expression-2 :`
	`expression-3`
		Evaluate *expression-1* and interpret it as a Boolean. If it is true, evaluate *expression-2* and return it as the value of the entire expression. Otherwise, evaluate and return *expression-3*.

Switch	`switch(expression)`
	`{`
	`  case value-1:`
	`    statement-1`
	`    statement-2`
	
	`    …`
	`    [break;]`
	`  case value-2:`
	`    statement-3`
	`    statement-4`
	
	`    …`
	`    [break;]`
	`  …`
	`  [default:`
	`    default-statement]`
	`}`
		Evaluate *expression,* and compare its value to the value in each `case` clause. When a matching case is found, begin executing statements in sequence (including those from later cases), until the end of the `switch` statement or until a `break` statement is encountered. The optional `default` case will execute if no other case has matched the expression.

While	`while (condition)` `  statement`	Evaluate *condition* and interpret it as Boolean. If *condition* is false, the `while` construct terminates. If it is true, execute *statement,* and keep executing it until *condition* becomes false. Terminate the `while` loop if the special `break` command is encountered, and skip the rest of the current iteration if `continue` is encountered.

Name	Syntax	Behavior
Do-while	```do statement while (condition);```	Perform *statement* once unconditionally, then keep repeating statement until *condition* becomes false. (The `break` and `continue` commands are handled as in `while`.)
For	```for (initial-expression; termination-check; loop-end- expression) statement```	Evaluate *initial-expression* once unconditionally. Then if *termination-check* is true, evaluate *statement,* and then *loop-end-expression,* and repeat that loop until *termination-check* becomes false. Clauses may be omitted, or multiple clauses of the same kind can be separated with commas — a missing `termination-check` is treated as true. (The `break` and `continue` commands are handled as in `while`.)

Using Functions

The basic syntax for using (or *calling*) a function is:

```
function_name(expression_1, expression_2, ..., expression_n)
```

This includes the name of the function followed by a parenthesized and comma-separated list of input expressions (which are called the *arguments* to the function). Functions can be called with zero or more arguments, depending on their definitions.

When PHP encounters a function call, it first evaluates each argument expression and then uses these values as inputs to the function. After the function executes, the returned value (if any) is the result of the entire function expression.

All the following are valid calls to built-in PHP functions:

```
sqrt(9); // square root function, evaluates to 3
rand(10, 10 + 10); // random number between 10 and 20
strlen("This has 22 characters"); // returns the number 22
pi(); // returns the approximate value of pi
```

These functions are called with 1, 2, 1, and 0 arguments, respectively.

Return values versus side effects

Every function call is a PHP expression, and (just as with other expressions) there are only two reasons why you might want to include one in your code: for the *return value* or for the *side effects*.

The *return value* of a function is the value of the function expression itself. You can do exactly the same things with this value as with the results of evaluating any other expression. For example, you can assign it to a variable, as in:

```
$my_pi = pi();
```

Or you can embed it in more complicated expressions, as in:

```
$approx = sqrt($approx) * sqrt($approx)
```

Functions are also used for a wide variety of *side effects*, including writing to files, manipulating databases, and printing things to the browser window. It's okay to make use of both return values and side effects at the same time — for example, it is very common to have a side-effecting function return a value that indicates whether or not the function succeeded.

The result of a function may be of any type, and it is common to use the array type as a way for functions to return multiple values.

Function Documentation

The architecture of PHP has been cleverly designed to make it easy for other developers to extend. The basic PHP language itself is very clean and flexible, but there is not a lot there — most of PHP's power resides in the large number of built-in functions. This means that developers can contribute simply by adding new built-in functions, which is nice especially because it does not change anything that PHP users may be relying on.

Although this book covers many of these built-in functions, explaining some of them in greater detail than the online manual can, the manual at `www.php.net` is the authoritative source for function information. In this book, we get to choose our topics to some extent, whereas the PHP documentation group has the awesome responsibility of covering every aspect of PHP in the manual. Also, although we hope to keep updating this book in future editions, the manual will have the freshest information on new additions to the ever-growing PHP functionality. It's worth looking at some of the different resources that the PHP site and manual offer.

NOTE   **Although the following information is correct at this writing, some details may become dated or inapplicable if the online manual is reorganized.**

To find the manual, head to `www.php.net`. A handy search bar at the top offers quick and easy access to any individual part of the online documentation. Alternatively, find the Documentation item at the top of the page. The Documentation page that this tab leads to has links to manual information in a wide variety of formats and languages.

The largest section of the manual is the function reference, where each built-in function gets its own page of documentation. Typically, each group of functions has a page of general explanation, leading to pages for individual functions. Each function page starts off with the name of the function and a one-line description. This is followed by a C-style header declaration of the function (explained in

the next section), followed by a slightly longer description and possibly an example or two, and then (in the annotated manual) clarifications and gotcha reports from users.

Headers in documentation

For those unfamiliar with C function headers, the very beginning of a function documentation page might be confusing. The format is:

```
return-type function-name(type1 arg1, type2 arg2, . . .);
```

This specifies the type of value the function is expected to return, the name of the function, and the number and expected types of its arguments.

Here is a typical header description:

```
string substr(string string, int start[, int length]);
```

This says that the function `substr()` will return a string and expects to be given a string and two integers as its arguments. Actually, the square brackets around `length` indicate that this argument is optional — so `substr()` should be called either with a string and an int, or a string and two ints.

Unlike in C, the argument types in these documentary headers are not absolute requirements. If you call `substr()` with a number as its first argument, you will not get an error. Instead, PHP will convert the first argument to a string as it begins to execute the function. However, the argument types do document the intent of the function's author, and it is a good idea either to use the function as documented or to understand the type conversion issues well enough that you are sure the result will be what you expect.

In general, the type names used in function documentation will be those of the basic types or of their aliases: integer (or `int`), double (or `float`, `real`), Boolean, string, array, object, resource, and NULL. In addition, you may see the types `void` and `mixed`. The `void` return type means that the function does not return a value at all, whereas the `mixed` argument type means that the argument might be of any type.

Finding function documentation

What's the best way to find information about a function in the manual? That is likely to depend on what kind of curiosity you have. The most common questions about functions are:

- I want to use function X. Now, how does X work again?
- I'd really like to do task Y. Is there a function that handles that for me?

For the first type of curiosity, the full version of the online manual offers an automatic lookup by function name. You can simply type `http://php.net/functionName` and the `functionName` will be searched for automatically. Alternately, the "Search For" box in the upper-right corner of the manual pages defaults to a mode where it searches for specific function names and displays the

corresponding function page if found. (You can also make other choices, including searching the mailing list or the entire online documentation — the latter is a good choice when you don't know the name of the function you want, but can guess at words that appear on its manual page.)

For the second type of curiosity, your best bet is probably to use the hierarchical organization of the function reference. For example, the `substr` function shown in the "Headers in Documentation" section is found in the "String Functions" section. You can browse the chapter list of the function reference for the best fit for the task you want to do.

Defining Your Own Functions

User-defined functions are not a requirement in PHP. You can produce interesting and useful web sites simply with the basic language constructs and the large body of built-in functions. If you find that your code files are getting longer, harder to understand, and more difficult to manage, however, it may be an indication that you should start wrapping some of your code up into functions.

What is a function?

A *function* is a way of wrapping up a chunk of code and giving that chunk a name, so that you can use that chunk later in just one line of code. Functions are most useful when you will be using the code in more than one place, but they can be helpful even in one-use situations, because they can make your code much more readable.

Function definition syntax

Function definitions have the following form:

```
function function-name ($argument-1, $argument-2, ..)
{
  statement-1;
  statement-2;
  ...
}
```

That is, function definitions have four parts:

- The special word `function`
- The name that you want to give your function
- The function's parameter list — dollar-sign variables separated by commas
- The function body — a brace-enclosed set of statements

Just as with variable names, the name of the function must be made up of letters, numbers, and underscores, and it must not start with a number. Unlike variable names, function names are

converted to lowercase before they are stored internally by PHP, so a function is the same regardless of capitalization.

The short version of what happens when a user-defined function is called is:

1. PHP looks up the function by its name (you will get an error if the function has not yet been defined).

2. PHP substitutes the values of the calling arguments (or the *actual parameters*) into the variables in the definition's parameter list (or the *formal parameters*).

3. The statements in the body of the function are executed. If any of the executed statements are `return` statements, the function stops and returns the given value. Otherwise, the function completes after the last statement is executed, without returning a value.

NOTE   **The alert and experienced programmer will have noticed that the preceding description implies call-by-value, rather than call-by-reference. In Chapter 26, we explain the difference and show how to get call-by-reference behavior.**

Function definition example

As an example, imagine that we have the following code that helps decide which size of bottled soft drink to buy. (This is sometime next year, when supermarket shoppers routinely use their wearable wireless web browsers to get to our handy price-comparison site.)

```
$liters_1 = 1.0;
$price_1 = 1.59;
$liters_2 = 1.5;
$price_2 = 2.09;

$per_liter_1 = $price_1 / $liters_1;
$per_liter_2 = $price_2 / $liters_2;
if ($per_liter1 < $per_liter2)
  print("The first deal is better!<BR>");
else
  print("The second deal is better!<BR>");
```

Because this kind of comparison happens in our web site code all the time, we would like to make part of this a reusable function. One way to do this would be the following rewrite:

```
function better_deal ($amount_1, $price_1,
                      $amount_2, $price_2)
{
  $per_amount_1 = $price_1 / $amount_1;
  $per_amount_2 = $price_2 / $amount_2;
  return($per_amount_1 < $per_amount_2);
}

$liters_1 = 1.0;
```

```
$price_1 = 1.59;
$liters_2 = 1.5;
$price_2 = 2.09;

if (better_deal($liters_1, $price_1,
               $liters_2, $price_2))
  print("The first deal is better!<BR>");
else
  print("The second deal is better!<BR>");
```

Our `better_deal` function abstracts out the three lines in the previous code that did the arithmetic and comparison. It takes four numbers as arguments and returns the value of a Boolean expression. As with any Boolean value, we can embed it in the test portion of an `if` statement. Although this function is longer than the original code, there are two benefits to this rewrite: We can use the function in multiple places (saving lines overall), and if we decide to change the calculation, we have to make the change in only one place.

Alternatively, if the only way we ever use these price comparisons is to print which deal is preferred, we can include the printing in the function, like this:

```
function print_better_deal ($amount_1, $price_1,
                            $amount_2, $price_2)
{
  $per_amount_1 = $price_1 / $amount_1;
  $per_amount_2 = $price_2 / $amount_2;
  if ($per_amount_1 < $per_amount_2)
    print("The first deal is better!<BR>");
  else
    print("The second deal is better!<BR>");
}

$liters_1 = 1.0;
$price_1 = 1.59;
$liters_2 = 1.5;
$price_2 = 2.09;

print_better_deal($liters_1, $price_1,
                  $liters_2, $price_2);
```

Our first function used the `return` statement to send back a Boolean result, which was used in an `if` test. The second function has no `return` statement, because it is used for the side effect of printing text to the user's browser. When the last statement of this function is executed, PHP simply moves on to executing the next statement after a function call.

Formal parameters versus actual parameters

In the preceding examples, the arguments we passed to our functions happened to be variables, but this is not a requirement. The actual parameters (that is, the arguments in the function call) may

be any expression that evaluates to a value. In our examples, we could have passed numbers to our function calls rather than variables, as in:

```
print_better_deal(1.0, 1.59, 1.5, 2.09);
```

Also, notice that in the examples we had a couple of cases where the actual parameter variable had the same name as the formal parameter (for example, `$price_1`), and we also had cases where the actual and formal names were different. (`$liters_1` is not the same as `$amount_1`.) As we will see in the next section, this name agreement doesn't matter either way — the names of a function's formal parameters are completely independent of the world outside the function, including the function call itself.

Argument number mismatches

What happens if you call a function with fewer arguments than appear in the definition, or with more? As you might have come to expect by now, PHP handles this without anything crashing, but it may print a warning depending on your settings for error reporting.

Too few arguments

If you supply fewer actual parameters than formal parameters, PHP will treat the unfilled formal parameters as if they were unbound variables. However, under the usual settings for error reporting in PHP6, you will also see a warning printed to the browser.

The default error-reporting setting in PHP6 reports on every kind of error except runtime notices, which are the least serious condition that is detected. The reason you see warnings about too few arguments to a function is that this is treated as a runtime-warning situation (the next most serious category). If you really need function calls that sometimes provide too few arguments and seeing warnings is unacceptable, you have two options for suppressing the warnings:

- You can temporarily change the value of error reporting in your script, with a statement like `error_reporting(E_ALL ^ E_NOTICE ^ E_WARNING;`. This will turn off both runtime notices and runtime warnings from the point where it appears in your script up to the next `error_reporting()` statement (if any). (Note that this is dangerous, as lots of other problems might produce warnings besides the one you're interested in.)

- You can suppress errors for any single expression by using the error-control operator @, which you can put in front of any expression to suppress errors from that expression only. For example, if the function call `my_function()` is producing a warning, `@my_function()` will not. Note that this is dangerous as well because all types of errors except for parse errors will be suppressed.

We don't advise using either of these workarounds, but we provide them because we are such non-judgmental people by nature. PHP actually provides ways to write functions that expect variable numbers of arguments (see the "Variable Numbers of Arguments" section in Chapter 26), and using them is a much better idea than shooting the messenger.

TIP   **Rather than decreasing PHP's reportage of errors, we advise increasing it to the maximum level possible when you are developing new code. You can do this globally by changing the** `php.ini` **file (see Chapter 29) or simply by including the statement** `error_reporting(E_ALL);` **at the top of your scripts. Among other things, this increase in reportage will mean that you will be warned about variables you have forgotten to assign, which is one of the most frequent causes of time-wasting bugs.**

Too many arguments

If you hand too many arguments to a function, the excess arguments will simply be ignored, even when error reporting is set to `E_ALL`. As you will see in Chapter 26, this tolerance turns out to be helpful in defining functions that can take a variable number of arguments.

Functions and Variable Scope

As we said in Chapter 4, outside of functions, the rules about variable scope are simple: Assign a variable anywhere in the execution of a PHP code file, and the value will be there for you later in that file's execution. The rules become somewhat more complicated in the bodies of function definitions, but not much.

The basic principle governing variables in function bodies is: Each function is its own little world. That is, barring some special declarations, the meaning of a variable name inside a function has nothing to do with the meaning of that name elsewhere. (This is a feature, not a bug — you want functions to be reusable in different contexts, and so having the behavior be independent of the context is a good thing. If not for this kind of scoping, you would waste a lot of time chasing down bugs caused by using the same variable name in different parts of your code.)

NOTE   **As of PHP 4.1, there is a small set of global variables that *are* automatically visible from within function definitions, in contradiction to the previous paragraph and the following one. These are the *superglobal arrays* (**`$_POST`**, **`$_GET`**, **`$_SESSION`**, and so on), which contain keys and values corresponding to variable bindings from different sources. For more on these variables and their uses, see Chapter 6.**

The only variable values that a function has access to are the formal parameter variables (which have the values copied from the actual parameters), plus any variables assigned inside the function. This means that you can use local variables inside a function without worrying about their effects on the outside world. For example, consider this function and its subsequent use:

```
function SayMyABCs ()
{
  $count = 0;
  while ($count < 10)
    {
      print(chr(ord('A') + $count));
      $count = $count + 1;
    }
```

```
    print("<BR>Now I know $count letters<BR>");
  }
$count = 0;
SayMyABCs();
$count = $count + 1;
print("Now I've made $count function call(s).<BR>");
SayMyABCs();
$count = $count + 1;
print("Now I've made $count function call(s).<BR>");
```

The intent of `SayMyABCs()` is to print a sequence of letters. (The functions `chr()` and `ord()` translate between letters and their numeric ASCII codes — we use them here just as a trick to generate letters in sequence.) The output of this code is:

```
ABCDEFGHIJ
Now I know 10 letters
Now I've made 1 function call(s).
ABCDEFGHIJ
Now I know 10 letters
Now I've made 2 function call(s).
```

Both the function definition and the code outside the function make use of variables called `$count`, but they refer to different variables and do not clash.

The default behavior of variables assigned inside functions is that they do not interact with the outside world; they act as though they are newly created each time the function is called. Both of these behaviors, however, can be overridden with special declarations.

Global versus local

The scope of a variable defined inside a function is *local* by default, meaning that (as we explained in the previous section) it has no connection with the meaning of any variables outside the function. Using the `global` declaration, you can inform PHP that you want a variable name to mean the same thing as it does in the context outside the function. The syntax of this declaration is simply the word `global`, followed by a comma-delimited list of the variables that should be treated that way, with a terminating semicolon. To see the effect, consider a new version of the previous example. The only difference is that we have declared `$count` to be global, and we have removed its initial assignment to zero inside the function:

```
function SayMyABCs2 ()
{
  global $count;
  while ($count < 10)
    {
      print(chr(ord('A') + $count));
      $count = $count + 1;
    }
  print("<BR>Now I know $count letters<BR>");
```

```
}
$count = 0;
SayMyABCs2();
$count = $count + 1;
print("Now I've made $count function call(s).<BR>");
SayMyABCs2();
$count = $count + 1;
print("Now I've made $count function call(s).<BR>");
```

Our revised version prints the following browser output:

```
ABCDEFGHIJ
Now I know 10 letters
Now I've made 11 function call(s).

Now I know 11 letters
Now I've made 12 function call(s).
```

This is buggy behavior, and the global declaration is to blame. There is now only one $count variable, and it is being increased both inside and outside the function. When the second call to SayMyABCs() happens, $count is already 11, so the loop that prints letters is never entered.

Although this example shows global to bad advantage, it can be quite useful, especially because (as we'll see in Chapter 6) PHP provides some variable bindings to every page even before any of your own code is executed. It can be helpful to have a way for functions to see these variables without the bother of passing them in as arguments with each call.

Static variables

By default, functions retain no memory of their own execution, and with each function call local variables act as though they have been newly created. The static declaration overrides this behavior for particular variables, causing them to retain their values in between calls to the same function. Using this, we can modify our earlier function SayMyABCs2() to give it some memory:

```
function SayMyABCs3 ()
{
   static $count = 0; //assignment only if first time called
   $limit = $count + 10;
   while ($count < $limit)
      {
         print(chr(ord('A') + $count));
         $count = $count + 1;
      }
   print("<BR>Now I know $count letters<BR>");
}
$count = 0;
SayMyABCs3();
$count = $count + 1;
```

```
    print("Now I've made $count function call(s).<BR>");
    SayMyABCs3();
    $count = $count + 1;
    print("Now I've made $count function call(s).<BR>");
```

This memory-enhanced version gives us the following output:

```
ABCDEFGHIJ
Now I know 10 letters
Now I've made 1 function call(s).
KLMNOPQRST
Now I know 20 letters
Now I've made 2 function call(s).
```

The static keyword allows for an initial assignment, which has an effect only if the function has not been called before. The first time SayMyABCs3() executes, the local version of $count is set to zero. The second time the function is called, it has the value it had at the end of the last execution, so we are able to pick up our studies where we left off. Notice that changes to $count outside the function still have no effect on the local value.

Exceptions

You've already seen some fairly primitive error handling in the form of die(), and you might well imagine the custom error handling possibilities implied by the combination of control structures and basic use of print() or printf() commands (more on this in Chapter 26). However, in prior versions of PHP, a chief complaint was the lack of standardized means for handling errors, and separating that means from the application code itself. Enter Exceptions.

Exceptions use the try, catch syntax similar to Java or Python, although programmers using those languages will note the absence of finally.

Let's start with a simple example that has no error handling at all:

```
function print_header($title, $keywords, $description) {
  print("<HTML><HEAD>");
  print("<TITLE>$title</TITLE>");
  print("<META NAME=\"Keywords\" CONTENT=\"$keywords\">");
  print("<META NAME=\"Description\" CONTENT=\"$description\">");
  print("</HEAD><BODY>");
}

print_header('My Page',
  'PHP, Programming, Beer',
  '');
```

The custom function print_header() is designed to make it easy for us to place a standardized, search engine–friendly header at the top of each page. However, we've left the description variable undefined, which will not yield an error, but will leave us without a meaningful description for our

page. Unfortunately, because the function is essentially called correctly and PHP is forgiving in nature, we may never know that we've left off this important detail. Some form of error handling is necessary to point this out, and *Exceptions* provide a handy way of dong so. Consider this revised code:

```
function print_header($title, $keywords, $description) {
  if(strlen($description) < 40)
    throw new Exception('A reasonable description length is
required<BR>');
  print("<HTML><HEAD>");
  print("<TITLE>$title</TITLE>");
  print("<META NAME=\"Keywords\" CONTENT=\"$keywords\">");
  print("<META NAME=\"Description\" CONTENT=\"$description\">");
  print("</HEAD><BODY>");
}

try {
  print_header('My Page',
    'PHP, Programming, Beer',
    '');
} catch (Exception $e) {
  echo($e->getMessage());
}
```

The first new thing in our revised function is a simple test in line 2 suggesting an appropriate minimum length for the `$description` variable. The line immediately following initiates an instance of the `Exception` class with the message suggested by the quoted value.

NOTE **You can create your own classes and extensions of existing classes, including those for exception handling. PHP gives you Exception for free. We'll go into much greater depth on the subject of classes in Chapter 20 and exception handling itself in Chapter 30.**

Next, instead of simply calling the function, we've enclosed the function in a new control structure, the `try...catch` block. If we execute the code as written, PHP first tries to execute the function as described, then it terminates execution almost immediately, because the `$description` variable has failed our simple test. At this point, the script can continue execution after the `try...catch` block, or it can be terminated with `die()` or `exit()`.

Multiple exceptions can be defined in a single function. This is a good idea because it yields more specific information about what exactly happened. Because execution stops with the first exception, only this exception will be caught.

CROSS-REF **Exceptions are a huge topic; they're outlined here so that you can start using them immediately. You'll find nods to exceptions throughout this book, but they are covered in depth in Chapter 30.**

Function Scope

Although the rules about the scope of variable names are fairly simple, the scoping rules for function names are even simpler. There is just one rule in PHP6: Functions must be defined once (and only once) somewhere in the script that uses them. (See the following note about differences between this behavior and PHP3.) The scope of function names is implicitly global, so a function defined in a script is available everywhere in that script. For clarity's sake, however, it is often a good idea to define all your functions before any code that calls those functions.

> **NOTE** In PHP3, functions could be used only after they were defined. This meant that the safest practice was to define (or include the definitions of) all functions early in a given script, before actually using any of them. Beginning with PHP4, scripts are precompiled before being run, and one effect of this precompilation is that the compiler discovers all function definitions before actually running the code. This means that functions and code can appear in any order in a script, as long as all functions are defined once (and only once).

Include and require

It's very common to want to use the same set of functions across a set of web site pages, and the usual way to handle this is with either `include` or `require`, both of which import the contents of some other file into the file being executed. Using either one of these forms is vastly preferable to *cloning* your function definitions (that is, repeating them at the beginning of each page that uses them); when you want to modify your functions, you will have to do it only once. (We covered these forms in Chapter 3, but they are worth reviewing here in the context of including function definitions.)

For example, at the top of a PHP code file we might have lines like:

```
include "basic-functions.inc";
include "advanced-function.inc";
(.. code that uses basic and advanced functions ..)
```

which import two different files of function definitions. (Note that parentheses are optional with both `include()` and `require()`.) As long as the only things in these files are function definitions, the order of their inclusion does not matter.

Both `include` and `require` have the effect of splicing in the contents of their file into the PHP code at the point that they are called. The only difference between them is how they fail if the file cannot be found. The `include` construct will cause a warning to be printed, but processing of the script will continue; `require`, on the other hand, will cause a fatal error if the file cannot be found.

> **NOTE** Note that `include` and `require` are now more similar in their behavior than they used to be. Prior to PHP 4.0.2, `require` had its file contents spliced in statically, before the actual execution of the page; whereas the contents from `include` were spliced in dynamically as the page executed. Among other things, this led to subtle differences in behavior when the `include/require` form was in conditional code. Now, however, both `include` and `require` have the same dynamic behavior. This means, for example, that if an `include/require` form is in a loop executed 10 times, 10 inclusions will be made.

95

Including only once

Sometimes you really want a file to be included once, but not more than once. This is true most often in the case of function definitions. For example, two different function definition files might, in turn, include the same file of utility functions — if a top-level page includes both of these files, the utility functions might be included twice, leading to complaints from PHP that functions are being defined twice.

To the rescue come `include_once` and `require_once`, which act just like their counterparts except that they will not include a file named by a given string if that file has already been included. It's usually better to use the `_once` version, in general, for including function and class definition files.

The include path

When you `include` a filename, PHP searches for a file by that name in the directories specified in the `include_path` (which is settable in your `php.ini` file). The default path includes the same directory as the one the top-level code page is in. See Chapter 29 for details about how to add locations to your `include` path.

In situations where a single instance of PHP serves several virtual sites, it's generally easier and less confusing to PHP to use the `$_SERVER` superglobal array to specify the location of an `include` file:

```
include_once($_SERVER['DOCUMENT_ROOT']."/path/to/include_file");
```

CAUTION   **Remember that included (and required) files are parsed by default in HTML mode rather than in PHP mode. This means that any included file meant to be interpreted as PHP needs to have the usual PHP tags at the beginning and end, though the end tags aren't technically required.**

Recursion

Some compiled languages, like C and C++, impose somewhat complex ordering constraints on how functions are defined. To know how to compile a function, the compiler must know about all the functions that the function calls, which means the called functions must be defined first. So what do you do if two functions each call the other or if one function calls itself? Issues like this led the designers of C to a separation of function declarations (or prototypes) from function definitions (or implementations). The idea is that you use declarations to inform the compiler in advance about the types of arguments and return types of the functions you plan to use, which is enough information for the compiler to handle the actual definitions in any order.

In PHP, this problem goes away, and so there is no need for separate function prototypes. As long as each function that is called is defined once (and only once) in the current code file or one that is included in the course of the current script's execution, PHP will have no problem resolving function calls, regardless of the interleaving of function calls and definitions.

This means that *recursive functions* (functions that call themselves) are no problem in PHP4. For example, we can define a recursive function and then immediately call it:

```
function countdown ($num_arg)
```

```
      {
        if ($num_arg > 0)
          {
            print("Counting down from $num_arg<BR>");
            countdown($num_arg - 1);
          }
      }
    countdown(10);
```

This produces the browser output:

```
Counting down from 10
Counting down from 9
Counting down from 8
Counting down from 7
Counting down from 6
Counting down from 5
Counting down from 4
Counting down from 3
Counting down from 2
Counting down from 1
```

As with all recursive functions, it's important to be sure that the function has a *base case* (a nonrecursive branch) in addition to the recursive case, and that the base case is certain to eventually occur. If the base case is never invoked, the situation is much like a `while` loop where the test is always true — we will have an infinite loop of function calling. In the case of the preceding function, we know that the base case will happen, because every invocation of the recursive case reduces the countdown number, which must eventually become zero. Of course, this assumes that the input is a positive integer rather than a negative number or a double. Notice that our "greater than zero" test guards against infinite recursion even in these cases, whereas a "not equal to zero" test would not.

Similarly, *mutually recursive functions* (functions that call each other) work without a hitch. For example, the following definitions plus function call:

```
function countdown_first ($num_arg)
  {
    if ($num_arg > 0)
      {
        print("Counting down (first) from $num_arg<BR>");
        countdown_second($num_arg - 1);
      }
  }
function countdown_second ($num_arg)
  {
    if ($num_arg > 0)
      {
        print("Counting down (second) from $num_arg<BR>");
        countdown_first($num_arg - 1);
      }
```

```
    }

countdown_first(5);
```

produce the browser output:

```
Counting down (first) from 5
Counting down (second) from 4
Counting down (first) from 3
Counting down (second) from 2
Counting down (first) from 1
```

Summary

PHP has a C-like set of control structures, which branch or loop depending on the value of Boolean expressions, which in turn can be combined using Boolean operators (and, or, xor, !, &&, ||). The structures if and switch are used for simple branching; while, do-while, and for are used for looping, and exit() or die() terminates script execution.

Most of the power of PHP resides in the large number of built-in functions provided by PHP's benevolent army of open source developers. Each of these functions should be documented (albeit briefly) in the online manual at www.php.net.

You can also write your own functions, which are then used in exactly the same way as the built-in functions. Functions are written in a simple C-style syntax, as in the following:

```
function my_function ($arg1, $arg2, ..)
{
statement1;
statement2;
..
return($value);
}
```

User-defined functions can use arguments of any PHP type and can also return values of any type. The types of arguments and return values do not need to be declared.

In PHP, the ordering of function definitions and function calls makes no difference, as long as every function that is called is defined exactly once. There is no need for separate function declarations or prototypes. Variables assigned inside a function are local to that function, unless specified otherwise with the global declaration. Local variables may be declared to be *static,* which means that they hold onto their values in between function calls.

Finally, with our brief treatment of exceptions, we're well on our way to writing thoughtful friendly code that uses standardized error handling.

Chapter 6

Passing Information with PHP

I n this chapter, we'll briefly discuss some things you need to know about passing data between web pages. Some of this information is not specific to PHP but is a consequence of the PHP/HTML interaction or of the HTTP protocol itself.

HTTP Is Stateless

The most important thing to recall about the way the web works is that the HTTP protocol itself is stateless. If you are a poetic soul, you might say that each HTTP request is on its own, with no direction home, like a complete unknown . . . you know how the rest goes.

For the less lyrical among us, this means that each HTTP request — in most cases, this translates to each resource (HTML page, `.jpg` file, style sheet, and so on) being asked for and delivered — is independent of all the others, knows nothing substantive about the identity of the client, and has no memory.

Even if you design your site with very strict one-way navigation (Page 1 leads only to Page 2, which leads only to Page 3, and so on), the HTTP protocol will never know or care that someone browsing Page 2 must have come from Page 1. You cannot set the value of a variable on Page 1 and expect it to be imported to Page 2 by the exigencies of HTTP itself. You can use HTTP to display a form, and someone can enter some information using it — but unless you employ some extra means to pass the information to another page or program, the variable will simply vanish into the ether as soon as you move to another page.

99

This is where a form-handling technology like PHP comes in. PHP will catch the variable tossed from one page to the next and make it available for further use. PHP happens to be unusually good at this type of data-passing function, which makes it fast and easy to employ for a wide variety of web site tasks.

HTML forms are mostly useful for passing a few values from a given page to one single other page of a web site. There are more persistent ways to maintain state over many pageviews, such as cookies and sessions, which we cover in Chapter 24. This chapter will focus on the most basic techniques of information-passing between web pages, which utilize the GET and POST methods in HTTP to create dynamically generated pages and to handle form data.

GET Arguments

The GET method passes arguments from one page to the next as part of the *Uniform Resource Indicator* (you may be more familiar with the term *Uniform Resource Locator,* or *URL*) query string. When used for form handling, GET appends the indicated variable name(s) and value(s) to the URL designated in the ACTION attribute with a question mark separator and submits the whole thing to the processing agent (in this case a web server).

This is an example HTML form using the GET method (save the file under the name sportselect .html):

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN" "http://www.w3.org/
TR/html4/strict.dtd">
<HTML>
<HEAD>
<TITLE>A GET method example, part 1</TITLE>
</HEAD>

<BODY>
<FORM ACTION="sports.php" METHOD="GET">
<P>Choose your favorite sport:<BR>
<SELECT NAME="Sport">
<OPTION VALUE="Baseball">Baseball</OPTION>
<OPTION VALUE="Basketball">Basketball</OPTION>
<OPTION VALUE="Football">Football</OPTION>
<OPTION VALUE="Ice Hockey">Ice Hockey</OPTION>
<OPTION VALUE="Racing">Auto Racing</OPTION>
<OPTION VALUE="Soccer">Soccer</OPTION>
</SELECT>
<P><INPUT TYPE="submit" NAME="Submit" VALUE="Select"></P>
</FORM>
</BODY>
</HTML>
```

When the user makes a selection and clicks the Submit button, the browser agglutinates these elements in this order, with no spaces between the elements:

■ The URL in quotes after the word ACTION (http://localhost/baseball.php)

■ A question mark (?) denoting that the following characters constitute a GET string.

■ A variable NAME, an equal sign, and the matching VALUE (Team=Cubbies)

■ An ampersand (&) and the next NAME-VALUE pair (Submit=Select); further name-value pairs separated by ampersands can be added as many times as the server query-string-length limit allows.

The browser thus constructs the URL string:

```
http://<your-server-name>/sports.php?Sport=Ice+Hockey&Submit=Select
```

It then forwards this URL into its own address space as a new request. The PHP script to which the preceding form is submitted (sports.php) will grab the GET variables from the end of the request string, stuff them into the $_GET superglobal array (explained in a moment), and do something useful with them — in this case, plug one of two values into a text string.

The following code sample shows the PHP form handler for the preceding HTML form:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN" "http://www.w3.org/
TR/html4/strict.dtd">
<HTML>
<HEAD>
<TITLE>A GET method example, part 2</TITLE>
<STYLE TYPE="text/css">
<!--
BODY  {font-size: 24pt;}
-->
</STYLE>
</HEAD>

<BODY>
<P>You've indicated that you like
<?php echo $_GET['Sport']; ?>!</P>
</BODY>
</HTML>
```

Note that the value inputted into the previous page's HTML form field named "Sport" is now available in a PHP variable called $_GET['Sport']. Finally, you should see a page that says You've indicated that you like Ice Hockey! in big type.

NOTE   At this point, it makes some sense to explain just how to access values submitted from page to page. This chapter discusses the two main methods for passing values: `GET` and `POST` (there are others, but they are not covered until Part III). Each method has an associated super-global array, explained in more depth in Chapter 8, which can be distinguished from other arrays by the underscore that begins its name. Each item submitted via the `GET` method is accessed in the handler via the `$_GET` array; each item submitted via the `POST` method is accessed in the handler via the `$_POST` array. The syntax for referencing an item in a superglobal array is simple and 100 percent consistent:

```
$_ARRAY_NAME['index_name']
```

where the `index_name` is the name part of a name-value pair (for the `GET` method), or the name of an HTML form field (for the `POST` method). As in the preceding example, `$_GET['Sport']`, indicates the value of the form select field called `'Sport'`, sent by the `GET` operation in the original file. You must use the array appropriate to the method used to send data. In this case, `$_POST['Sport']` is undefined because no data was `POST`ed by the original form.

The `GET` method of form handling offers one big advantage over the `POST` method: It constructs an actual new and differentiable URL query string. Users can now bookmark this page. The result of forms using the `POST` method is not bookmarkable.

Just because you *can* achieve the desired functionality with `GET` arguments doesn't mean you *should*. The disadvantages of `GET` for most types of form handling are so substantial that the original HTML 4.0 draft specification deprecated its use in 1997. These flaws include:

■ The `GET` method is not suitable for logins because the username and password are fully visible onscreen as well as potentially stored in the client browser's memory as a visited page.

■ Every `GET` submission is recorded in the web server log, data set included.

■ Because the `GET` method assigns data to a server environment variable, the length of the URL is limited. You may have seen what seem like very long URLs using `GET` — but you really wouldn't want to try passing a 300-word chunk of HTML-formatted prose using this method.

CAUTION   The original HTML spec called for query strings to be limited to 255 characters. Although this stricture was later loosened to mere encouragement of a 255-character limit, using a longer string is asking for trouble.

The `GET` method of form handling had to be reinstated by the W3C after much outcry, largely because of the bookmarkability factor. Despite that it's still implemented as the default choice for form handling in all browsers, `GET` now comes with a strong recommendation to deploy it in *idempotent* usages only — in other words, those that have no permanent side effects. Putting two and two together, the single most appropriate form-handling use of `GET` is the search box. Unless you have a compelling reason to use `GET` for non-search-box form handling, use `POST` instead.

A Better Use for GET-Style URLs

Although the actual `GET` method of form handling is deprecated, the style of URL associated with it turns out to be very useful for site navigation. This is especially true for dynamically generated sites such as those often constructed with PHP, because the appended-variable style of URL works particularly smoothly with a template-based content-development system.

As an illustration, imagine you are the proud proprietor of an information-rich web site about solar cars. You've toiled long and hard over informative and attractive pages such as these:

```
Suspension_design.html
Windtunnel_testing.html
friction_braking.html
```

But as your site grows, a flat-file site structure like this can take a lot of time to administer, as even the most trivial changes must be repeated on every page. If the structure of these pages is very similar, you might want to move to a template-based system with PHP.

You might decide to utilize a single template with separate text files for each topic (containing information, photos, comments, and so on):

```
topic.php
  suspension_design.inc
  windtunnel_testing.inc
  friction_braking.inc
```

Or you might decide you needed a larger, more specialized choice of template files:

```
Vehicle_structure.php
  Tubular_frames.inc
Mechanical_systems.php
  Friction_braking.inc
Electrical_systems.php
  Solar_array.inc
racing.php
  race_strategy.inc
```

A simple template file might look something like this (because we haven't included the necessary .inc text files, this example will not actually work):

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN" "http://www.w3.org/
TR/html4/strict.dtd">
<HTML>
<HEAD>
<TITLE>Solar-car topics</TITLE>
<STYLE TYPE="text/css">
<!--
BODY  {font: verdana; font-size: 12pt}
-->
</STYLE>
</HEAD>

<BODY>
<TABLE BORDER=0 CELLPADDING=0 WIDTH="100%">
<TR>
<!-- Navbar, with Get-style URLs. -->
<TD ALIGN=CENTER VALIGN=TOP>
  <P>
```

103

```
    <A HREF="mechanical_systems.php?Name=friction_braking">
<B>Friction braking</B></A>
    <BR>
    <A HREF="mechanical_systems.php?Name=steering">
<B>Steering</B></A>
    <BR>
    <A HREF="mechanical_systems.php?Name=suspension">
<B>Suspension</B></A>
    <BR>
    <A HREF="mechanical_systems.php?Name=tires">
<B>Tires and wheels</B></A>
    <BR>
    </P>
</TD>

<!-- Main body of content -->
<TD ALIGN=LEFT VALIGN=TOP>
<?php include($_GET['Name'] . "inc"); ?>
</TD></TR></TABLE>
</BODY>
</HTML>
```

Notice that the links on the navbar, when clicked, will be handled by the browser as if they were the product of a GET submission.

But even with this solution, you still have to tend part of your garden by hand: making sure that each include file is properly formatted in HTML, adding a new link to the navbar each time you add a new page to the site, and other such chores. Following the general rule to separate form and content as much as is feasible, you might choose to go to another level of abstraction with a database. In that case, a URL such as http://www.example.com/topic.php?topicID=2 would point to a PHP template that makes database calls. (Using a number variable rather than a word makes for faster database interaction.) This system could also automatically add a link to the navbar whenever you added new topics to the database, so it could produce web pages entirely without ongoing human intervention (all right, maybe *entirely* is an exaggeration — but with significantly fewer person-hours of grunt labor).

POST Arguments

POST is the preferred method of form submission today, particularly in *nonidempotent* usages (those that will result in permanent changes), such as adding information to a database. The form data set is included in the body of the form when it is forwarded to the processing agent (in this case, PHP). No visible change to the URL will result according to the different data submitted.

The POST method has one primary advantage:

- There is a much larger limit on the amount of data that can be passed (a couple of megabytes rather than a couple of hundred characters).

`POST` has these disadvantages:

- The results at a given moment cannot be bookmarked.
- Browsers exhibit different behavior when the visitor uses their Back and Forward navigation buttons within the browser.

There is a misguided belief that `POST` is more secure than `GET`. In reality, neither offers any more security than the other. The visitor can still view variables and data being sent with a `POST` just as they can with a `GET`. The only difference is that the data doesn't show up in the address bar. This doesn't mean that it's hidden. Data sent with a `POST` can be viewed and altered by the web site user.

The first and most important rule of programming, especially web programming is:

Never Trust Input

Always assume that the visitor has either maliciously or accidentally altered the data being passed into your application, and validate the data.

Only when the request is secured using SSL or TLS or some other form of encryption is the form data somewhat secure. Nevertheless, the end user or visitor can still see and alter the data. SSL merely encrypts the data in transit, preventing prying eyes on the network from looking at it. SSL does nothing to prevent the visitor from changing form data.

I'll cover much more about security throughout the book. I believe security needs to be included in every aspect of programming, and, therefore, you'll see security tips when appropriate and within context, rather than trying to make sense of them in a specific chapter. Chapter 28 will examine PHP security, concentrating on overall best practices and also server security, as well.

Get and Post Both

Did you know that with PHP you can use both `GET` and `POST` variables on the same page? You might want to do this for a dynamically generated form, for example.

But what if you (deliberately or otherwise) use the same variable name in both the `GET` and the `POST` variable sets? PHP keeps all `ENVIRONMENT`, `GET`, `POST`, `COOKIE`, and `SERVER` variables in the `$GLOBALS` array if you have set the `register_globals` configuration directive to "on" in your `php.ini` file (doing so creates a security risk). If there is a conflict, it is resolved by overwriting the variable values in the order you set, using the `variables_order` option in `php.ini`. Later trumps earlier, so if you use the default `"EGPCS"` value, cookies will triumph over `POST`s that will themselves obliterate `GET`s. You can control the order of overwriting by simply changing the order of the letters on the appropriate line of this file, or even better, turning `register_globals` off and using the new PHP superglobal arrays instead. See the section on superglobals later in this chapter.

Formatting Form Variables

PHP is so efficient at passing data around because the developers made a very handy but (in theory) slightly sketchy design decision. PHP automatically, but invisibly, assigns the variables for you on the new page when you submit a data set using GET or POST. Most of PHP's competitors make you explicitly do this assignment yourself on each page; if you forget to do so or make a mistake, the information will not be available to the processing agent. PHP is faster, simpler, and mostly more goof-proof.

But because of this automatic variable assignment, you need to always use a good NAME attribute for each INPUT. NAME attributes are not strictly necessary in HTML proper — your form will render fine without them — but the data will be of little use because the HTML form-field NAME attribute will be the variable name in the form handler.

In other words, in this form:

```
<FORM ACTION="<?php echo $_SERVER['PHP_SELF']; ?>"
METHOD="POST">
<INPUT TYPE="text" NAME="email">
<INPUT TYPE="submit" NAME="submit" VALUE="Send">
</FORM>
```

the text field named email will cause the creation of a PHP variable called $_POST['email'] when the form is submitted. Similarly, the submit button will lead to the creation of a variable called $_POST['submit'] on the next page. The name you use in the HTML form will be the name of your variable in the PHP form handler.

CAUTION $HTTP_POST_VARS, $HTTP_SERVER_VARS, **and the whole family of these long-form predefined variables were deprecated in PHP5. If you are already an experienced PHP programmer, perhaps with a large body of previously written code lying around, you might want to think about rewriting now for backward compatibility. They are supported for the time being, but their days are numbered. Use** $_POST, $_GET, **and friends instead.**

Remember that you cannot use a variable name beginning with a number — so you should not name your form field something like 5 (you laugh, but we've seen people try to do it) — and PHP variable names are case sensitive. Also, *please* try to use informative variable names rather than a succession of form fields named myvar and e.

TIP   **It's a good idea to standardize how you name form variables, to make your code more readable and so that you spend less time flipping back to the form itself when you are supposed to be writing code to process that form. For example, you might precede all form variables with** frm **to indicate their source. You might then consistently use the first few letters of each identifying word for what a field does, for example,** frmNameFirst, frmOfficeAdd, frmHomeAdd, **and so on. The specific standard you set is less important than having a standard to begin with.**

Another thing to keep in mind when creating your HTML forms is that, if you ever want this form to be displayed with prefilled inputs, you need to set the VALUE attribute. This is particularly relevant to two kinds of forms: those that are used to edit data from a database, and those that are intended to possibly be submitted more than once. The latter case is very common in situations where a form should redisplay on error with values already prefilled — for instance, a registration form that will not work until the user provides a valid e-mail address or other required data.

For example, the form in Listing 6-1 (which represents a retirement savings calculator) is designed to be submitted multiple times while the user fiddles around with the values. Every time you submit the form, the values from the previous go-round will be filled in for you automatically. Note the use of the VALUE attribute in the form fields in this code sample.

LISTING 6-1

Form with prefilled values (retirement_calc.php)

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01//EN" "http://www.w3.org/TR/html4/
strict.dtd">
<HTML>
<HEAD>
<TITLE>A POST example: retirement savings worksheet</TITLE>
<STYLE TYPE="text/css">
<!--
BODY    {font-size: 14pt}
.heading    {font-size: 18pt; color: red}
-->
</STYLE>
</HEAD>

<?php

// This test, along with the Submit button value in the form
// below, will check to see if the form is being rendered for
// the first time (in which case it will display with only the
// default annual gain filled in).

if (!IsSet($_POST['Submit']) || $_POST['Submit'] != 'Calculate')
{
  $_POST['CurrentAge'] = "";
  $_POST['RetireAge'] = "";
  $_POST['Contrib'] = "";
  $Total = 0;
  $AnnGain = 7;
} else {
  $AnnGain = $_POST['AnnGain'];
  $Years = $_POST['RetireAge'] - $_POST['CurrentAge'];
  $YearCount = 0;
```

```
  $Total = $_POST['Contrib'];

  while ($YearCount <= $Years) {
     $Total = round($Total * (1.0 + $AnnGain/100) +
$_POST['Contrib']);
     $YearCount = $YearCount + 1;
  }
}
?>
<BODY>

<DIV ID="Div1" class="heading">
A retirement-savings calculator</DIV>

<P class=blurb>Fill in all the values (except "Nest Egg")
and see how much money you'll have for your retirement
under different scenarios.  You can change the values and
resubmit the form as many times as you like. You must fill
in the two "Age" variables. The "Annual return" variable has
a default inflation-adjusted value (7% = 8% growth minus 1%
inflation) which you can change to reflect your greater
optimism or pessimism.</P>

<FORM ACTION="<?php echo $_SERVER['PHP_SELF']; ?>"
METHOD="POST">
<P>Your age now:
<INPUT TYPE="text" SIZE=5 NAME="CurrentAge"
VALUE="<?php echo $_POST['CurrentAge']; ?>">
<P>The age at which you plan to retire:
<INPUT TYPE="text" SIZE=6 NAME="RetireAge"
VALUE="<?php echo $_POST['RetireAge']; ?>">
<P>Annual contribution:
<INPUT TYPE="text" SIZE=15 NAME="Contrib"
VALUE="<?php echo $_POST['Contrib']; ?>">
<P>Annual return:
<INPUT TYPE="text" SIZE=5 NAME="AnnGain"
VALUE="<?php echo $AnnGain; ?>"> %
<BR><BR>
<P><B>NEST EGG</B>:  <?php echo $Total; ?>
<P><INPUT TYPE="submit" NAME="Submit" VALUE="Calculate">
</FORM>
</BODY>
</HTML>
```
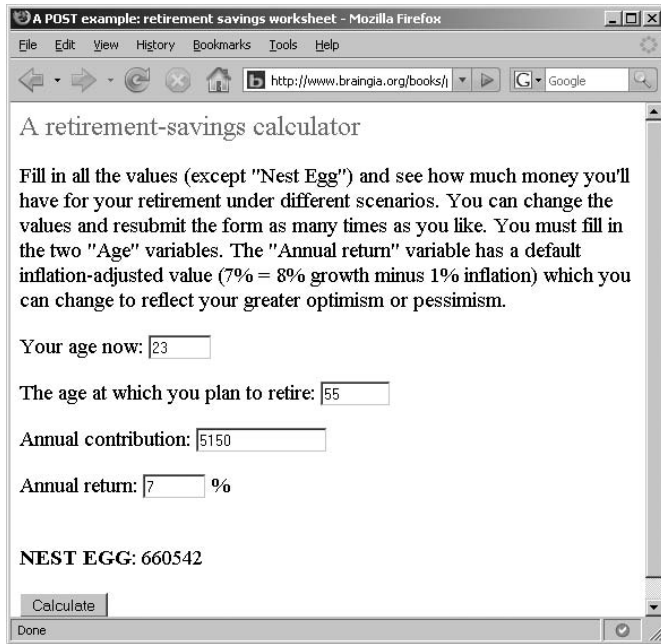
Figure 6-1 shows the result of the Listing 6-1.

A form using the POST method with VALUE attributes

Consolidating forms and form handlers

As you can see in the preceding example, it is often handy to make the HTML form and the form handler into one script. This practice has many advantages, such as making it easier to change the name of the file without harming functionality, making it easier to display error messages and pre-filled form fields, and achieving better control over your variable namespace. Suppose that you are making a login form that redisplays with an error message if the login is unsuccessful. If you have separate forms and form handlers, you'll probably have to do something yucky with GET vars and redirection. If you consolidate, it's very simple to control the display without these machinations.

CROSS-REF   To see how these techniques can be used with data from MySQL, see Chapter 17.

When you consolidate, generally the form-handling code should come before the form display. This order may be something of a shift in thinking for those who are used to writing the form before the handler, but if you think it through, you will see the logic of the practice. You have to give yourself an opportunity to set variables and make choices before you can decide what to show the user. This

is especially relevant if you will be redirecting the user to a different page under certain circumstances, via the `header()` function, because this decision point must come before any HTML output has been displayed to the browser.

PHP Superglobal Arrays

A change that has been coming for a long time in PHP is the gradual phasing out of automatic global variables in favor of *superglobal* arrays, which were introduced in PHP4. Understanding superglobal arrays before you understand arrays may present difficulties; if so, we recommend that you read Chapter 8 and come back to this section later.

In the good old days before PHP4.1, you could write a piece of code like this and expect it to work:

```php
<?php
if (isSet($submit)) {
  echo $email;
} else {
?>
<FORM ACTION="<?php echo $PHP_SELF; ?>" METHOD="POST">
<INPUT TYPE="text" NAME="email">
<INPUT TYPE="submit" NAME="submit" VALUE="Send">
</FORM>
```

All `GET`, `POST`, `COOKIE`, `ENVIRONMENT`, and `SERVER` variables were made global by the `register_globals` directive in `php.ini` and were directly accessible by their names by default.

The PHP team decided to phase out the practice of registering globals, forcing everyone to call these variables as indices in an array (for example, $_POST['secretpassword']). This had already been possible in PHP4, via arrays named $HTTP_GET_VARS, $HTTP_POST_VARS, $HTTP_POST_VARS, and so on, but few developers had used this syntax; frankly, it was a lot of extra keystrokes for a small increase in security. So the PHP team also took this opportunity to rename these arrays with shorter names: $_GET, $_POST, $_COOKIE, $_ENV, and $_SERVER.

These superglobal arrays also have one cool feature that may ameliorate some pain: They are automatically global everywhere. This means, for instance, that you no longer have to pass cookie values into a function or declare the $HTTP_COOKIE_VARS array global before you can access those values in a function. This will help those who functionalize to the max and will be a small amelioration for everyone else.

As of PHP6, `register_globals` is officially gone.

Summary

The HTTP protocol is stateless. This means a plain HTML page is incapable of receiving information from any other page. It can be used to pass values via a URL or an HTML form, but a separate program called a *form handler* must step in to recognize and perform actions on the passed values. In first-generation web development, these form handlers were Perl or C CGI scripts, but nowadays web developers are more likely to use an HTML-embedded programming language such as PHP. PHP makes it particularly easy to write form handlers and even to combine them with HTML display on a single web page.

Information is passed between web pages using one of four main methods: GET, POST, a cookie, or sessions. GET is mainly used to construct complex URL strings for use with dynamically generated pages. Forms are a good way to pass information from one web page to a single other web page. We deal with the persistent state methods, cookies, and sessions in Chapter 24.

Chapter 7

Learning PHP String Handling

Although images, sound files, videos, animations, and applets make up an important portion of the World Wide Web, much of the web is still text — one character's worth after another, like this sentence. The basic PHP data type for representing text is the *string*.

In this chapter, we cover almost all PHP's capabilities for manipulating strings (although we leave more advanced string functions and the pattern-matching power of regular expressions for separate treatment in Chapter 22). We start with the basics of strings, then move to the most commonly used operators and functions.

Strings in PHP

Strings are sequences of characters that can be treated as a unit — assigned to variables, given as input to functions, returned from functions, or sent as output to appear on your user's web page. The simplest way to specify a string in PHP code is to enclose it in quotation marks, whether single quotation marks (') or double quotation marks ("), like this:

```
$my_string = 'A literal string';
$another_string = "Another string";
```

The difference between single and double quotation marks lies in how much interpolation PHP does of the characters between the quote signs before creating the string itself. If you enclose a string in single quotation marks, almost no interpolation will be performed; if you enclose it in

double quotation marks, PHP will splice in the values of any variables you include, as well as make substitutions for certain special character sequences that begin with the backslash (\) character. For example, if you evaluate the following code in the middle of a web page:

```
$statement = 'everything I say';
$question_1 =
  "Do you have to take $statement so literally?\n<BR>";
$question_2 =
  'Do you have to take $statement so literally?\n<BR>';
echo $question_1;
echo $question_2;
```

you should expect to see the browser output:

```
Do you have to take everything I say so literally?
Do you have to take $statement so literally?\n
```

CROSS-REF   For the details on exactly how PHP interprets both singly and doubly quoted strings, see the "Strings" section of Chapter 4.

Interpolation with curly braces

In most situations, you can simply include a variable in a doubly quoted string, and the variable's value will be spliced into the string when it is interpreted. There are two situations where the string parser might very reasonably get confused and need more guidance from you. The first situation is when your notion of where the variable name should stop is not the same as the parser's, and the other occurs when the expression you want to have interpolated is not a simple variable. In these cases, you can clear things up by enclosing the value you want interpolated in curly braces: { }.

For example, PHP has no difficulty with the following code:

```
$sport = 'volleyball';
$plan = "I will play $sport in the summertime";
```

The parser in this case encounters the $ symbol, and then begins collecting characters for a variable name until it runs into the space after $sport. Spaces cannot be part of a variable name, so it is clear that the variable in question is $sport, and PHP successfully finds a value for that variable ('volleyball'), and splices the value in.

Sometimes, though, it is not convenient to stop a variable name with a space. Take this example:

```
$sport1 = 'volley';
$sport2 = 'foot';
$sport3 = 'basket';
$plan1 = "I will play $sport1ball in the summertime";  //wrong
$plan2 = "I will play $sport2ball in the fall";  //wrong
$plan3 = "I will play $sport3ball in the winter";  //wrong
```

You will not get the desired effect here, because PHP interprets $sport1 as part of the variable name $sport1ball, which is probably unbound. Instead, you need something like:

```
$plan1 = "I will play {$sport1}ball in the summertime"; //right
```

which asks PHP to evaluate only the variable expression within the braces before interpolating.

For similar reasons, PHP has difficulty interpolating complex variable expressions, such as multi-dimensional arrays and object variables, unless curly braces are used. The general rule is that if you have a { immediately followed by a $, PHP will evaluate the variable expression up until the closing } and will interpolate the resulting value into the string. (If you need a literal {$ to appear in your string, you can accomplish it by escaping either character with a backslash (\).)

> **TIP** See the "Concatenation and Assignment" section later in this chapter for ideas on other ways to address challenges like this.

Characters and string indexes

Unlike some programming languages, PHP has no distinct character type different from the string type. In general, functions that would take character arguments in other languages expect strings of length 1 in PHP.

You can retrieve the individual characters of a string by including the number of the character, starting at 0, enclosed in curly braces immediately following a string variable. These characters will actually be one-character strings. For example, the following code:

```
$my_string = "Doubled";
for ($index = 0; $index < 7; $index++) {
  $string_to_print = $my_string{$index};
  print("$string_to_print$string_to_print");
}
```

gives the browser output:

```
DDoouubblleedd
```

with each character of the string being printed twice per loop. (The number 7 is hardcoded in this example only because we haven't yet covered how to find out the length of a string — see the function strlen() in the later section "Inspecting strings.")

String operators

PHP offers two string operators: the dot (.) or concatenation operator and the .= concatenating assignment operator. The concatenating assignment operator is discussed in the next section. The concatenation operator, when placed between two string arguments, produces a new string that is the result of putting the two strings together in sequence. For example:

```
$my_two_cents = "I want to give you a piece of my mind ";
```

```
$third_cent = " And another thing";
print($my_two_cents . "..." . $third_cent);
```

gives the output:

```
I want to give you a piece of my mind ... And another thing
```

Note that we are not passing multiple string arguments to the print statement — we are handing it one string argument, which was created by concatenating three strings together. The first and third strings are variables, but the middle one is a literal string enclosed in double quotation marks.

> **NOTE**   Note that the concatenation operator is not + as in Java, and it does not overload any-
> thing else. If you forget this and add strings using +, they will be interpreted as numbers,
> with the result that 'one' + 'two' equals 0 (because no successful string-to-number conversion can
> be made).

Concatenation and assignment

Just as with arithmetic operators, PHP has a shorthand operator (.=) that combines concatenation with assignment. The following statement:

```
$my_string_var .= $new_addition;
```

is exactly equivalent to:

```
$my_string_var = $my_string_var . $new_addition;
```

Note that, unlike commutative addition and multiplication, with this shorthand operator it matters that the new string is added to the right. If you want the new string tacked on to the left, there's no alternative shorter than:

```
$my_string_var = $new_addition . $my_string_var;
```

Note also that unassigned variables are treated as empty strings for the purposes of concatenation, so $my_string_var will end up unchanged if $new_addition has never been given a value.

The heredoc syntax

In addition to the single-quote and double-quote syntaxes, PHP offers another way to specify a string, called the *heredoc syntax*. This syntax turns out to be extremely useful for specifying large chunks of variable-interpolated text, because it spares you from the need to escape internal quotation marks. It is especially useful in creating pages that contain HTML forms.

The operator in the heredoc syntax is <<<. What is expected immediately after this is a label (unquoted) that indicates the beginning of a multiline string. PHP will continue including subsequent lines in this string until it sees the same label again, beginning a line. The ending label may optionally be followed by a semicolon but by nothing else.

For example:

```
$my_string_var = <<<EOT
Everything in this rather unnecessarily wordy
ramble of prose will be incorporated into the
string that we are building up inevitably, inexorably,
character by character, line by line, until we reach that
blessed final line which is this one.
EOT;
```

Note that the preceding final EOT must *not* be indented at all — otherwise it will be taken to be just more text to be included. The label need not be literally EOT — it can be whatever you like within the normal rules for variable names in PHP.

Interpolation of variables happens exactly the same way as with double-quoted strings. The nice thing about heredoc, though, is that quote signs can be included without any escaping and without prematurely terminating the string. Here's another example:

```
echo <<<ENDOFFORM
<FORM METHOD=POST ACTION="{$_ENV['PHP_SELF']}">
<INPUT TYPE=TEXT NAME=FIRSTNAME VALUE=$firstname>
<INPUT TYPE=SUBMIT NAME=SUBMIT VALUE=SUBMIT>
</FORM>
ENDOFFORM;
```

This has the effect of echoing a very simple form to the browser.

String Functions

PHP gives you a huge variety of functions for the munching and crunching of strings. If you're ever tempted to roll your own function that reads strings character by character to produce a new string, pause for a moment to think whether the task might be common. If so, there is probably a built-in function that handles it.

For more information on string functions see `http://php.net/manual/en/ref.strings.php`.

In this section, we present the basic functions for inspecting, comparing, modifying, and printing strings. If you want to be really comfortable with string manipulation in PHP, you should probably have at least a passing acquaintance with everything in this section. Both the regular expression functions and the more abstruse string functions can be found in Chapter 22.

NOTE   **A note for C programmers: Many of the PHP string function names should be familiar to you. Just keep in mind that, because PHP takes care of memory management for you, the functions that return strings are allocating the string storage on their own and do not need to be given a preallocated string to write into.**

Inspecting strings

What kinds of questions can you ask strings? First on the list is how long the string is, using the `strlen()` function (the name is short for *string length*).

```
$short_string = "This string has 29 characters";
print("It does have " . strlen($short_string) .
         " characters");
```

This code gives the following output:

```
It does have 29 characters
```

Knowing the string's length is particularly useful in form validation or for situations in which we'd like to loop through a string character by character. A useless but illustrative example, using the preceding example string, is:

```
for ($index = 0; $index < strlen($short_string); $index++)
  print($short_string{$index});
```

This simply prints:

```
This string has 29 characters
```

which is the string we started with.

Finding characters and substrings

The next question you can ask your strings is what they contain. For example, the `strpos()` function finds the numerical position of a particular character in a string, if it exists.

```
$twister = "Peter Piper picked a peck of pickled peppers";
print("Location of 'p' is " . strpos($twister, 'p') .'<BR>');
print("Location of 'q' is "  . strpos($twister, 'q') .'<BR>');
```

This gives us the browser output:

```
Location of 'p' is 8
Location of 'q' is
```

The 'q' location is apparently blank because `strpos()` returns false if the character in question cannot be found, and a false value prints as the empty string. You should note that the `strpos()` function is case sensitive.

CAUTION   The `strpos()` **function is one of those cases where PHP's type-looseness can be problematic. If no match can be found, the function returns a false value; if the very first character is a match, the function returns** 0 **(because the indexing count starts with 0 rather than 1). Both of these values look false if used in a Boolean test. One way to distinguish them is to use the identity comparison operator (===, introduced as of PHP4), which is true only if its arguments are the same and of the same type — you can use it to test if the returned value is** 0 **(or is** FALSE**) without risk of confusion with other values that might be the same after type coercion.**

The `strpos()` function can also be used to search for a substring rather than a single character, simply by giving it a multicharacter string rather than a single-character string. You can also supply an extra integer argument specifying the position to begin searching forward from.

Searching in reverse is also possible, using the `strrpos()` function. (Note the extra *r*, which you can think of as standing for *reverse*.) This function takes a string to search and a single-character string to locate, and it returns the last position of occurrence of the second argument in the first argument. (Unlike with `strpos()`, the string searched for must have only one character.) If we use this function on our example sentence, we find a different position:

```
$twister = "Peter Piper picked a peck of pickled peppers";
printf("Location of 'p' is " . strrpos($twister, 'p') .'<BR>');
```

Specifically, we find the third *p* in `peppers`:

```
Location of 'p' is 40
```

Are strings immutable?

In some programming languages (such as C), it is common to manipulate strings by directly changing them — that is, storing new characters into the middle of an existing string, replacing old characters. Other languages try to keep the programmer out of certain kinds of trouble by making string classes that are *immutable* (or unchangeable) — you can make new strings by creating modified copies of old ones, but once you have made a string, you are not allowed to change it by directly changing the characters that make it up.

Where does PHP fit in? As it turns out, PHP strings can be changed, but the most common practice seems to be to treat strings as immutable.

Strings can be changed by treating them as character arrays and assigning directly into them, like this:

```
$my_string = "abcdefg";
$my_string[5] = "X";

print($my_string . "<BR>");
```

which will give the browser output:

```
abcdeXg
```

This modification method seems to be undocumented, however, and shows up nowhere in the online manual, even though the corresponding extraction method (now updated to use curly braces) is highlighted. Also, almost all PHP string-manipulation functions return modified copies of their string arguments rather than making direct changes, which seems to indicate that this is the style that the PHP designers prefer. Our advice is not to use this direct-modification method to change strings, unless you know what you are doing and there is some large benefit in terms of memory savings.

Comparison and searching

Is this string the same as that string? It's a question that your code is likely to have to answer frequently, especially when dealing with input typed by the end user.

> **NOTE**   For the == operator, two strings are the same if they contain exactly the same sequence of characters. It does not test any stricter notion of being the same, such as being stored at the same memory address, but it does pay attention to case (or capitalization).

The simplest method to find an answer is to use the basic comparison operator (==), which does equality testing on strings as well as numbers.

> **CAUTION**   Comparing two strings using == (or the corresponding < and > operators) is trustworthy if both the arguments are strings and if you know that no type conversion is being performed. (See Chapter 4 for more on this.) Using strcmp() (described next) is always trustworthy.

The most basic workhorse string-comparison function is strcmp(). It takes two strings as arguments and compares them byte by byte until it finds a difference. It returns a negative number if the first string is less than the second and a positive number if the second string is less. It returns 0 if they are identical.

The strcasecmp() function works the same way, except that the equality comparison is case insensitive. The function call strcasecmp("hey!", "HEY!") should return 0.

Searching

The comparison functions just described tell you whether one string is equal to another. To find out if one string is contained within another, use the strpos() function (covered earlier) or the strstr() function (or one of its relatives).

The strstr() function takes a string to search in and a string to look for (in that order). If it succeeds, it returns the portion of the string that starts with (and includes) the first instance of the string it is looking for. If the string is not found, a false value is returned. Here is a successful search followed by an unsuccessful search:

```
$string_to_search = "showsuponceshowsuptwice";
$string_to_find = "up";
print("Result of looking for $string_to_find" .
        strstr($string_to_search, $string_to_find) . "<br>");
$string_to_find = "down";
print("Result of looking for $string_to_find" .
        strstr($string_to_search, $string_to_find));
```

which gives us:

```
Result of looking for up: uponceshowsuptwice
Result of looking for down:
```

The blank space after the colon in the second line is the result of trying to print a false value, which prints as the empty string. The strstr() function also has an alias by the name of strchr(). Other than the name, the two functions are identical. Just as with strcmp(), strstr() has a case-insensitive version, by the name of stristr(). (That i in the middle stands for *insensitive*.) It is identical to strstr() in every way, except that the comparison treats lowercase letters as indistinguishable from their uppercase counterparts. The string functions we have covered so far are summarized in Table 7-1.

TABLE 7-1

Simple Inspection, Comparison, and Searching Functions

Function	Behavior
strlen()	Takes a single string argument and returns its length as an integer.
strpos()	Takes two string arguments: a string to search, and the string being searched for. Returns the (0-based) position of the beginning of the first instance of the string if found and a false value otherwise. It also takes a third optional integer argument, specifying the position at which the search should begin.
strrpos()	Like strpos(), except that it searches backward from the end of the string, rather than forward from the beginning. The search string must only be one character long, and there is no optional position argument.
strcmp()	Takes two strings as arguments and returns 0 if the strings are exactly equivalent. If strcmp() encounters a difference, it returns a negative number if the first different byte is a smaller ASCII value in the first string, and a positive number if the smaller byte is found in the second string.
strcasecmp()	Identical to strcmp(), except that lowercase and uppercase versions of the same letter compare as equal.
strstr()	Searches its first string argument to see if its second string argument is contained in it. Returns the substring of the first string that starts with the first instance of the second argument, if any is found — otherwise, it returns false.
strchr()	Identical to strstr().
stristr()	Identical to strstr() except that the comparison is case independent.

Substring selection

Many of PHP's string functions have to do with slicing and dicing your strings. By *slicing,* we mean choosing a portion of a string; by *dicing,* we mean selectively modifying a string. Keep in mind that (most of the time) even dicing functions do not change the string you started out with. Usually, such functions return a modified copy, leaving the original argument intact.

The most basic way to choose a portion of a string is the `substr()` function, which returns a new string that is a subsequence of the old one. As arguments, it takes a string (that the substring will be selected from), an integer (the position at which the desired substring starts), and an optional third integer argument that is the length of the desired substring. If no third argument is given, the substring is assumed to continue until the end. (Remember that, as with all PHP arguments that deal with numerical string positions, the numbering starts with 0 rather than 1.)

For example, the statement:

```
echo(substr("Take what you need, and leave the rest behind",
        23));
```

prints the string `leave the rest behind`, whereas the statement:

```
echo(substr("Take what you need, and leave the rest behind",
        5, 13));
```

prints `what you need` — a 13-character string starting at (0-based) position 5.

Both the start-position argument and the length argument can be negative, and in each case the negativity has a different meaning. If the start position is negative, it means that the starting character is determined by counting backward from the end of the string, rather than forward from the beginning. (A start position of `-1` means start with the last character, `-2` means second to last, and so on.)

Now, you might expect that a negative length would similarly imply that the substring should be determined by counting backward from the start character rather than forward. This is not the case — it is always true that the character at the start position is the first character in the returned string (not the last). Instead, a negative-length argument means that the final character is determined by counting backward from the end rather than forward from the start position.

Here are some examples, with positive and negative arguments:

```
$alphabet_test = "abcdefghijklmnop";
print("3: " . substr($alphabet_test, 3) . "<BR>");
print("-3: " . substr($alphabet_test, -3) . "<BR>");
print("3, 5: " . substr($alphabet_test, 3, 5) . "<BR>");
print("3, -5: " . substr($alphabet_test, 3, -5) . "<BR>");
print("-3, -5: " . substr($alphabet_test, -3, -5) . "<BR>");
print("-3, 5: " . substr($alphabet_test, -3, 5) . "<BR>");
```

This gives us the output:

```
3: defghijklmnop
-3: nop
3, 5: defgh
3, -5: defghijk
-3, -5:
-3, 5: nop
```

Notice that there is an intimate relationship between the functions `substr()`, `strstr()`, and `strpos()`. The `substr()` function selects a substring by numerical position, `strstr()` selects a substring by its content, and `strpos()` finds the numerical position of a given substring. In the case where we're sure in advance that the string `$containing` has the string `$contained` as a substring, the expression:

```
strstr($containing, $contained)
```

should be equivalent to the code:

```
substr($containing, strpos($containing, $contained))
```

String cleanup functions

Although they are technically substring functions, just like the others in this chapter, the functions `chop()`, `ltrim()`, and `trim()` are really used for cleaning up untidy strings. They trim whitespace off the end, the beginning, and the beginning and end, respectively, of their single string argument. Some examples:

```
$original = "  More than meets the eye   ";
$chopped = chop($original);
$ltrimmed = ltrim($original);
$trimmed = trim($original);
print("The original is '$original'<BR>");
print("Its length is " . strlen($original) . "<BR>");
print("The chopped version is '$chopped'<BR>");
print("Its length is " . strlen($chopped) . "<BR>");
print("The ltrimmed version is '$ltrimmed'<BR>");
print("Its length is " . strlen($ltrimmed) . "<BR>");
print("The trimmed version is '$ltrimmed'<BR>");
print("Its length is " . strlen($trimmed) . "<BR>");
```

The result as viewed by a browser is:

```
The original is ' More than meets the eye '
Its length is 28
The chopped version is ' More than meets the eye'
Its length is 25
The ltrimmed version is 'More than meets the eye '
Its length is 26
The trimmed version is 'More than meets the eye'
Its length is 23
```

The original string had three spaces at the end (subject to removal by `chop()` or `trim()`) and two at the beginning (removed by `ltrim()` and `trim()`). We were careful to describe our result as viewed by a browser because the multiple spaces have apparently been collapsed to one in the output, as browsers will do. If we viewed the HTML source produced by PHP originally, we would still see sequences of two and three spaces.

123

In addition to spaces, these functions remove whitespace like that denoted by the escape sequences \n, \r, \t, and \0 (end-of-line characters, tabs, and the null character used to terminate strings in C programs).

You will hear the name chop() more frequently, but the identical function can also be called with the more logical name of rtrim(). Finally, notice that although chop() sounds extremely destructive, it does not harm the $original argument, which retains the same value.

String replacement

The substring functions we've seen so far are all about choosing a portion of the argument rather than building a genuinely new string. Enter the functions str_replace() and substr_replace().

The str_replace() function enables you to replace all instances of a particular substring with an alternate string. It takes three arguments: the string to be searched for, the string to replace it with when it is found, and the string to perform the replacement on. For example:

```
$first_edition =
    "Burma is similar to Rhodesia in at least one way.";
$second_edition = str_replace("Rhodesia", "Zimbabwe",
                               $first_edition);
$third_edition = str_replace("Burma", "Myanmar",
                               $second_edition);
print($third_edition);
```

gives us:

```
Myanmar is similar to Zimbabwe in at least one way.
```

This replacement will happen for all instances found of the search string. If our outdated encyclopedia could be snarfed into a single PHP string, we could update it in one pass.

One subtlety to be aware of: What happens when multiple instances of the search string overlap? For example, with code like:

```
$tricky_string = "ABA is part of ABABA";
$maybe_tricked = str_replace("ABA", "DEF", $tricky_string);
print("Substitution result is '$maybe_tricked'<BR>");
```

the behavior we see is:

```
Substitution result is 'DEF is part of DEFBA'
```

which is probably as reasonable as any other alternative.

As you've seen, str_replace() picks out portions to replace by matching to a target string; by contrast, substr_replace() chooses a portion to replace by its absolute position. The function takes up to four arguments: the string to perform the replacement on, the string to replace it with,

the starting position for the replacement, and (optionally) the length of the section to be replaced. For example:

```
print(substr_replace("ABCDEFG", "-", 2, 3));
```

gives us:

```
AB-FG
```

The `CDE` portion of the string has been replaced with the single `-`. Notice that you are allowed to replace a substring with a string of a different length. If the length argument is omitted, it is assumed that you want to replace the entire portion of the string after the start position.

The `substr_replace()` function also takes negative arguments for starting position and length, which are treated exactly the same way as in the `substr()` function (described in the earlier section "Substring selection"). It is important to remember with both `str_replace` and `substr_replace` that the original string remains unchanged by these operations.

Finally, we have a couple more whimsical functions that produce new strings from old. The `strrev()` function simply returns a new string with the characters of its input in reverse order. The `str_repeat()` function takes a string argument and an integer argument and returns a string that is the appropriate number of copies of the string argument tacked together. For example:

```
print(str_repeat("cheers ", 3));
```

gives us:

```
cheers cheers cheers
```

for the end of this section at long last.

The substring search and replacement functions are summarized in Table 7-2.

Substring and String Replacement Functions

Function	Behavior
substr()	Returns a subsequence of its initial string argument, as specified by the second (position) argument and optional third (length) argument. The substring starts at the indicated position and continues for as many characters as specified by the length argument or until the end of the string, if there is no length argument.
	A negative position argument means that the start character is located by counting backward from the end, whereas a negative length argument means that the end of the substring is found by counting back from the end, rather than forward from the start position.

continued

TABLE 7-2	*(continued)*

<div align="center">

Substring and String Replacement Functions
</div>

Function	Behavior
chop(), or rtrim()	Returns its string argument with trailing (right-hand side) whitespace removed. Whitespace is a blank space, \n, \r, \t, and \0.
ltrim()	Returns its string argument with leading (left-hand side) whitespace removed.
Trim()	Returns its string argument with both leading and trailing whitespace removed.
Str_ replace()	Used to replace target substrings with another string. Takes three string arguments: a substring to search for, a string to replace it with, and the containing string. Returns a copy of the containing string with *all* instances of the first argument replaced by the second argument.
Substr_ replace()	Puts a string argument in place of a position-specified substring. Takes up to four arguments: the string to operate on, the string to replace with, the start position of the substring to replace, and the length of the string segment to be replaced. Returns a copy of the first argument with the replacement string put in place of the specified substring.
	If the length argument is omitted, the entire tail of the first string argument is replaced. Negative position and length arguments are treated as in substr().

Case functions

These functions change lowercase to uppercase and vice versa. The first two (de)capitalize entire strings, whereas the second two operate only on first letters of words.

strtolower()

The strtolower() function returns an all-lowercase string. It doesn't matter if the original is all uppercase or mixed. This fragment:

```php
<?php
$original = "They DON'T KnoW they're SHOUTING";
$lower = strtolower($original);
echo $lower;
?>
```

returns the string "they don't know they're shouting".

TIP   If you have been faced with extensive form-validation needs before, you might already have noticed that strtolower() is extremely handy for use with those that still think their e-mail addresses contain capital letters. Subsequent functions in this category will prove similarly useful.

strtoupper()

The `strtoupper()` function returns an all-uppercase string, regardless of whether the original was all lowercase or mixed:

```php
<?php
$original = "make this link stand out";
echo("<B>strtoupper($original)</B>");
?>
```

ucfirst()

The `ucfirst()` function capitalizes only the first letter of a string:

```php
<?php
$original = "polish is a word for which pronunciation depends on
capitalization";
echo(ucfirst($original));
?>
```

ucwords()

The `ucwords()` function capitalizes the first letter of each word in a string:

```php
<?php
$original = "truth or consequences";
$capitalized = ucwords($original);
echo "While $original is a parlor game, $capitalized is a town in New
Mexico.";
?>
```

NOTE   **Neither** ucwords() **nor** ucfirst() **converts anything into lowercase. Each makes only the appropriate leading letters into uppercase. If there are inappropriate capital letters in the middle of words, they will not be corrected.**

Escaping functions

One of the virtues of PHP is that it is willing to talk to almost anybody. In its role as a glue language, PHP talks to database servers, to LDAP servers, over sockets, and over the HTTP connection itself. Frequently, it accomplishes this communication by first constructing a message string (like a database query) and then shipping it off to the receiving program. Often, however, the program attaches special meanings to certain characters, which therefore have to be *escaped,* meaning that the receiving program is told to take them as a literal part of the string rather than treating them specially.

Many users deal with this issue by enabling magic-quotes, which ensures that quotation marks are escaped before strings are inserted into databases. If that's not feasible or desirable, there are good old-fashioned strip-slashing and add-slashing by hand. The `addslashes()` function

escapes quotation marks, double quotation marks, backslashes, and NULLs with backslashes, because these are the characters that typically need to be escaped for database queries.

```php
<?php
$escapedstring = addslashes("He said, 'I'm a dog.'");
$query = "INSERT INTO test (quote) values ('$escapedstring')";
$result = mysql_query($query) or die(mysql_error());
?>
```

This will prevent the SQL statement from thinking it's finished right before the letter I. When you pull the data back out, you'll need to use stripslashes() to get rid of the slashes.

```php
<?php
$query = "SELECT quote FROM test WHERE ID=1";
$result = mysql_query($query) or die(mysql_error());
$new_row = mysql_fetch_array($result);
$quote = stripslashes($new_row[0]);

echo $quote;
```

The quotemeta() function escapes a wider variety of characters, all of which usually have a special meaning in the Unix command line: '.', '\' '+', '*', '?', '[', '^', ']', '(', '$', and ')'. For example, the code:

```php
$literal_string =
    'These characters ($, *) are very special to me\n<BR>';
$qm_string = quotemeta($literal_string);
echo $qm_string;
```

will print:

```
These characters \(\$, \*\) are very special to me\\n
```

CROSS-REF   For escaping functions specific to HTML, see the "Advanced String Functions" section in Chapter 22.

Printing and output

The workhorse constructs for printing and output are print and echo, which we cover in detail in Chapter 4. The standard way to print the value of variables to output is to include them in a doubly quoted string (which will interpolate their values) and then give that string to print or echo.

If you need even more tightly formatted output, PHP also offers printf() and sprintf(), which are modeled on C functions of the same name. The two functions take identical arguments: a special format string (described later in this section) and then any number of other arguments, which will be spliced into the right places in the format string to make the result.

The only difference between printf() and sprintf() is that printf() sends the resulting string directly to output, whereas sprintf() returns the result string as its value.

NOTE **To C programmers:** **This** `sprintf()` **function is slightly different from C's version in that you need not supply an allocated string for** `sprintf()` **to write into — PHP allocates the result string for you.**

The complicated bit about these functions is the format string. Every character that you put in the string will show up literally in the result, except the % character and characters that immediately follow it. The % character signals the beginning of a *conversion specification*, which indicates how to print one of the arguments that follow the format string.

After the %, there are six elements that make up the conversion specification, some of which are optional: padding, alignment, minimum width, precision, and type.

- An optional *sign* character used for numbers to indicate whether the number will be negative (-).

- The single (optional) *padding* character is either a 0 or a space ( ). This character is used to fill any space that would otherwise be unused but that you have insisted (with the minimum width argument) be filled with something. If this padding character is not given, the default is to pad with spaces.

- The optional *alignment* character (-) indicates whether the printed value should be left- or right-justified. If present, the value will be left-justified; if absent, it will be right-justified.

- An optional *minimum width* number that indicates how many spaces this value should take up, at a minimum. (If more spaces are needed to print the value, it will overflow beyond its bounds.)

- An optional precision specifier is written as a dot (.) followed by a number. It indicates how many decimal points of *precision* a double should print with. (This has no effect on printing things other than doubles.)

- A single character indicating how the *type* of the value should be interpreted. The f character indicates printing as a double, the s character indicates printing as a string, and then the rest of the possible characters (b, c, d, o, x, X) mean that the value should be interpreted as an integer and printed in various formats. Those formats are b for binary, c for printing the character with the corresponding ASCII values, o for octal, x for hexadecimal (with lowercase letters) and X for hexadecimal with uppercase letters.

Here's an example of printing the same double in several different ways:

```
<pre>
<?php
$value = 3.14159;
printf("%f,%10f,%-010f,%2.2f\n",
        $value, $value, $value, $value);
?>
</pre>
```

gives us:

```
3.141590, 3.141590,3.141590000000000, 3.14
```

129

The `<pre></pre>` construct is HTML that tells the browser to format the enclosed block literally, without collapsing many spaces into one, and so on.

Summary

Strings are sequences of characters, and the string is one of the eight basic data types in PHP. Unlike in some other languages, there is no distinct character type, since single characters behave as strings of length 1. Literal strings are specified in code by either single (`'`) or double (`"`) quotation marks. Singly quoted strings are interpreted nearly literally, while doubly quoted strings interpret a number of escape sequences and automatically interpolate variable values.

The main string operator is `'.'`, which concatenates two strings together. In addition, there is a dizzying array of string functions, which help you inspect, compare, search, extract, chop, replace, slice, and dice strings to your heart's content. For the most sophisticated string-manipulation needs, PHP supports both POSIX and Perl-compatible regular expressions (covered in Chapter 22).

Chapter 8

Learning Arrays

Arrays are definitely one of the coolest and most flexible features of PHP. Unlike vector arrays from other languages (C, C++, Pascal), PHP arrays can store data of varied types and automatically organize it for you in a large variety of ways.

CROSS-REF This chapter treats arrays and array functions in some depth. For a very quick introduction to the syntax and use of arrays, see Chapter 4. For a more complete survey of advanced array functions, see Chapter 21.

The Uses of Arrays

An array is a collection of variables indexed and bundled into a single, easily referenced supervariable that offers an easy way to pass multiple values between lines of code, functions, and even pages. Throughout much of this chapter, we will be looking at the inner workings of arrays and exploring all the built-in PHP functions that manipulate them. Before we get too deep into that, however, it's worth listing the common ways that arrays are used in real PHP code.

Many built-in PHP environment variables are in the form of arrays (for example, $_SESSION, which contains all the variable names and values being propagated from page to page via PHP's session mechanism). If you want access to them, you need to understand, at a minimum, how to reference arrays.

Almost any situation that calls for a number of pieces of data to be packaged and handled as one is appropriate for a PHP array.

What Are PHP Arrays?

PHP arrays are *associative* arrays with a little extra machinery thrown in. The *associative* part means that arrays store element values in association with key values rather than in a strict linear index order. (If you have seen arrays in other programming languages, they are likely to have been *vector* arrays rather than associative arrays — see the related sidebar for an explanation of the difference.) If you store an element in an array, in association with a key, all you need to retrieve it later from that array is the key value. For example, storage is as simple as this:

```
$state_location['San Mateo'] = 'California';
```

which stores the element `'California'` in the array variable $state_location, in association with the lookup key `'San Mateo'`. After this has been stored, you can look up the stored value by using the key, like so:

```
$state = $state_location['San Mateo'];  // equals 'California'
```

Simple, no?

If all you want arrays for is to store key/value pairs, the preceding information is all you need to know. Similarly, if you want to associate a numerical ordering with a bunch of values, all you have to do is use integers as your key values, as in:

```
$my_array[1] = "The first thing";
$my_array[2] = "The second thing"; // and so on ...
```

NOTE   **For Perl programmers*: Arrays in PHP are much like hashes in Perl, with some syntactic differences. For one thing, all variables in PHP are denoted with a leading $, not just scalar variables. Second, even though the array is associative, the indices are grouped by square brackets ([ ]) rather than curly braces ({ }). Finally, there is no array or list type indexed only by integers. The convention is to use integers as associative indices, and the array itself maintains an internal ordering for iteration purposes.**

In addition to the machinery that makes this kind of key/value association possible, arrays track some other things behind the scenes. Because of this, we sometimes treat them as other kinds of data structures. As you will see, arrays can be *multidimensional*. They can store values in association with a sequence of key values rather than a single key. Also, arrays automatically maintain an ordered list of the elements that have been inserted in them, independent of what the key values happen to be. This makes it possible to treat arrays as linked lists. In general, we will reveal the workings of this extra machinery as we explore the functions that use it.

NOTE   **A note for C++ programmers*: You should be aware that arrays can handle some of the same tasks that require the use of template libraries in C++. Much of the reason for having templates in the first place is to get around restrictions having to do with strict typing of data. PHP's looser typing system makes it possible, for example, to write general algorithms that iterate over the contents of arrays without committing to the type of the array elements themselves.**

Associative Arrays versus Vector Arrays

If you have programmed in languages like C, C++, and Pascal, you are probably used to a particular usage of the word array, one that doesn't match the PHP usage very well at all. A more specific term for a C-style array is a vector array, whereas a PHP-style array is an associative array.

In a vector array, the contained elements all need to be of the same type, and usually the language compiler needs to know in advance how many such elements there are likely to be. For example, In C you might declare an array of 100 double-precision floating-point numbers with a statement like:

```
double my_array[100]; // This is C, not PHP!
```

The restriction on types and the advance declaration of size have an associated benefit: Vector arrays are very fast, both for storage and lookup. The reason is that the compiler will usually lay out the array in a contiguous block of computer memory, as large as the size of the element type multiplied by the number of elements. This makes it very easy for the programming language to locate a particular array slot — all it needs to know is the starting memory address of the array, the size of the element type, and the index of the element it wants to look up, and it can directly compute the memory address of that slot.

By contrast, PHP arrays are associative (and so some would call them *hashes,* rather than arrays). Rather than having a fixed number of slots, PHP creates array slots as new elements that are added to the array. Rather than requiring elements to be of the same type, PHP arrays have the same type-looseness that PHP variables have — you can assign arbitrary PHP values to be array elements. Finally, because vector arrays are all about laying out their elements in numerical order; the keys used for lookup and storage must be integer numbers. PHP arrays can have keys of arbitrary type, instead, including string keys. So, you could have successive array assignments like:

```
$my_array[1] = 1;
$my_array['orange'] = 2;
$my_array[3] = 3;
```

without any paradox. The result is that your array has three values (1, 2, 3), each of which is stored in association with a key (1, 'orange', and 3, respectively).

The extra flexibility of associative arrays comes at a price, because there is a little bit more going on between your code and the actual computation of a memory address than is true with vector arrays. For most web programming purposes, however, this extra access time is not a significant cost.

The fact that integers are legal keys for PHP arrays means that you can easily imitate the behavior of a vector array, simply by restricting your code to use only integers as keys.

> **NOTE** A general note for programmers familiar with other languages*:* PHP does not need very many different kinds of data structures, in part because of the great flexibility offered by PHP arrays. By careful choice of a subset of array functions, you can make arrays pretend to act like vector arrays, structure/record types, linked lists, hash tables, or stacks and queues — data structures that in other languages either require their own data types or less common language features such as pointers and explicit memory management.

Creating Arrays

There are three main ways to create an array in a PHP script: by assigning a value into one (and thereby implicitly creating it), by using the `array()` construct, and by calling a function that happens to return an array as its value.

Direct assignment

The simplest way to create an array is to act as though a variable is already an array and assign a value into it, like this:

```
$my_array[1] = "The first thing in my array that I just made";
```

If `$my_array` was an unbound variable (or bound to a nonarray variable) before this statement, it will now be a variable bound to an array with one element. If instead `$my_array` was already an array, the string will be stored in association with the integer key 1. If no value was associated with that number before, a new array slot will be created to hold it; if a value was associated with 1, the previous value will be overwritten. (You can also assign into an array by omitting the index entirely as in `$my_array[]`, described later in this chapter.)

The array() construct

The other way to create an array is via the `array()` construct, which creates a new array from the specification of its elements and associated keys. In its simplest version, `array()` is called with no arguments, which creates a new empty array. In its next simplest version, `array()` takes a comma-separated list of elements to be stored, without any specification of keys. The result is that the elements are stored in the array in the order specified and are assigned integer keys beginning with zero. For example, the statement:

```
$fruit_basket = array('apple', 'orange', 'banana', 'pear');
```

causes the variable `$fruit_basket` to be assigned to an array with four string elements (`'apple'`, `'banana'`, `'orange'`, `'pear'`), with the indices 0, 1, 2, and 3, respectively. In addition (as you'll see in the "Iteration" section later in this chapter), the array will remember the order in which the elements were stored.

The assignment to `$fruit_basket`, then, has exactly the same effect as the following:

```
$fruit_basket[0] = 'apple';
$fruit_basket[1] = 'orange';
$fruit_basket[2] = 'banana';
$fruit_basket[3] = 'pear';
```

assuming that the `$fruit_basket` variable was unbound at the first assignment. The same effect could also have been accomplished by omitting the indices in the assignment, like so:

```
$fruit_basket[] = 'apple';
$fruit_basket[] = 'orange';
```

```
$fruit_basket[] = 'banana';
$fruit_basket[] = 'pear';
```

In this case, PHP again assumes that you are adding sequential elements that should have numerical indices counting upward from zero.

> **NOTE** Yes, the default numbering for array indices starts at zero, not one. This is the convention for arrays in most programming languages. We're not sure why computer scientists start counting at zero (mathematicians, like everyone else in the world, start with one), but it probably has its origin in the kind of pointer arithmetic that calculates memory addresses for vector arrays. Addresses for successive elements of such arrays are found by adding successively larger offsets to the array's address, but the offset for the first element is zero (because the first element's address is the same as the array's address).

Specifying indices using array()

The simple example of `array()` in the preceding section assigns indices to our elements, but those indices will be the integers, counting upward from zero — we're not getting a lot of choice in the matter. As it turns out, `array()` offers us a special syntax for specifying what the indices should be. Instead of element values separated by commas, you supply key/value pairs separated by commas, where the key and value are separated by the special symbol =>.

Consider the following statement:

```
$fruit_basket = array(0 => 'apple', 1 => 'orange',
                      2 => 'banana', 3 => 'pear');
```

Evaluating it will have exactly the same effect as our earlier version — each string will be stored in the array in succession, with the indices 0, 1, 2, 3 in order. Instead, however, we can use exactly the same syntax to store these elements with different indices:

```
$fruit_basket = array('red' => 'apple', 'orange' => 'orange',
                      'yellow' => 'banana', 'green' => 'pear');
```

This gives us the same four elements, added to our new array in the same order, but indexed by color names rather than numbers. To recover the name of the yellow fruit, for example, we just evaluate the expression:

```
$fruit_basket['yellow'] // will be equal to 'banana'
```

Finally, as we said earlier, you can create an empty array by calling the `array` function with no arguments. For example:

```
$my_empty_array = array();
```

creates an array with no elements. This can be handy for passing to a function that expects an array as argument.

Functions returning arrays

The final way to create an array in a script is to call a function that returns an array. This may be a user-defined function, or it may be a built-in function that makes an array via methods internal to PHP.

Many database-interaction functions, for example, return their results in arrays that the functions create on the fly. Other functions exist simply to create arrays that are handy to have as grist for later array-manipulating functions. One such is `range()`, which takes two integers as arguments and returns an array filled with all the integers (inclusive) between the arguments. In other words:

```
$my_array = range(1,5);
```

is equivalent to:

```
$my_array = array(1, 2, 3, 4, 5);
```

Retrieving Values

After we have stored some values in an array, how do we get them out again?

Retrieving by index

The most direct way to retrieve a value is to use its index. If we have stored a value in `$my_array` at index 5, `$my_array[5]` should evaluate to the stored value. If `$my_array` has never been assigned, or if nothing has been stored in it with an index of 5, `$my_array[5]` will behave like an unbound variable.

The list() construct

There are a number of other ways to recover values from arrays without using keys, most of which exploit the fact that arrays are silently recording the order in which elements are stored. We cover this in more detail in this chapter's "Iteration" section, but one such example is `list()`, which is used to assign several array elements to variables in succession. Suppose that the following two statements are executed:

```
$fruit_basket = array('apple', 'orange', 'banana');
list($red_fruit, $orange_fruit) = $fruit_basket;
```

This will assign the string `'apple'` to the variable `$red_fruit` and the string `'orange'` to the variable `$orange_fruit` (with no assignment of `'banana'`, because we didn't supply enough variables). The variables in `list()` will be assigned to elements of the array in the order they were originally stored in the array. Notice the unusual behavior here — the `list()` construct is on the left-hand side of the assignment operator (=), where we normally find only variables.

In some sense, `list()` is the opposite or inverse of `array()` because `array()` packages its arguments into an array, and `list()` takes the array apart again into individual variable assignments. If we evaluate:

```
list($first, $second) = array($first, second);
```

the original values of `$first` and `$second` will be assigned to those variables again, after having been briefly stored in an array.

> **NOTE** We have been careful to refer to both `array()` and `list()` as constructs, rather than functions. This is because they are not in fact functions — like certain other specialized PHP language features (`if`, `while`, `function`, and so on) they are interpreted specially by the language itself and are not run through the usual routine of function-call interpretation. Remember that the arguments to a function call are evaluated before the function is really invoked on those arguments, so constructs that need to do other kinds of interpretation on what they are given cannot be implemented as function calls. It's a useful exercise to look hard at the example uses of both `array()` and `list()` to figure out why treating them as function calls could not result in the behavior advertised.

Multidimensional Arrays

So far, the array examples we have looked at have all been one-dimensional, with only one level of bracketed keys. However, PHP can easily support multidimensional arrays, with arbitrary numbers of keys. And just as with one-dimensional arrays, there is no need to declare our intentions in advance — the first reference to an array variable can be an assignment like:

```
$multi_array[1][2][3][4][5] = "deeply buried treasure";
```

That is a five-dimensional array with successive keys that happen, in this case, to be five successive integers.

Actually, in our opinion, thinking of arrays as multidimensional makes matters more confusing than they need to be. Instead, just remember that the values that are stored in arrays can themselves be arrays, just as legitimately as they can be strings or numbers. The multiple-index syntax in the preceding example is simply a concise way to refer to a (four-dimensional) array that is stored with a key of 1 in `$multi_array`, which in turn has a (three-dimensional) array stored in it, and so on. Note also that you can have different depths of reference in different parts of the array, like this:

```
$multi_level_array[0] = "a simple string";
$multi_level_array[1]['contains'] = "a string stored deeper";
```

The integer key of 0 stores a string, and the key of 1 stores an array that, in turn, has a string in it. However, you cannot continue on with this assignment:

```
$multi_level_array[0]['contains'] = "another deep string";
```

without the result of losing the first assignment to `'a simple string'`. The key of 0 can be used to store a string or another array, but not both at once.

137

If we remember that multidimensional arrays are simply arrays that have other arrays stored in them, it's easier to see how the `array()` creation construct generalizes. In fact, even this seemingly complicated assignment is not that complicated:

```
$cornucopia = array('fruit' =>
                 array('red' => 'apple',
                       'orange' => 'orange',
                       'yellow' => 'banana',
                       'green' => 'pear'),
                 'flower' =>
                   array('red' => 'rose',
                         'yellow' => 'sunflower',
                         'purple' => 'iris'));
```

It is simply an array with two values stored in association with keys. Each of these values is an array itself. After we have made the array, we can reference it like this:

```
$kind_wanted = 'flower';
$color_wanted = 'purple';
print("The $color_wanted $kind_wanted is " .
      $cornucopia[$kind_wanted][$color_wanted]);
```

See the browser output:

```
The purple flower is iris
```

NOTE **There's a reason that we used the string concatenation operator, `.`, in the preceding print statement, rather than simply embedding the** `$cornucopia[$kind_wanted]` `[$color_wanted]` **in our print string as we do with other variables. PHP3 string parsing can be confused by multiple array indices within a double-quoted string, so it needs to be concatenated separately. PHP since version 4 handles this in a better way — you are safe embedding array references in a string as long as you enclose the reference in curly braces, like this:**

```
print( "The thing we want is
{$cornucopia[$kind_wanted][$color_wanted]}");
```

Finally, notice that there is no great penalty for misindexing into a multidimensional array when we are trying to retrieve something; if no such key is found, the expression is treated like an unbound variable. So, if we try the following instead:

```
$kind_wanted = 'fruit';
$color_wanted = 'purple'; //uh-oh, we didn't store any plums
print("The $color_wanted $kind_wanted is " .
      $cornucopia[$kind_wanted][$color_wanted]);
```

The worst that happens is the unsatisfying:

```
The purple fruit is
```

This is the worst thing that happens, of course, unless you have raised your error_reporting level to `E_ALL`, as we advise you to do at some points in this book. In that case, you will get a notice message about an undefined index (`'purple'`) just as you would if you had an unbound variable.

Inspecting Arrays

Now we can make arrays, store values in arrays, and then pull the values out again when we want them. Table 8-1 summarizes a few other functions we can use to ask questions of our arrays.

TABLE 8-1

Simple Functions for Inspecting Arrays

Function	Behavior
`is_array()`	Takes a single argument of any type and returns a true value if the argument is an array, and false otherwise.
`count()`	Takes an array as argument and returns the number of nonempty elements in the array. (This will be 1 for strings and numbers.)
`sizeof()`	Identical to `count()`.
`in_array()`	Takes two arguments: an element (that might be a value in an array), and an array (that might contain the element). Returns `true` if the element is contained as a value in the array, `false` otherwise. (Note that this does not test for the presence of keys in the array.)
`isset($array[$key])`	Takes an *array[key]* form and returns `true` if the key portion is a valid key for the array. (This is a specific use of the more general function `isset()`, which tests whether a variable is bound.)

Note that all of these functions work on only the depth of the array specified, so that testing for values layers deep in a multidimensional array requires that you specify out that number of places. In the case of our preceding `$cornucopia` example, for instance:

```
count($cornucopia); // what do you expect here? 2? 7? 9?
```

returns a 2, while

```
count($cornucopia[fruit]);
```

returns 4.

Deleting from Arrays

Deleting an element from an array is simple, exactly analogous to getting rid of an assigned variable. Just call `unset()`, as in the following:

```
$my_array[0] = 'wanted';
$my_array[1] = 'unwanted';
$my_array[2] = 'wanted again';
unset($my_array[1]);
```

Assuming that `$my_array` was unbound when we started, at the end it has two values (`'wanted'`, `'wanted again'`), in association with two keys (0 and 2, respectively). It is as though we had skipped the original `'unwanted'` assignment (except that the keys are numbered differently).

Note that this is *not* the same as setting the contents to an empty value. If, instead of calling `unset()`, we had the following statement:

```
$my_array[1] = '';
```

at the end we would have three stored values (`'wanted'`, `''`, `'wanted again'`) in association with three keys (0, 1, and 2, respectively).

Iteration

We've seen how to put things into arrays, how to find them once we have put them there, and how to delete them when we don't want them anymore. What we need next is a technique for dealing with array elements in bulk. Iteration constructs help us do this by letting us step or loop through arrays, element by element or key by key.

We'll first delve briefly into the internal representation of arrays to understand how PHP supports iteration. (Although important, this subsection is skippable — if you want to use it but don't want to know how it works, you can jump down to the section titled "Using iteration functions.")

Support for iteration

In addition to storing values in association with their keys, PHP arrays silently build an ordered list of the key/value pairs that are stored, in the order that they are stored. The reason for this is to support operations that iterate over the entire contents of an array. (Notice that this is difficult to do simply by building a loop that increments an index, because array indices are not necessarily numerical.)

There is, in fact, sort of a hidden pointer system built into arrays. Each stored key/value pair points to the next one, and one side effect of adding the first element to an array is that a current pointer points to the very first element, where it will stay unless disturbed by one of the iteration functions.

NOTE **Each array remembers a particular stored key/value pair as being the current one, and array iteration functions work in part by shifting that current marker through the internal list of keys and values. Although we will call this marker the *current pointer*, PHP does not support full pointers in the sense that C and C++ programmers may be used to, and this usage of the word will turn up only in the context of iterating through arrays.**

This linked-list pointer system is an alternative way to inspect and manipulate arrays, which exists alongside the system that allows key-based lookup and storage. Figure 8-1 shows an abstract view (not necessarily reflecting the real implementation) of how these systems locate elements in an array.

FIGURE 8-1

Internal structure of an array
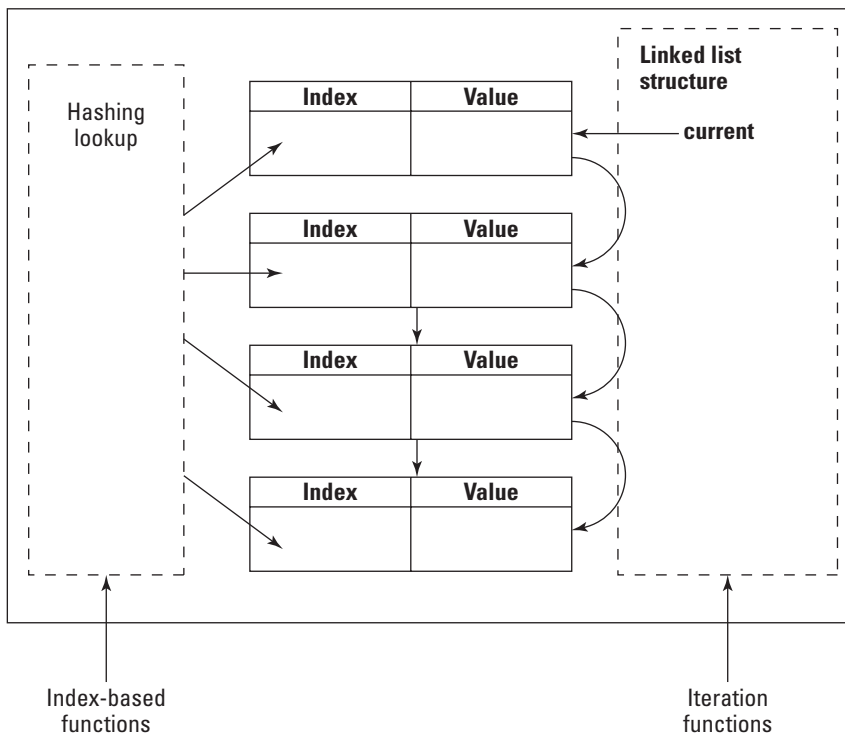
Using iteration functions

To explore the iteration functions, let's construct a sample array that we can iterate over.

```
$major_city_info = array();
$major_city_info[0] = 'Chicago';
```

```
$major_city_info['Chicago'] = 'United States';
$major_city_info[1] = 'Stockholm';
$major_city_info['Stockholm'] = 'Sweden';
$major_city_info[2] = 'Montreal';
$major_city_info['Montreal'] = 'Canada';
```

In this example, we created an array and stored some names of cities in it, in association with numerical indices. We also stored the names of the relevant countries into the array, indexed by the city names. (We could have accomplished all this with one big call to `array()`, but the separate statements make the structure of the array somewhat clearer.)

Now, we can use the array key system to pull out the data we have stored. If we want to rely on the convention in the preceding example (cities stored with numerical indices, countries stored with city-name indices), we can write a function that prints the city and the associated country, like this:

```
function city_by_number ($number_index, $city_array)
{
  if (IsSet($city_array[$number_index]))
    {
      $the_city = $city_array[$number_index];
      $the_country = $city_array[$the_city];
      print("$the_city is in $the_country<BR>");
    }
}
city_by_number(0, $major_city_info);
city_by_number(1, $major_city_info);
city_by_number(2, $major_city_info);
```

If we have set `$major_city`, as in the previous block of code, the browser output we should expect is:

```
Chicago is in United States
Stockholm is in Sweden
Montreal is in Canada
```

Now, this method of retrieval is fine when we know how the array is structured and we know what all the keys are, but what if you would simply like to print everything that an array contains?

Our favorite iteration method: foreach

Our favorite construct for looping through an array is `foreach`. Although it is probably inherited from Perl's `foreach`, it has a somewhat odd syntax (which is not the same as Perl's odd syntax). It comes in two flavors — which one you decide to use will depend on whether you care about the array's keys or just the values.

```
foreach ($array_variable as $value_variable) {
  // .. do something with the value in $value_variable
}  // Note that this is an example template, not real PHP code
```

```
foreach ($array_variable as $key_var => $value_var) {
// .. do something with $key_var and/or $value_var
}
```

Although in the preceding pseudocode we assume that the array of interest is in the variable $array_variable, you can have any expression that evaluates to an array in that position, for example:

```
foreach (function_returning_array() as $value_variable) {
  // .. do something with the value in $value_variable
}
```

> **NOTE** Like array() and list(), but unlike the genuine iteration functions in the rest of this section, foreach is a language construct, not a function. (See the earlier note about list() for an explanation of the difference.)

As an example, let's write a function to print all the names from our sample array:

```
function print_all_foreach ($city_array)
{
  foreach ($city_array as $name_value) {
    print("$name_value<BR>");
  }
}
print_all_foreach($major_city_info);
print_all_foreach($major_city_info);// again, as an experiment
```

As output, we get all the names, in the order we stored them, twice over:

```
Chicago
United States
Stockholm
Sweden
Montreal
Canada
Chicago
United States
Stockholm
Sweden
Montreal
Canada
```

We printed the contents twice to show that calling the function is repeatable.

Iterating with current() and next()

We like foreach, but it is really only good for situations where you want to simply loop through an array's values. For more control, let's look at current() and next().

The current() function returns the stored value that the current pointer points to. (Refer back to Figure 8-1 for a diagram of the array internals.) When an array is newly created with elements, the element pointed to will always be the first element. The next() function first advances that pointer and then returns the current value pointed to. If the next() function is called when the current pointer is already pointing to the last stored value and, therefore, runs off the end of the array, the function returns a false value.

As an example, we can print out an array's contents with the iteration functions current() and next(). (Notice that the final function call is repeated.)

```
function print_all_next($city_array)
{  // warning--doesn't quite work. See the function each()
  $current_item = current($city_array);
  if ($current_item)
    print("$current_item<BR>");
  else
    print("There's nothing to print");
  while($current_item = next($city_array))
    print("$current_item<BR>");
}
print_all_next($major_city_info);
print_all_next($major_city_info);// again, to see what happens
```

NOTE   There is a gotcha lurking in the preceding code example, which doesn't bite us in this particular example but makes this function untrustworthy as a general method for finding everything in an array. The problem is that we may have stored a false value in the array, which our while loop won't be able to distinguish from the false value that next() returns when it has run out of array elements. See the discussion of the each() function later in this chapter under "Empty values and the each() function" for a solution.

When we execute this array-printing code, we get the following again:

```
Chicago
United States
Stockholm
Sweden
Montreal
Canada
Chicago
United States
Stockholm
Sweden
Montreal
Canada
```

Now, how is it that we are seeing the same thing from the second call to print_all_next()? How did the current pointer get back to the beginning to start all over again the second time? The answer

lies in the fact that PHP function calls are *call-by-value*, meaning that they copy their arguments rather than operating directly on them. Both of the function calls, then, are getting a fresh copy of their array argument, which has never itself been disturbed by a call to next().

We can test this explanation by passing the arrays by reference rather than by value. If we define the same function but call it with ampersands (&) like this:

```
print_all_next(&$major_city_info);
print_all_next(&$major_city_info);// again
```

We get the following printing behavior:

```
Chicago
United States
Stockholm
Sweden
Montreal
Canada
There's nothing to print
```

NOTE The trick we used to test the array behavior (passing a variable reference to a function) has been deprecated, so you may get a warning when running this code, in addition to seeing the results printed above.

The reason is that this time the current pointer of the global version of the array was moved by the first function call.

NOTE Most of the iteration functions have both a returned value and a side effect. In the case of the functions next(), prev(), reset(), and end(), the side effect is to change the position of the internal pointer, and what is returned is the value from the key/value pair pointed to *after* the pointer's position is changed.

Starting over with reset()

In the preceding section, we wrote a function intended to print out all the values in an array, and we saw how it could fail if the array's internal pointer did not start off at the beginning of the list of key/value pairs. The reset() function gives us a way to "rewind" that pointer to the beginning — it sets the pointer to the first key/value pair and then returns the stored value. We can use it to make our printing function more robust by replacing the call to current() with a call to reset().

```
function print_all_array_reset($city_array)
{  // warning--still not reliable. See the function each()
  $current_item = reset($city_array); //rewind, return value
  if ($current_item)
    print("$current_item<BR>");
```

```
  else
    print("There's nothing to print");
  while($current_item = next($city_array))
    print("$current_item<BR>");
}
```

This function is somewhat more predictable in that it will always start with the first element, regardless of the pointer's location in the array it is handed. (Whether this is a good idea depends, of course, on what the function is used for and whether its arguments are passed by value or by reference.)

Perhaps confusingly, we use our call to `reset()` in the preceding example both for its side effect (rewinding the pointer) and for its return value (the first value stored). Alternatively, we could replace the first real line of the function body with these two lines:

```
reset($city_array);  // rewind to the first element
$current_item = current($city_array); // the first value
```

Reverse order with end() and prev()

We have seen the functions `next()`, which moves the current pointer ahead by one, and `reset()`, which rewinds the pointer to the beginning. Analogously, there are also the functions `prev()`, which moves the pointer back by one, and `end()`, which jumps the pointer to the last entry in the list. We can use these, for example, to print our array entries in reverse order.

```
function print_all_array_backwards($city_array)
{  // warning--still not reliable. See the function each()
  $current_item = end($city_array); //fast-forward to last
  if ($current_item)
    print("$current_item<BR>");
  else
    print("There's nothing to print");
  while($current_item = prev($city_array))
    print("$current_item<BR>");
}
print_all_array_backwards($major_city_info);
```

If we call this on the same `$major_city_info` data as in previous examples, we get the same print-out in reverse order:

```
Canada
Montreal
Sweden
Stockholm
United States
Chicago
```

Extracting keys with key()

So far, we have printed only the values stored in arrays, even though we are storing keys as well. The keys are also retrievable from the internal linked list of an array by using the key() function — this acts just like current() except that it returns the key of a key/value pair, rather than the value. (Refer to Figure 8-1.) Using the key() function, we can modify one of our earlier printing functions to print keys as well as values.

```
function print_keys_and_values($city_array)
{   // warning--See the discussion of each() below
  reset($city_array);
  $current_value = current($city_array);
  $current_key = key($city_array);
  if ($current_value)
    print("Key: $current_key; Value: $current_value<BR>");
  else
    print("There's nothing to print");
  while($current_value = next($city_array))
  {
      $current_key = key($city_array);
      print("Key: $current_key; Value: $current_value<BR>");
  }
}
print_keys_and_values($major_city_info);
```

With the same data as before, this gives us the browser output:

```
Key: 0; Value: Chicago
Key: Chicago; Value: United States
Key: 1; Value: Stockholm
Key: Stockholm; Value: Sweden
Key: 2; Value: Montreal
Key: Montreal; Value: Canada
```

Empty values and the each() function

We have written several functions that print the contents of arrays by iterating through them and, as we have pointed out, all but the foreach version have the same weakness. Each one of them tests for completion by seeing whether next() returns a false value. This will reliably happen when the array runs out of values, but it will also happen if and when we encounter a false value that we have actually stored. False values include the empty string (""), the number 0, and the Boolean value FALSE, any or all of which we might reasonably store as a data value for some task or other.

To the rescue comes each(), which is somewhat similar to next() but has the virtue of returning false only after it has run out of array to traverse. Oddly enough, if it has not run out, each() returns an array itself, which holds both keys and values for the key/value pair it is pointing at. This

147

characteristic makes each() confusing to talk about because you need to keep two arrays straight: the array that you are traversing and the array that each() returns every time that it is called. The array that each() *returns* has the following four key/value pairs:

- Key: 0; Value: *current-key*

- Key: 1; Value: *current-value*

- Key: 'key'; Value: *current-key*

- Key: 'value'; Value: *current-value*

The *current-key* and *current-value* are the key and value from the array being traversed. In other words, the returned array packages up the current key/value pair from the traversed array and offers both numerical and string indices to specify whether you are interested in the key or the value.

> **NOTE**   **In addition to having a different type of return value,** each() **differs from** next() **in that** each() **returns the value that was pointed to** *before* **moving the current pointer ahead, whereas** next() **returns the value** *after* **the pointer is moved. This means that if you start with a current pointer pointing to the first element of an array, successive calls to** each() **will cover each array cell, whereas successive calls to** next() **will skip the first value.**

We can use each() to write a more robust version of a function to print all keys and values in an array:

```
function print_keys_and_values_each($city_array)
{  // reliably prints everything in array
  reset($city_array);
  while ($array_cell = each($city_array))
  {
    $current_value = $array_cell['value'];
    $current_key = $array_cell['key'];
    print("Key: $current_key; Value: $current_value<BR>");
  }
}
print_keys_and_values_each($major_city_info);
```

Applying this function to our standard sample array gives the following browser output:

```
Key: 0; Value: Chicago
Key: Chicago; Value: United States
Key: 1; Value: Stockholm
Key: Stockholm; Value: Sweden
Key: 2; Value: Montreal
Key: Montreal; Value: Canada
```

This is exactly the same as was produced by our earlier function `print_keys_and_values()`. The difference is that our new function will not stop prematurely if one of the values is false or empty.

Walking with array_walk()

Our last iteration function lets you pass an arbitrary function of your own design over an array, doing whatever your function pleases with each key/value pair. The `array_walk()` function takes two arguments: an array to be traversed and the name of a function to apply to each key/value pair. (It also takes an optional third argument, discussed later in this section.)

The function that is passed in to `array_walk()` should take two (or three) arguments. The first argument will be the value of the array cell that is visited, and the second argument will be the key of that cell. For example, here is a function that prints a descriptive statement about the string length of an array value:

```
function print_value_length($array_value, $array_key_ignored)
{
  $the_length = strlen($array_value);
  print("The length of $array_value is $the_length<BR>");
}
```

(Notice that this function intentionally does nothing with the second argument.) Now let's pass this function over our standard sample array using `array_walk()`:

```
array_walk($major_city_info, 'print_value_length');
```

which gives the browser output:

```
The length of Chicago is 7
The length of United States is 13
The length of Stockholm is 9
The length of Sweden is 6
The length of Montreal is 8
The length of Canada is 6
```

The final flexibility that `array_walk()` offers is accepting an optional third argument that, if present, will be passed on, in turn, as a third argument to the function that is applied. This argument will be the same throughout the array's traversal, but it offers an extra source of runtime control for the passed function's behavior.

CAUTION   **You should not alter an array while you are iterating through the array using** `array_walk()`**. There is no guarantee how** `array_walk()` **will behave if you do this.**

Table 8-2 shows a summary of the behavior of the array iteration functions that we covered in this section. Notice that `foreach` and `list` are not included; they are not functions.

TABLE 8-2

Functions for Iterating over arrays

Function	Arguments	Side Effect	Return Value
current()	One array argument	None.	The value from the key/value pair currently pointed to by the internal "current" pointer (or false if no such value).
next()	One array argument	Advances the pointer by one. If already at the last element, it will move the pointer "past the end," and subsequent calls to current() will return false.	The value pointed to after the pointer has been advanced (or false if no such value).
prev()	One array argument	Moves the pointer back by one. If already at the first element, will move the pointer "before the beginning."	The value pointed to after the pointer has been moved back (or false if no such value).
reset()	One array argument	Moves the pointer back to point to the first key/value pair, or "before the beginning" if the array is empty.	The first value stored in the array, or false for an empty array.
end()	One array argument	Moves the pointer ahead to the last key/value pair.	The last value that is currently in the list of key/value pairs.
pos()	One array argument	None. (This function is an alias for current().)	The value of the key/value pair that is currently pointed to.
each()	One array argument	Moves the pointer ahead to the next key/value pair.	An array that packages the keys and values of the key/value pair that was current before the pointer was moved (or false if no such pair). The returned array stores the key and value under its own keys 0 and 1, respectively, and also under its own keys 'key' and 'value'.

Function	Arguments	Side Effect	Return Value
`array_walk()`	1) An array argument, 2) the name of a two- (or three-) argument function to call on each key/value, and 3) an optional third argument.	This function invokes the function named by its second argument on each key/value pair. Side effects depend on the side effects of the passed function.	(Returns 1.)

Summary

The array is a basic PHP data type and plays the role of both record types and vector array types in other languages. PHP arrays are associative, meaning that they store their values in association with unique *keys* or *indices*. Indices can be either strings or numbers, and are denoted as indices by square brackets. (The expression `$my_array[4]` refers to the value stored in `$my_array` in association with the integer index 4, and not necessarily to the 4th element of `$my_array`.)

The loose typing of PHP means that any PHP value can be stored as an array. In turn, this means that arrays can be stored as array elements. *Multidimensional arrays* are simply arrays that contain other arrays as elements, with a reference syntax of successive brackets. (The expression `$my_array[3][4]` refers to the element (indexed by 4) of an array that is an element [indexed by 3] of `$my_array`.)

The array is the standard vehicle for PHP functions that return structured data, so PHP programmers should learn to unpack arrays, even if they are not interested in constructing them. PHP also offers a huge variety of functions for manipulating data after you have it stored in an array, including functions for counting, summarizing, and sorting.

Chapter 9

Learning PHP Number Handling

I f you need to do serious numerical, scientific, or statistical computa-tion, a web-scripting language is probably not where you want to be doing it. With that said, however, PHP does offer a generous array of functions that nicely cover most of the mathematical tasks that arise in web scripting. It also offers some more advanced capabilities such as arbitrary-precision arithmetic and access to hashing and cryptographic libraries.

The PHP designers have, quite sensibly, not tried to reinvent any wheels in this department. Instead, they found about 18 perfectly good wheels by the side of the road and built a lightweight fiberglass chassis to connect them all together. Many of the more basic math functions in PHP are simple wrappers around their C counterparts (for more on this, see the sidebar "A Glimpse behind the Curtain" in Chapter 27, which will cover PHP's math-ematics capabilities in greater detail).

Numerical Types

PHP has only two numerical types: *integer* (also known as *long*), and *double* (aka *float*), which correspond to the largest numerical types in the C lan-guage. PHP does automatic conversion of numerical types, so they can be freely intermixed in numerical expressions, and the "right thing" will typi-cally happen. PHP also converts strings to numbers where necessary.

In situations where you want a value to be interpreted as a particular numerical type, you can force a typecast by prepending the type in parentheses, such as:

```
(double) $my_var
(integer) $my_var
```

Or you can use the functions `intval()` and `doubleval()`, which convert their arguments to integers and doubles, respectively.

CROSS-REF   **For more details on the integer and double types, see Chapter 4.**

Mathematical Operators

Most of the mathematical action in PHP is in the form of built-in functions rather than in the form of operators. In addition to the comparison operators covered in Chapter 5, PHP offers five operators for simple arithmetic, as well as some *shorthand* operators that make incrementing and assigning statements more concise.

Arithmetic operators

The five basic arithmetic operators are those you would find on a four-function calculator, plus the modulus operator (%). (If you are unfamiliar with modulus, see the discussion following Table 9-1.) The operators are summarized in Table 9-1.

TABLE 9-1

Arithmetic Operators

Operator	Behavior	Examples
+	Sum of its two arguments.	`4 + 9.5` evaluates to `13.5`
-	If there are two arguments, the right-hand argument is subtracted from the left-hand argument. If there is just a right-hand argument, then the negative of that argument is returned.	`50 - 75` evaluates to `-25` `- 3.9` evaluates to `-3.9`
*	Product of its two arguments.	`3.14 * 2` evaluates to `6.28`
/	Floating-point division of the left-hand argument by the right-hand argument.	`5 / 2` evaluates to `2.5`

%	Integer remainder from division of left-hand argument by the absolute value of the right-hand argument. (See discussion in the following section.)	`101 % 50` evaluates to `1`
		`999 % 3` evaluates to `0`
		`43 % 94` evaluates to `43`
		`-12 % 10` evaluates to `-2`
		`-12 % -10` evaluates to `-2`

154

Arithmetic operators and types

With the first three arithmetic operators (+, -, *), you should expect *type contagion* from doubles to integers; that is, if both arguments are integers, the result will be an integer, but if either argument is a double, then the result will be a double. With the division operator, there is the same sort of contagion, and in addition the result will be a double if the division is not even.

> **TIP** If you want integer division rather than floating-point division, simply coerce or convert the division result to an integer. For example, `intval(5 / 2)` evaluates to the integer 2.

Modular arithmetic is sometimes taught in school as *clock arithmetic*. The process of taking one number modulo to another amounts to "wrapping" the first number around the second, or (equivalently) taking the remainder of the first number after dividing by the second. The result of such an operation is always less than the second number.

Roughly speaking, a conventional civilian analog clock displays hours elapsed modulo 12, while military time is modulo 24. (The *roughly* in the previous sentence is because the real modulus function converts numbers to the range 0 to *n*-1, rather than the range 1 to *n*. If bell-tower clocks respected this, noontime would be marked by silence, rather than by 12 chimes.)

The modulus operator in PHP (%) expects integer arguments — if it is given doubles, they will simply be converted to integers (by truncation) first. The result is always an integer.

Most programming languages have some form of the modulus operator, but they differ in how they handle negative arguments. In some languages, the result of the operator is always positive, and `-2 % 26` equals 24. In PHP, though, `-2 % 26` is `-2`, and, in general, the statement `$mod = $first_num % $second_num` is exactly equivalent to the expression:

```
if ($first_num >= 0)
    $mod = $first_num % abs($second_num);
else
    $mod = - (abs($first_num) % abs($second_num));
```

where `abs()` is the absolute value function.

Incrementing operators

PHP inherits a lot of its syntax from C, and C programmers are famously proud of their own conciseness. The incrementing/decrementing operators taken from C make it possible to more concisely represent statements like `$count = $count + 1`, which tend to be typed frequently.

The increment operator (++) adds one to the variable it is attached to, and the decrement operator (--) subtracts one from the variable. Each one comes in two flavors, *postincrement* (which is placed immediately after the affected variable), and *preincrement* (which comes immediately before). Both flavors have the same side effect of changing the variable's value, but they have different values as expressions. The postincrement operator acts as if it changes the variable's value after the expression's value is returned, whereas the preincrement operator acts as though it makes the change

first and then returns the variable's new value. You can see the difference by using the operators in assignment statements, like this:

```
$count = 0;
$result = $count++;
print("Post ++: count is $count, result is $result<BR>");
$count = 0;
$result = ++$count;
print("Pre ++: count is $count, result is $result<BR>");
$count = 0;
$result = $count--;
print("Post --: count is $count, result is $result<BR>");
$count = 0;
$result = --$count;
print("Pre --: count is $count, result is $result<BR>");
```

which gives the browser output:

```
Post ++: count is 1, result is 0
Pre ++: count is 1, result is 1
Post --: count is -1, result is 0
Pre --: count is -1, result is -1
```

In this example, the statement $result = $count++; is exactly equivalent to:

```
$result = $count;
$count = $count + 1;
```

while $result = ++$count; is equivalent to:

```
$count = $count + 1;
$result = $count;
```

Assignment operators

Incrementing operators like ++ save keystrokes when adding one to a variable, but they don't help when adding another number or performing another kind of arithmetic. Luckily, all five arithmetic operators have corresponding assignment operators (+=, -=, *=, /=, and %=) that assign to a variable the result of an arithmetic operation on that variable in one fell swoop. The statement:

```
$count = $count * 3;
```

can be shortened to:

```
$count *= 3;
```

and the statement:

```
$count = $count + 17;
```

becomes:

```
$count += 17;
```

Comparison operators

PHP includes the standard arithmetic comparison operators, which take simple values (numbers or strings) as arguments and evaluate to either TRUE or FALSE:

CROSS-REF **For examples of using the comparison operators and also some gotcha issues with comparing doubles and strings, see Chapter 5.**

- The < (less than) operator is true if its left-hand argument is strictly less than its right-hand argument but false otherwise.

- The > (greater than) operator is true if its left-hand argument is strictly greater than its right-hand argument but false otherwise.

- The <= (less than or equal) operator is true if its left-hand argument is less than or equal to its right-hand argument but false otherwise.

- The >= (greater than or equal) operator is true if its left-hand argument is greater than or equal to its right-hand argument but false otherwise.

- The == (equal to) operator is true if its arguments are exactly equal but false otherwise.

- The != (not equal) operator is false if its arguments are exactly equal and true otherwise. This operator is the same as <>.

- The === operator (identical to) is true if its two arguments are exactly equal and of the same type.

- The !== operator (not identical to) is true if the two arguments are not equal or not of the same type.

TIP **The** identical to **operator (===)can, at times, be a necessary antidote to PHP's automatic type conversions. None of the following expressions will have a** true **value:**

```
2 === 2.0
2 === "2"
"2.0" === 2.0
0 === FALSE
```

This behavior can be invaluable, for example, if you have a function that returns a string when it succeeds (which might be the empty string) and a FALSE **value when it fails. Testing the truth of the return value would confuse** FALSE **with the empty string, whereas the identical operator can distinguish them.**

Precedence and parentheses

Operator *precedence* rules govern the relative stickiness of operators, deciding which operators in an expression get first claim on the arguments that surround them. You can find a complete table of all operator precedences in the manual at `www.php.net`, but the important precedence rules for arithmetic are:

- Arithmetic operators have higher precedence (that is, bind more tightly) than comparison operators.
- Comparison operators have higher precedence than assignment operators.
- The *, /, and % arithmetic operators have the same precedence.
- The + and – arithmetic operators have the same precedence.
- The *, /, and % operators have higher precedence than + and –.
- When arithmetic operators are of the same precedence, associativity is from left to right (that is, a number will associate with an operator to its left in preference to the operator on its right).

If you find the precedence rules difficult to remember, the next person who reads your code may have the same problem, so feel free to parenthesize when in doubt. For example, can you easily figure out the value of this expression?

```
1 + 2 * 3 - 4 - 5 / 4 % 3
```

As it turns out, the value is 2, as you can see more easily when we add parentheses that are not, strictly speaking, necessary:

```
((1 + (2 * 3)) - 4) - ((5 / 4) % 3)
```

Simple Mathematical Functions

The next step up in sophistication from the arithmetic operators consists of miscellaneous functions that perform tasks like converting between the two numerical types (which we discussed in Chapter 4) and finding the minimum and maximum of a set of numbers (see Table 9-2).

For example, the result of the following expression:

```
min(3, abs(-3), max(round(2.7), ceil(2.3), floor(3.9)))
```

is 3, because the value of every function call is also 3.

TABLE 9-2

Simple Math Functions

Function	Behavior
floor()	Takes a single argument (typically a double) and returns the largest integer that is less than or equal to that argument.
ceil()	Short for ceiling — takes a single argument (typically a double) and returns the smallest integer that is greater than or equal to that argument.
round()	Takes a single argument (typically a double) and returns the nearest integer. If the fractional part is exactly 0.5, it returns the nearest even number.
abs()	Short for absolute value — if the single numerical argument is negative, the corresponding positive number is returned; if the argument is positive, the argument itself is returned.
min()	Takes any number of numerical arguments (but at least one) and returns the smallest of the arguments.
max()	Takes any number of numerical arguments (but at least one) and returns the largest of the arguments.

Randomness

PHP's functions for generating pseudo-random numbers are summarized in Table 9-3. (If you are new to random number generation and are wondering what the *pseudo* is all about, please see the accompanying sidebar.)

There are two random number generators (invoked with rand() and mt_rand(), respectively), each with the same three associated functions: a seeding function, the random number function itself, and a function that retrieves the largest integer that might be returned by the generator.

The particular pseudo-random function that is used by rand() may depend on the particular libraries that PHP was compiled with. By contrast, the mt_rand() generator always uses the same random function (the Mersenne Twister), and the author of mt_rand()'s online documentation argues that it is also faster and "more random" (in a cryptographic sense) than rand(). We have no reason to believe that this is not correct, so we prefer mt_rand() to rand().

TABLE 9-3

Random Number Functions

Function	Behavior
`srand()`	Takes a single positive integer argument and seeds the random number generator with it.
`rand()`	If called with no arguments, returns a "random" number between 0 and `RAND_MAX` (which can be retrieved with the function `getrandmax()`). The function can also be called with two integer arguments to restrict the range of the number returned — the first argument is the minimum and the second is the maximum (inclusive).
`getrandmax()`	Returns the largest number that may be returned by `rand()`. This number is limited to 32768 on Windows platforms.
`mt_srand()`	Like `srand()`, except that it seeds the "better" random number generator.
`mt_rand()`	Like `rand()`, except that it uses the "better" random number generator.
`mt_getrandmax()`	Returns the largest number that may be returned by `mt_rand()`.

NOTE   **On some PHP versions and some platforms, you can apparently get seemingly random numbers from** `rand()` **and** `mt_rand()` **without seeding first — this should not be relied upon, however, both for reasons of portability and because the unseeded behavior is not guaranteed.**

Seeding the generator

The typical way to seed either of the PHP random number generators (using `mt_srand()` or `srand()`) looks like this:

```
mt_srand((double)microtime()*1000000);
```

This sets the seed of the generator to be the number of microseconds that have elapsed since the last whole second. (Yes, the typecast to `double` is necessary here, because `microtime()` returns a string, which would treated as an integer in the multiplication but for the cast.) Please use this seeding statement even if you don't understand it — just place it in any PHP page, once only, before you use the corresponding `mt_rand()` or `rand()` functions, and it will ensure that you have a varying starting point and therefore random sequences that are different every time. This particular seeding technique has been thought through by people who understand the ins and outs of pseudo-random number generation and is probably better than any attempt an individual programmer might make to try something trickier.

TIP   **Although the random number functions only return integers, it is easy to convert a random integer in a given range to a corresponding floating-point number (say, one between 0.0 and 1.0 inclusive) with an expression like** `rand() / getrandmax()`**. You can then scale and shift the range as desired (to, say, a number between 100.0 and 120.0) with an expression like** `100.0 + 20.0 * (rand() / getrandmax())`**.**

Pseudo-Random Number Generators

As with all programming languages, the "random" number functions offered by PHP are really implemented by pseudo-random number generators. This is because conventional computer architectures are deterministic machines that will always produce the same results given the same starting conditions and inputs and have no good source of randomness. (Here we're talking about the ideal computer as it is supposed to work, not the actual physically embodied, power-interruptible, cosmic-ray flippable, seemingly very random machines we all struggle with daily!) You could imagine connecting a conventional computer to a source of random bits such as a mechanical coin-flip reader, or a device that observed quantum-level events, but such peripherals don't seem to be widely available at this time.

So we must make do with pseudo-random generators, which produce a deterministic sequence of numbers that looks random enough for most purposes. They typically work by running their initial input number (the *seed*) through a particular mathematical function to produce the first number in the sequence; each subsequent number in the sequence is the result of applying that same function to the previous number in the sequence. The sequence will repeat at some point (once it generates a particular number for the second time, it is doomed to follow the same sequence as it did the first time around), but a good iteration function will generate a very long sequence of numbers that have little apparent pattern before the loop occurs.

How do you choose a seed to start off with? Because of the generator's determinism, if you hardcode a PHP page to have a particular seed, that page will always see the same sequence from the generator. (Although this is not usually what you want, it can be an invaluable trick when you are trying to debug behavior that depends on the particular numbers that are generated.) The typical seeding technique is to use a fast-changing digit from the system clock as the initial seed — although those numbers are not exactly random, they are likely to vary quickly enough that subsequent page executions will start with a different seed every time.

Here's some representative code that uses the pseudo-random functions:

```
print("Seeding the generator<BR>");
mt_srand((double)microtime() * 1000000);
print("With no arguments: " . mt_rand() . "<BR>");
print("With no arguments: " . mt_rand() . "<BR>");
print("With no arguments: " . mt_rand() . "<BR>");
print("With two arguments: " .
        mt_rand(27, 31) . "<BR>");
print("With two arguments: " .
        mt_rand(27, 31) . "<BR>");
print("With two arguments: " .
        mt_rand(27, 31) . "<BR>");
```

with the browser output:

```
Seeding the generator
With no arguments: 1962311688
With no arguments: 1494083765
With no arguments: 1224081997
With two arguments: 31
With two arguments: 27
With two arguments: 30
```

Obviously, if you run exactly this code, you will get numbers that differ from those in the output shown here, because the point of seeding the generator this way is to ensure that different executions produce different sequences of numbers.

CAUTION   In some old versions of PHP3, the `rand()` function buggily ignored its arguments, returning numbers between `0` and `getrandmax()` regardless of restrictions. We have also heard some reports of that behavior under more recent Windows implementations. If you suspect that you are suffering from such a bug, you can define your own restricted version of `rand()` like this:

```
function my_rand ($min, $max)
{
  return(rand() % (($max - $min) + 1)
        + $min);
}
```

Unlike `rand()`, this version requires the `min` and `max` arguments.

Example: Making a random selection

Now let's use the random functions for something useful (or, at least, something that could be used for something useful). The following two functions let you construct a random string of letters, which could, in turn, be used as a random login or password string:

```
function random_char($string)
{
  $length = strlen($string);
  $position = mt_rand(0, $length - 1);
  return($string[$position]);
}
function random_string ($charset_string, $length)
{
    $return_string = ""; // the empty string
  for ($x = 0; $x < $length; $x++)
    $return_string .= random_char($charset_string);
  return($return_string);
}
```

The `random_char()` function chooses a character (or, actually, a substring of length 1) from its input string. It does this by restricting the `mt_rand()` function to positions within the length of the string (with chars numbered starting at zero), and then returning the character that is at that random position. The `random_string()` function calls `random_char()` a number of times on a string representing the universe of characters to be chosen from and concatenates a string of the desired length.

Now, to demonstrate this code, we first seed the generator, define our universe of allowable characters, and then call `random_string()` a few times in a row:

```
mt_srand((double)microtime() * 1000000);
$charset = "abcdefghijklmnopqrstuvwxyz";
```

```
$random_string = random_string($charset, 8);
print("random_string: $random_string<BR>");
$random_string = random_string($charset, 8);
print("random_string: $random_string<BR>");
$random_string = random_string($charset, 8);
print("random_string: $random_string<BR>");
```

with the result:

```
random_string: eisexkio
random_string: mkvflwfy
random_string: gpulbwth
```

In this example, we seed the generator only once, and we draw that seed value from the system clock. Notice what happens if we make the mistake of repeatedly seeding the generator with the same value:

```
mt_srand(43);
$random_string = random_string($charset, 8);
print("random_string: $random_string<BR>");

mt_srand(43);
$random_string = random_string($charset, 8);
print("random_string: $random_string<BR>");

mt_srand(43);
$random_string = random_string($charset, 8);
print("random_string: $random_string<BR>");
```

Because the sequence that is generated depends deterministically on the seed, we get the same behavior each time:

```
random_string: qgkxvurw
random_string: qgkxvurw
random_string: qgkxvurw
```

In these examples, we chose to draw random characters from strings, but this kind of selection process is generalizable to draw items from arrays or to be used in any situation that requires choosing random members from a set. All you need is the universe of items, a way to put them in numerical order, and a way to retrieve them by order number, and you can then use the rand() or mt_rand() function to choose a random order number for the retrieval.

Summary

The highlights of PHP math are summarized in Table 9-4. Refer to Chapter 27 for more advanced mathematical concepts as they are handled by PHP.

TABLE 9-4

Summary of PHP Math Operators and Functions

Category	Description
Arithmetic operators	Operators +, -, *, /, % perform basic arithmetic on integers and doubles.
Incrementing operators	The ++ and -- operators change the values of numerical variables, increasing them by one or decreasing them by one (respectively). The value of the postincrement form (`$var++`) is the same as the variable's value before the change; the value of the preincrement form (`++$var`) is the variable's value after the change.
Assignment operators	Each arithmetic operator (like +) has a corresponding assignment operator (+=). The expression `$count += 5` is equivalent to `$count = $count + 5`.
Comparison operators	These operators (`<`, `<=`, `>`, `>=`, `==`, `!=`) compare two numbers and return either `true` or `false`. The === operator is true if and only if its arguments are equal and of the same type while the !== is true if the arguments are not equal or aren't of the same type.
Basic math functions	`floor()`, `ceil()`, and `round()` convert doubles to integers, `min()` and `max()` take the minimum and maximum of their numerical arguments, and `abs()` is the absolute value function.

Chapter 10

PHP Gotchas

E ven though we've tried to give clear instructions, and you've no
doubt followed them to the letter, problems can still arise. This
chapter lays out some of the most common problems by symptom
and suggest some frequent causes.

CROSS-REF **There is a whole other universe of gotchas involving database
connectivity. This chapter deals with PHP-only problems. You
may want to skip ahead to Chapter 19 if you're having problems with PHP
and a database. Also, problems specific to certain more advanced features
(including sessions, cookies, building graphics, e-mail, and XML) are dealt
with in their individual chapters in Parts III and IV.**

Installation-Related Problems

Instead of getting moralistic about people who rush through their installs
without understanding the documentation, we'll point out a few common
symptoms that characteristically appear when you've just installed PHP for
the first time.

TIP **If you are seeing similar errors but are confident that your
installation is stable, follow the cross-references to later parts of
this chapter.**

Symptom: Text of file displayed in browser window

If you are seeing the text of your PHP script instead of the resulting HTML, the PHP engine is clearly not being invoked. Check that you are accessing the site through the web server and not via the file-system. Use this:

```
http://localhost/mysite/mypage.php
```

rather than this:

```
file://home/httpd/html/mysite/mypage.php
```

Symptom: PHP blocks showing up as text under HTTP or browser prompts you to save file

The PHP engine is not being invoked properly. If you're properly requesting the file via HTTP as explained previously, the most common reason for this error is that you haven't specified all the file extensions you want to be served by the web server and parsed with the PHP interpreter. Go back to Chapter 2, and review how to configure your Web server to recognize PHP file extensions. The second most common reason is that your `php.ini` file is in the wrong place or has a bad configuration directive.

CROSS-REF   **If you see PHP code in your Web browser and you have a stable installation, your problem is probably due to missing PHP tags. See the "Rendering Problems" section later in this chapter.**

Symptom: Server or host not found/Page cannot be displayed

If your browser can't find your server, you may have a DNS (Domain Name Service) or Web-server configuration issue.

If you can get to the site via IP address rather than domain name, your problem is probably DNS-related.

If you cannot get to the site via IP address for a new installation, it's likely you haven't successfully bound the IP address to your network interface or configured the web server to handle requests for a particular domain (see Chapter 2). If you can't get to the site via IP address for a previously working installation, most likely your Web server is down or unreachable for a reason not related to PHP.

Rendering Problems

This section covers problems where PHP does not report an error per se, but what you see is not what you thought you would get.

Symptom: Totally blank page

A blank page could be caused by any number of issues. Usually, it's caused by a fatal error in the PHP code from which the PHP interpreter cannot recover. Begin by debugging at the top of the PHP file that you're trying to visit by placing a die() after the opening <?php tag:

```
<?php
die(print "hello");
```

If you refresh the page, and see the word hello in the browser, then you've ruled out problems with the web server and the PHP module itself. Continue to move the die() statement further down into the PHP code until you reproduce the blank page error. Don't forget that any files included through a "require," "require_once," "include," or the like could also be causing the script to fail. If you place the die() statement just before an included file and it works and then move the die() just after the included file and the script fails, then you've determined that the problem (or at least a problem) lies in the included file.

Of course, another possible answer in this case is that the PHP module is not working at all. Test by browsing a different page in the same directory that you've previously verified is being correctly handled by PHP.

Also see the "Timeouts" section near the end of this chapter for more information on what happens when you write code that runs "forever."

Finally, you might be seeing a blank screen if your PHP hits a more or less fatal error but you have error reporting turned off. Error reporting should probably be turned off for production servers for security reasons, but error reporting to the browser is actually a huge help for development servers. Check your php.ini file's display_errors setting and make sure the settings are what you expected. If you really dislike error reporting to the browser, you need to make heavy use of the error_log function in exception handling. See Chapters 30 and 31 for more debugging tips.

Symptom: PHP code showing up in Web browser

If you are seeing literal PHP code in your browser, rather than a rendering of the HTML it should be producing, you may have omitted a PHP start tag somewhere. (This assumes that you have had PHP running successfully and that you are using the correct tags for your installation. If not, see the "Installation-Related Problems" section near the beginning of this chapter.)

It's easy to forget that PHP treats included files as HTML, not as PHP, unless you tell it otherwise with a start tag at the beginning of the file. For example, assume that we load the following PHP file:

```
<HTML><HEAD></HEAD><BODY>
<?php include("secret.php");
secret_function(); ?>
</BODY></HTML>
```

167

which includes the file `secret.php`, which in turn looks like this:

```
function secret_function ()
{
  echo "Open sesame!";
}
```

The result is shown in Figure 10-1.

A PHP include appearing as HTML
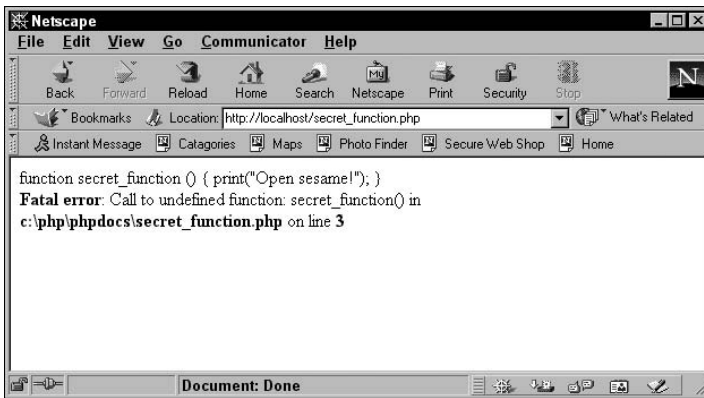
This can be fixed by adding PHP tags to the included file like this:

```
<?php
function secret_function ()
{
  echo "Open sesame!";
}
?>
```

Failures to Load Page

A couple of different kinds of errors are seen when PHP is unable to find a file that you have asked it to load.

Symptom: Page cannot be found

If your browser can't find a PHP page you've created, and you have recently installed PHP, please see the section "Installation-Related Problems" earlier in this chapter. If you get this message when you have

been loading other PHP files without incident, it's quite likely you are just misspelling the filename or path. Alternatively, you may be confused about where the web server document root is located.

Symptom: Failed opening [file] for inclusion

When including files from PHP files, we sometimes see errors like this (on a Unix platform, the file paths would be different):

```
Warning Failed opening 'C:\InetPub\wwwroot\asdf.php' for
inclusion (include_path='') in [no active file] on line 0
```

It turns out that this is the included-file version of `Page cannot be found` — that is, PHP hasn't even gotten to loading the first line of the active file. There is no active file because no file by that name could be found.

It's also possible that you will see this message as a result of incorrect permissions on the file you are trying to load.

Parse Errors

The most common category of error arises from mistyped or syntactically incorrect PHP code, which confuses the PHP parsing engine.

Symptom: Parse error message

Although the causes of parsing problems are many, the symptom is almost always the same: a parse error message like that in Figure 10-2.

FIGURE 10-2

A parse error message
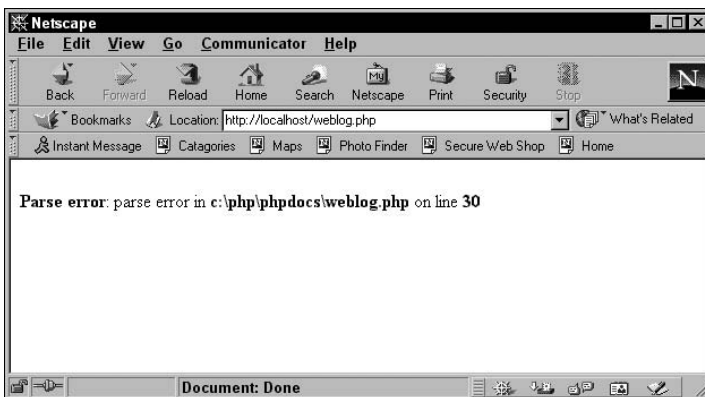
The most common causes of parse errors, detailed in the subsections that follow, are all quite minor and easy to fix, especially with PHP lighting the way for you. However, every parse error returns the identical message (except for filenames and line numbers) regardless of cause. Any HTML that may be in the file, even if it appears before the error-causing PHP fragment, will not be displayed or appear in the source code.

The missing semicolon

If each PHP instruction is not duly finished off with a semicolon, a parse error will result. In this sample fragment, the first line lacks a semicolon, and therefore, the variable assignment is never completed.

```
What we have here is
<?php
$Problem = "a silly misunderstanding"
echo $Problem; ?>.
```

No dollar signs

Another very common problem is that a dollar sign prepending a variable name is missing. If the dollar sign is missing during the initial variable assignment, like this:

```
What we have here is
<?php
Problem = "a big ball of earwax";
echo $Problem; ?>.
```

a parse error message will result. However, if instead the dollar sign is missing from a later output of the variable, like this:

```
What we have here is
<?php
$Problem = "a big ball of earwax";
print("Problem"); ?>.
```

PHP will not indicate a parse error. Instead, you will get the screen shown in Figure 10-3.

This is an excellent example of why you should not rely on PHP to tell you something is wrong. Although PHP's error messages are more informative than most, errors such as this are easily missed if your proofreading efforts aren't up to par.

TIP   **If you spend any significant portion of your time debugging PHP code, an editor that can jump to specific line numbers can be invaluable. Note that the actual mistake that caused the error may be on the line that PHP complains about, or before it, but never after it. For example, because there's nothing wrong with commands that span several lines, a missed semicolon won't cause a parse error until PHP tries to interpret subsequent lines as part of the same statement. Some integrated development environments (IDEs) will do on-the-fly syntax checking while you write. These can be helpful to spot the errors before they get to the server, while you're still coding.**

FIGURE 10-3

A missing dollar sign on variable output

Mode issues

Another family of glitches arises from faulty transitions in and out of PHP mode.

A parse error will result if you fail to close off a PHP block properly, as in:

```
What we have here is
<?php
$Problem = "Bad Code!";
echo $Problem; .
```

This particular mode issue is very common with short PHP blocks. Conversely, if you fail to begin the PHP block properly, the rest of the intended block will simply appear as HTML.

A slightly more tricky issue is engendered by the use of the minimal PHP style, which entails weaving in and out of HTML mode frequently. (See the discussion of minimal versus maximal style in Chapter 33.) For instance, this fragment (which omits the ?> after the first curly brace, when we intend to return to HTML mode) will return a parse error:

```
<?php if(!IsSet($stage))
{
What we have here is
<?php
$Problem = "an awful kerfuffle ";
print("$Problem"); ?>.
<?php
} else {
print("$Stage"); }
?>
```

Another instance of a very common problem is this one, which combines the short block and weaving-in-and-out-of-HTML issues neatly:

```
<FORM>
<INPUT TYPE="TEXT" SIZE=15 NAME="FirstName"
VALUE="<?php print("$FirstName"); ?>">
<INPUT TYPE="TEXT" SIZE=15 NAME="LastName"
VALUE="<?php print("$LastName"); ?>">
<INPUT TYPE="TEXT" SIZE=10 NAME="PhoneNumber"
VALUE="<?php print($PhoneNumber); ?>"
<INPUT TYPE="SUBMIT" NAME="Submit">
</FORM>
```

A PHP double-quote and the HTML closing bracket have been forgotten on the `PhoneNumber` input line here. This will both cause a parse error and prevent the Submit button from appearing on a client browser.

The sample code is meant to demonstrate how easy it can be to forget an element on a crowded page with lots of small but important symbols. You can reduce this type of error either by using a good programmer's text editor or by completing and testing the HTML first and adding the PHP later (or both).

Unescaped quotation marks

Another type of parse error is characteristic of maximal PHP: the unescaped quotation mark.

```
<?php
print("She said, /"What we have here is ");
$Problem = "a difference of opinion\"";
print("$Problem"); ?>.
```

In this case, the double-quote just before the word `What` is incorrectly, and therefore ineffectively, escaped by a forward slash rather than a backslash. If you simply forgot the backslash, the effect would be the same.

Unterminated strings

Failing to close off a quoted string can cause parse errors that refer to line numbers far away from the source of the problem. For example, a code file like this:

```
print("I am a guilty print statement!);  // line 5
// 47 lines of PHP code omitted ...
print("I am an innocent print statement!"); // line 53
```

might well produce a parse error that complains about line 53. This is because PHP is happy to include any text you might want in a quoted string, including many lines of your own code. This

inclusion finishes happily with the first double-quote in line 53, and then the parser finds the symbol I, which it can't figure out how to interpret as PHP code.

If the quotation mark symbol that begins the unterminated string happens to be the last one in the file, the line number in the complaint will be the last line in the file — again, probably far away from the scene of the crime.

Other parse error causes

The problems we have named are not an exhaustive list of the sources of parse errors. Anything that makes a PHP statement malformed will confuse the parser, including unclosed parentheses, unclosed brackets, operators without arguments, control structure tests without parentheses, and so on. Sometimes the parse error will include a statement about what PHP was expecting and didn't find, which can be a helpful clue. If the line of the parse error is the very last line of the file, it usually means that some kind of enclosure (quotation marks, parentheses, braces) was opened and never closed, and PHP kept on hoping until the very end.

Missing Includes

In addition to loading top-level source files, PHP needs to be able to load any files you bring in via `include()` or `require()`.

Symptom: Include warning

This kind of error is shown in Figure 10-4.

FIGURE 10-4

Include warning
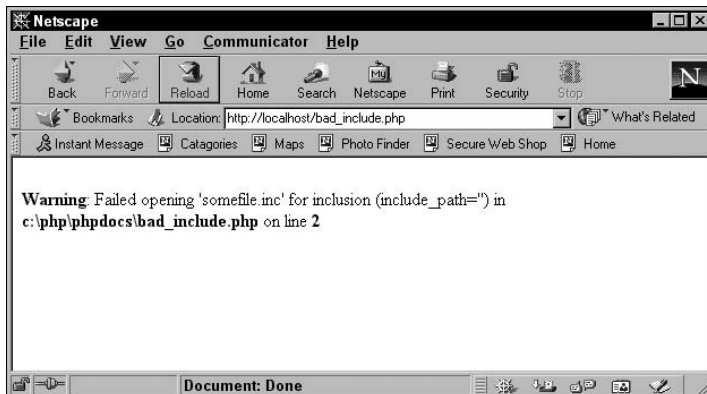
The problem is that you call somewhere in the script for a file to be included, but PHP can't find it. Check to see that the path is correct. You might also have a case sensitivity or other typographic issue. Note the important difference between `include()` and `require()`. If a file is included and PHP can't locate the file, execution of the script will continue with a PHP warning. If a file is required and PHP can't locate that file, execution will stop with an error.

Unbound Variables

PHP is different from many programming languages in that variables do not have to be declared before being assigned, and (under its default settings) PHP will not complain if they are used before being assigned (or *bound*) either. As a result, forgetting to assign a variable will not result in direct errors — either you will see puzzling, but error-free output, or you will see a downstream error that is a result of variables not having the values you expected. (If you would rather be warned, you can set the error-reporting level in `php.ini` or by evaluating `error_reporting(E_ALL)`.) Some symptoms of this kind of problem follow.

Symptom: Variable not showing up in print string

If you embed a variable in a double-quoted string (`"like $this"`) and then print the string using `print` or `echo`, the variable's value should show up in the string. If it seems to not be there at all in the output ("like "), the variable has probably never been assigned.

Symptom: Numerical variable unexpectedly zero

Although it's possible to have a math error or misunderstanding result in this symptom, it's much more likely that you believe that the variable has been assigned when it actually hasn't been.

Causes of unbound variables

PHP automatically converts the types of variables depending on the context in which they are used, and this is also true of unbound variables. In general, unbound variables are interpreted as `0` in a numerical context, `""` in a string context, `FALSE` in a Boolean context, and as an empty array in an array context. The following code shows the effect of forgetting to bind two variables (`$two_string` and `$three`); the resulting display appears in Figure 10-5:

```php
<?php
$one_string = "one";
$three_string = "three";
$one = 1;
$two = 2;
print("This math is as easy as $one_string, $two_string,
$three_string!<BR>");
print("$one_string is equal to $one<BR>");
print("$two_string is equal to $two<BR>");
```

```
print("$three_string is equal to $three<BR>");
print("$one_string divided by $two_string is " .
        ($one / $two) . "<BR>");
print("$one_string divided by $three_string is " .
        ($one / $three) . "<BR>");
?>
```

FIGURE 10-5

The effect of unbound variables
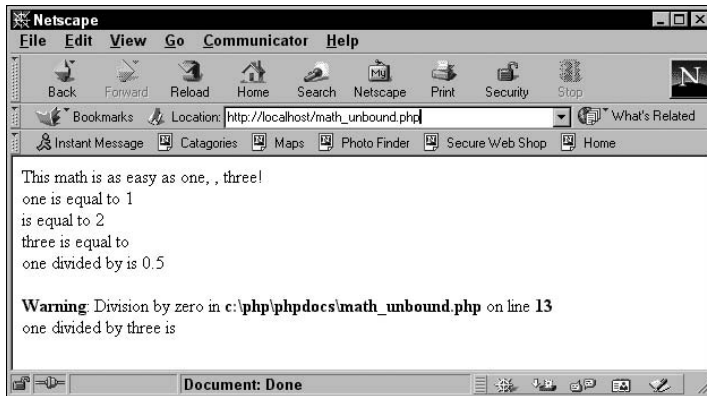
Case problems

Variables in PHP are case sensitive, so the same name with different capitalization results in a different variable. Even after a value is assigned to the variable $Mississippi, the variable $mississippi will still be unbound. (Capitalization aside, variables that are this difficult to spell are probably to be avoided for the same reason.)

Scoping problems

As long as no function definitions are involved, PHP variable scoping is simple: Assign a variable, and its value will be there for you from that point on in that script's execution (until the variable is reassigned). However, the only variables that are available inside a function body are the function's formal parameters and variables that have been declared to be global — if you have a puzzling, unbound variable inside a function, this is probably something you've forgotten. In the following code, for example, the variable $serial_no is neither passed in to the function nor declared to be global:

```
$name = "Steve Suehring";
$rank = "Intarweb Programmer";
$serial_no = "4";

function Answer($name)
```

```
{
global $rank;
print("Name: $name; Rank: $rank;
      serial no: $serial_no<BR>");
}
Answer($name);
```

The resulting browser output looks like:

```
Name: Steve Suehring, Rank: Intarweb Programmer, serial no:
```

because the variable is unbound inside the function.

Variable Naming Conventions

One way to avoid a lot of the gotchas in PHP is to decide on, and to rigorously use, a set of variable naming conventions for all of your code. In the frequent cases where variables will be assigned and used in widely separated places in the same script and even across scripts, such a set of standards will save lots of time referring back and forth. What conventions you decide on are less important than that you have some standard in the first place. That said, here are a few tips to help you decide what to do:

- A common mistake many new programmers make is thinking that variables must somehow be an abbreviation of the thing they represent. Remember, a variable is not an abbreviation, but rather a stand-in for some value that may change depending on circumstances or as a script executes. A longer, meaningful, and easy-to-remember variable name is better than a shorter variable name that is anybody's guess.

- Variable names that consist of multiple words strung together can be made more readable by using underscores (for example, `$office_address`) or initial capitalization (`$OfficeAddress`). There is some sense to the notion that the underscore solution can create confusion with function-naming conventions. Use what works best for you.

- In a more general sense, remember that you may not be the only person that has to read this code. You may get really excited about PHP and get involved in one of the many open source projects that use PHP. You may even start your own project (we'd be delighted to see that happen)! In either case, readable code will be a must, and good variable names are a foundation of producing readable code.

Function Problems

Many problems having to do with function calls result in *fatal errors*, which means that PHP gives up on processing the rest of the script.

Symptom: Call to undefined function my_function()

PHP is trying to call the function `my_function()`, which has not been defined. This could be because you misspelled the name of a function (built-in or user-defined) or because you have simply omitted the function definition. If you use `include/require` files to load user-defined functions, make sure that you are loading the appropriate files.

If the problem involves a fairly specialized, built-in function (for instance, it is related to XML or arbitrary-precision math), it may be that you did not enable the relevant function family when you installed or configured PHP.

Symptom: Call to undefined function ()

In this case, PHP is trying to call a function and doesn't even know the function's name. This is invariably because you have code of the form `$my_function()`, where the name of the function is itself a variable. Unless you are intentionally trying to exploit the variable-function-name feature of PHP, you probably accidentally put a $ in front of a sensible call to `my_function()`. Because `$my_function` is an unbound variable, PHP interprets it as the empty string — which is not the name of a defined function — and gives this uninformative error message.

Symptom: Call to undefined function array()

This problem has a cause that is similar to the cause of the previous problem, although it still baffled us completely the first time we ran into it. It can arise when you have code like the following:

```
$my_amendments = array();
$my_amendments(5) = "the fifth";
```

Unless you look closely, this looks like an innocent pair of statements to create an array and then store something in that array, with the number 5 as a key. And yet PHP is telling us that `array()` is an unbound function, even though we know that it is a very standard built-in function. What's going on?

The fault is actually with Line 2 above, rather than with Line 1. If we want to access an element of `$my_amendments`, the correct syntax is `$my_amendments[5]`, with square brackets. Instead, we used parentheses, which the parser interprets as an attempted function call. It takes what is immediately before the left parenthesis to be a function. Instead, what comes before the parenthesis is an array, which is not a function; PHP gives up on us, with this obscure complaint.

Symptom: Cannot redeclare my_function()

This is a simple one — somewhere in your code you have two definitions of `my_function()`, which PHP will not stand for. Make sure that you are not using `include` to pull in the same file of function definitions more than once. Use `include_once` or `require_once` to avoid seeing this error, with the caveat that, well, you won't see this error. Why might that be bad? It's conceivable that you could define two distinctly different functions and inadvertently give them the same name. This runs the risk of exposing your mistake at a somewhat inconvenient moment.

Symptom: Wrong parameter count

The function named in the error message is being called with either fewer or more arguments than it is supposed to handle. In the case of more parameters you're okay, but if you use fewer parameters than is expected you will get an error.

Math Problems

The problems that follow are specific to math and the numerical data types.

Symptom: Division-by-zero warning

Somewhere in your code, you have a division operator where the denominator is zero. The most common cause of this is an unbound variable, as in:

```
$numerator = 5;
$ratio = $numerator / $denominator;
```

where `$denominator` is unbound. It's also possible, of course, that the legitimate result of a computation is producing a zero denominator. In this case, the only thing to do is catch it with a test and do something reasonable if the test applies. See the following example:

```
$numerator = 5;
if (isset($denominator) && $denominator != 0)
    $ratio = $numerator / $denominator;
else
    print("I'm sorry, Dave, I cannot do that<BR>");
```

Symptom: Unexpected arithmetic result

Sometimes things just don't add up (or multiply up, or subtract up). If you are having this experience, check any complex arithmetic expressions for unbound variables (which would act as zeros) and for precedence confusions. If you have any doubt about the precedence of operators, add (possibly redundant) parentheses to make sure the grouping is as you intend.

Symptom: NaN (or NAN)

If you ever see this dreaded acronym, it means that some mathematical function you used has gone out of range or given up on its inputs. The value `NAN` stands for "Not a Number," and it has some special properties. Here's what happens if we try to take the arccosine of 45, even though arccosine is defined only when applied to numbers between –1.0 and 1.0:

```
$value = acos(45);
print("acos result is $value<BR>");
print("The type is " . gettype($value) . "<BR>");
```

```
    $value2 = $value + 5;
    print("Derived result is $value<BR>");
    print("The type is " . gettype($value2) . "<BR>");
    if ($value == $value)
      print("At least that much makes sense<BR>");
    else
      print("Hey, value isn't even equal to itself!<BR>");
```

The browser output looks like:

```
acos result is NAN
The type is double
Derived result is NAN
The type is double
Hey, value isn't even equal to itself!
```

Oddly enough, NAN *is* a number, at least in the sense that its PHP type in this example turns out to be double rather than string. It also infects other values with not-a-numberness when used in math expressions. (This behavior is a feature, not a bug, when used in very complex calculations that must be correct. It's better to have the whole value be tagged as untrustworthy than have one subexpression be silently bogus.) Finally, any equality comparison that involves NAN will be false — NAN is neither less than, nor greater than, nor equal to any other number, including itself. It is always unequal (!=) to all numbers, including itself. (The NAN value is not a PHP-specific feature — it is part of the IEEE standard for floating-point arithmetic, which is implemented by the C functions that underlie PHP.)

Because of the contagion of NAN values, this kind of problem can be difficult to debug. The best way to try to find the original offending NAN is with diagnostic print statements, especially because comparison tests will give counterintuitive results. You can explicitly test for NAN values using the built-in is_nan() function, which returns TRUE if the number submitted is not a number or FALSE otherwise. In earlier versions (you aren't using an earlier version, are you?), you can cobble together your own function for NAN testing like this:

```
    function is_nan($value)
    {
      return($value != $value);
    }
```

It uses the weird comparison properties of NAN as a type checker.

Timeouts

Of course any download can occasionally time out before a complete page can be delivered. However, this shouldn't be happening frequently on your local development server!

179

The most interesting reason for a timeout is an infinite loop. These can be difficult to track down quickly, as in this example:

```
//compute the factorial of 10
$Fact = 1;
for ($Index = 1; $Index <= 10; $index++)
  $Fact *= $Index;
```

This code shows a nasty little collaboration between a loop and a case confusion — the lowercase `$index` that is incremented has nothing to do with the `$Index` that is being tested, so the test will never become false.

Summary

In Table 10-1, we summarize the gotchas in this chapter by mapping symptoms to possible causes. We also offer some suggestions on how to fix the most common problems.

TABLE 10-1

From Symptoms to Causes

Symptom	Possible Causes	Advice
(New installation) Text of file displayed in browser window	The PHP engine is not being invoked, possibly because you are opening it via the local filesystem rather than as a request to your server.	Make sure that your request is to the web server, either via localhost (`http://localhost/[path]`) if testing on the server machine, or by the full URL (`www.example.com/[path]`).
(New installation) PHP blocks showing up as text, or browser prompts you to save file	PHP is not being invoked properly. Your web server may not be set up to map the right file extensions (for example, `.php`) to the PHP engine, or there may be a problem with the location or contents of `php.ini`.	Check your web server configuration, and the PHP init file (`php.ini`).
(New installation) Server or host not found/ Page cannot be displayed	Often due to Internet/DNS/web-server configuration problems, rather than PHP.	Try loading a pure HTML file with a file extension you have not set up for PHP (for example, `.html`) to rule out PHP problems.
Totally blank page	Usually due to PHP syntax errors.	Use `die()` to determine the location of the syntax error.

Symptom	Possible Causes	Advice
PHP code showing up in browser window	If the PHP engine is installed and functioning properly, this is usually due to a missing PHP start tag or misconfigured web server.	Check start and end tags and make sure that any `include` files of PHP code have correct tags at beginning and end also check web server functionality with a basic PHP page.
Parse error message	A variety of causes, including missing semicolons, variables without a $, unescaped quotation marks, unclosed quotation marks, brackets, or parentheses, and HTML being interpreted as PHP.	Locate the line with the parse error in the PHP file, and look for one of the causes in that line or the lines immediately preceding it. If the "error" is on the final line of the file, look for an unclosed quote, parenthesis, or bracket, possibly much earlier in the file.
`Include` warning	For one reason or another, PHP was not able to load a file named in an `include` statement.	Check that the file actually exists, the spelling of the filename, the pathname, and (on Unix systems) the case of the name. Also make sure that the file permissions allow the file to be read.
Variable value not showing up in `print` string	The variable has not been assigned, and so its value in a printed string is the empty string.	Check that you are assigning the variable before the `print` statement and compare spelling and case (capitalization). Make sure that you are not embedding any objects or multidimensional arrays in quoted strings. You can also use the statement `error-reporting(15)` to tell PHP to warn about any unbound variables.
Numerical variable unexpectedly zero	Often due to the variable never having been assigned.	(See preceding.)
Variable value is valid, but unexpected.	Often due to variable having been unexpectedly overwritten.	Use good variable names; search through all included files for variable name.
Call to undefined function `my_function()`	Function `my_function()` is being called without having been defined first.	If you are trying to call a function of your own, check that the definition (or inclusion of the file containing the definition) is before the use. If you are trying to call a built-in function, check the spelling. If it is correct, investigate whether that "family" of functions was included when you configured PHP (for example, either all the XML functions will work, or none will).

continued

181

TABLE 10-1   *(continued)*

Symptom	Possible Causes	Advice
Call to undefined `function ()`	An expression of the form `$my_function()` is being evaluated, and `$my_function` is not bound to the name of a defined function.	If you intend to use the variable-function feature, then add (or correct) the assignment of `$my_function`. If you are just trying to call `my_function()`, remove the $.
Call to undefined function `array()`	You probably have an expression of the form `$array_var_name(3)`, when what you want is `$array_var_name[3]`	Decide whether you want an array expression or a function call — if the former, then change parentheses to square brackets.
Cannot redeclare `my_function()`	The function `my_function()` is being defined twice in a page's execution.	Look for double definitions of `my_function` in the PHP file, or double-inclusions of the file that defines it.
Wrong parameter count	The named function (usually a built-in function) is being called with an incorrect number of arguments.	Compare the function call to the definition in the online PHP manual (`www.php.net`)
Division-by-zero warning	A / operator has a right-hand argument of zero. Can be due to an unbound variable in the denominator.	Assign the unbound variable if that's the cause. If the desired logic could actually result in zero denominators, install a test to catch that case.
Unexpected arithmetic result	Frequently due to an unbound variable in an arithmetic expression.	Check for unbound variables (see preceding), and make sure that arithmetic expressions are parenthesized appropriately.
`NAN` value	A built-in math function is being given inputs outside its acceptable range. If that function's results are used in arithmetic, the results are also `NAN`.	Trace backward from the `NAN` value to function calls that contribute to its computation. Test with `print` statements, or test for values that fail to be self-equal (a diagnostic for `NAN`).

Part II

MySQL Database Integration

Chapter 11

Introducing Databases and MySQL

Databases and PHP go together like cake and ice cream, Trinidad and Tobago, green eggs and ham — you get the picture.

After all, what's the Web about? Making vast stores of information available to a more or less wide public, that's what. Not that there aren't small brochureware sites galore, but the bigger and more frequently updated the data source, the more comparative value is provided by the Web over other media.

Perhaps the single greatest advantage of PHP over similar products is the unsurpassed choice and ease of database connectivity it offers. As detailed in the "Choosing a Database" section of this chapter, PHP supports native connections to a number of the most popular database server types, open source and commercial alike. Almost any database that will open its application programming interface (API) to the public seems to be included eventually. For any unsupported databases, there's generic ODBC (Open Database Connectivity) support.

What Is a Database?

A *database* is a collection of data. The term database usually indicates that the collection of data is stored on a computer. Regardless, it's the databases that are on computers that I'll concentrate on in this book.

Databases implemented through a computer are created within software. That software, commonly known as a database application, controls how

the actual data is stored and retrieved. Some database applications include Microsoft Access and OpenOffice.org's Base. Sometimes, databases are stored in a central location and managed by a database server. A database server is a database application built with multiple users in mind.

Most of the time when programming PHP you'll be accessing a database server. Some database servers include PostgreSQL, MySQL, Microsoft's SQL Server, and the Oracle suite of databases. You may also see database servers called RDBMS, which is an acronym for relational database management system.

Database servers usually have one or more distinct APIs for programmatically creating, accessing, managing, searching, and replicating the data they hold. It is through the API that you connect to and work with data stored in database servers when using PHP.

There is no requirement that an RDBMS be used to store data. Other data stores can be used such as a flat file or a table known as a hash table. These are perfectly fine for some applications, especially smaller applications; however, for larger applications or applications that require optimal speed for large data stores, an RDBMS is a requirement.

Why a Database?

If you're going to the trouble to use PHP at all, you're likely to need a database sooner or later — probably sooner. Even for something small, like a personal blog, you want to think hard about the advantages of using a database instead of static pages or included text files.

Maintainability and scalability

Having PHP assemble your pages on the fly from a template and a database is an addictive experience. Once you enjoy it, you'll never go back to managing a static HTML site of any size. For the effort of programming one page, you can produce an infinite number of uniform pages. Change one, and you've changed them all.

There are now web sites with hundreds of thousands of separate pages — you can rest assured that no one is maintaining them all by hand. If you have a web site that may eventually grow to more than a few dozen pages, you should think about moving to a database sooner rather than later.

Portability

Because a database is an application rather than a part of the operating system, you can easily transfer its structure and contents from one machine to another or (in certain cases) even from one platform to another. This is especially valuable for contractors, who may develop a project without being able to control the environment in which it will eventually be deployed — they can deliver a package of PHP plus a MySQL database schema dump.

Avoiding awkward programming

Certain things can be done with PHP but probably shouldn't, because they entail ugly or risky programming moves.

Say that you happen to be the commander of the starship *Enterprise* and are keeping a captain's log. Each log entry is contained in a text file identified by its unique stardate, which is plugged into a template by PHP — but hey, you're a busy spaceman with whole galaxies to explore; you don't always have time to write in your log every day. You want to put automatically generated Next and Previous links on each page for those who wish to read in straight chronological order. It's pretty easy to use PHP to find the previous stardated entry, but any attempt to locate the next entry can quickly become an infinite loop — because it's easier to prove something does exist than that it doesn't. On the other hand, if you put your log data in a database, the whole job becomes trivial. The database will tell you which is the latest entry at any given moment.

There are other types of programming tasks that a database is highly optimized to do, and given the option, you should take advantage of it to perform these chores. For instance, you should avoid sorting data sets on the PHP side in favor of writing queries so the data is returned presorted. We discuss these efficiency issues in greater detail in Chapter 18.

Searching

Although it's possible to search multiple text files for strings (especially on Unix platforms), it's not something most web developers will want to do often. After you search a few hundred files, the task becomes slow and hard to manage. Databases exist to make searching easy. With a single command, you can find anything from one ID number to a large text block to a JPEG-format image.

In some cases, information attains value only when put into a searchable database. For instance, relatively few people would want to read a long text list of movie directors and their films, but many might occasionally want to search a database of that information. You could argue that it's the searchability, as much as the information itself, that creates the value here.

PHP-Supported Databases

PHP Data Objects (PDO) was introduced back with the 5.1 release of PHP. PDO creates a consistent, abstracted interface to database servers and data. PHP offers several database-specific drivers for both PDO and non-PDO access. The PHP web site contains a list with the latest information about databases that can be integrated along with the PDO abstraction layer and other abstraction layers. See `www.php.net/pdo` for more information.

Our Focus: MySQL

MySQL, (officially pronounced my- S - Q - L and not "mysequel"), is an incredibly popular and powerful RDBMS. MySQL provides one of the letters in the ubiquitous acronym "LAMP," which is an abbreviation for Linux, Apache, MySQL, PHP/Perl/Python. MySQL has become so popular for several reasons. First, MySQL is free (as in price), although the licensing has changed (discussed later). Second, MySQL is also stable, meaning that it's not prone to crashing even under load. Third, MySQL is lightweight, meaning that it doesn't require many resources to install or run. Fourth, MySQL is fast and easy to use. Finally, MySQL is powerful, with all of the features required for web applications.

MySQL AB, which is the company behind MySQL (owned by Sun), changed the licensing for MySQL relatively recently. In the latest iteration as of this writing, MySQL offers a product called MySQL Server Community Edition, which is essentially the same as the MySQL Enterprise Server, but is lacking official MySQL support and some graphical user interface (GUI) tools. If your organization needs an officially supported product, where you can call for assistance with the database server at any time, then MySQL Enterprise is for you. MySQL AB's support is excellent; it's not unheard of to get responses from developers themselves. Otherwise, the MySQL Server Community Edition is your choice. For more information on the differences between the two versions, see `www.mysql.com/products/which-edition.html`.

I'll be concentrating on the MySQL Server Community Edition in this book, and the next chapter will show you how to obtain and install MySQL.

Summary

The great advantage of the Web is its capability to make large quantities of information publicly available quickly and cheaply. This functionality has been tremendously enhanced by the recent increase in availability of inexpensive, reliable databases.

PHP supports several types of databases, including flat-file, hash, and relational databases. Most large web sites (and even small sites, too) use some sort of relational database management system (RDBMS). MySQL is a common choice among PHP developers. MySQL is not only free but also light-weight, stable, and full of features necessary for both online and offline applications.

Chapter 12

Installing MySQL

Before jumping into MySQL installation you need to get the software. MySQL's database server can be downloaded from MySQL's web site at `www.mysql.com`. As of this writing, the free Community Edition server feels somewhat hidden on the web site. Therefore, with the caveat that the URL may change on a whim by the time you read this text, the download section for MySQL is currently located at `http://dev.mysql.com/downloads`. However, realize that most distributions of Linux include their own MySQL server package.

Obtaining MySQL

I strongly recommend using the MySQL server package directly from your Linux distribution rather than downloading from MySQL AB unless you have a very specific reason for using a different version. If you can't think what one of those specific reasons might be, then you probably don't have one, and you therefore should use the MySQL server available with your distribution.

Installing MySQL on Linux

There are several distributions upon which you might find yourself installing MySQL. It's always a challenge choosing which distributions to cover. No matter which ones we decide to cover there will always be someone installing on another distribution.

In this section I'll examine MySQL installation on Debian, CentOS, and Ubuntu. Additionally, I'll demonstrate compiling MySQL from source for those who don't have a MySQL server package available with their distribution. It should be noted that because MySQL 6 is so new it may not be available as a package in your distribution. If this is the case, I recommend sticking with the latest MySQL available for your distribution. For the most part, this book will use functions available in MySQL 5 and later, so MySQL 6 isn't a requirement. Where MySQL 6 is required, a special note will be shown.

Installing MySQL Server on Debian and Ubuntu

Debian's dpkg and apt installation and package management tools make installation of MySQL (and everything else for that matter) incredibly easy. Debian is a system administrator's dream because it's so stable, package installation is so easy, and the packages are maintained and configured with excellent defaults. But enough evangelizing; installation of MySQL server on Debian requires superuser privileges and is accomplished simply by running `apt-get`:

```
apt-get install mysql-server
```

Of course, that assumes that you have correctly configured sources in `/etc/apt/sources.list`. For more information on APT and configuration of the `sources.list` file, see `www.debian.org/doc/manuals/apt-howto/ch-basico.en.html`. Debian's package management system will install and configure any necessary prerequisites for you.

Debian separates MySQL into its components such as server, client, and libraries. Therefore, in order to use MySQL and PHP together, you should install the php5-mysql package:

```
apt-get install php5-mysql
```

As you can see by that installation command, the PHP5 version of the interface is being installed. That is the latest version available as of this writing.

Finally, you'll likely also want to install the MySQL command-line interface (CLI), which is accomplished by installing the mysql-client package:

```
apt-get install mysql-client
```

MySQL will now be installed and ready to use on your Debian server. However, by default the MySQL server won't listen on anything by localhost. To change this, edit `/etc/mysql/my.cnf` and comment out the `skip-networking` line with a pound sign or hash mark (*#*), so it looks like this:

```
#skip-networking
```

Now restart the MySQL server by typing this command:

```
/etc/init.d/mysql restart
```

Installing MySQL on Microsoft Windows

MySQL installation on Windows is much, much easier than it used to be thanks to fully automated installers

Installing MySQL on Windows

Default installation on any version of Windows is now much easier than it used to be, as MySQL now comes neatly packaged with a native Windows installer. Simply download the installer package, usually an msi, and run it. This will walk you through the trivial process and by default will install everything under `C:\Program Files\MySQL`, which is probably as good a place as any.

The MySQL installer will attempt to install itself as a service, which means you need Administrator rights on the computer upon which MySQL is being installed. Part of the installation process will configure the MySQL server. During this portion of the installation, you can configure things like the root password, the port on which MySQL will listen, and whether to include the MySQL utilities in the Windows path (I recommend that you do so).

The Windows install is now so simplified that for most cases you can simply click "Next" to continue and, where you have an exception, refer to the online manual for MySQL at `www.mysql.com`.

Summary

This chapter examined installation of MySQL on Linux and Windows. The Linux installation varies somewhat depending on the flavor of Linux on which MySQL is being installed. However, the Windows installation has been greatly refined and reduced to simply clicking through the installation and receiving a fully functional yet incredibly powerful database system. The online documentation for MySQL is available for assistance with installation issues, should they arise.

Chapter 13

Learning Structured Query Language (SQL)

This chapter is a basic introduction to SQL databases in which we discuss standards, database design, Data Manipulation Language, Data Definition Language, and database security procedures common to all SQL databases.

NOTE This chapter is in no way a comprehensive guide to SQL or to any particular SQL database. To go beyond the simplest common features, you will need to consult your particular manufacturer's documentation or specific books. You will also want to look at documentation and books relating to your specific SQL database.

Relational Databases and SQL

SQL is the language of relational databases. A simple query like a one-table SELECT will be more or less the same whether you're using a tiny database like mSQL or an expensive behemoth like Oracle.

The big advantage for you, the web developer, is that, after you learn SQL, you will be able to interact with numerous databases across all platforms without a steep retraining curve. Just imagine how horrible life would be if Oracle, MySQL, and SQL Server all had entirely different sets of commands for putting data in and getting data out of their stores — as if Oracle used SELECT to ask for data sets, MySQL used VALJ (the developers are Swedish, you know), and SQL Server used FIND IT IN THIS TABLE (to better match

the vocabulary of Windows). SQL is the common vocabulary and syntax that will save you from this nightmare. There are differences among products, and in their implementations of the SQL standard and the extensions they each define to that standard, but it's better to have 80 percent in common and 20 percent different than the other way around.

SQL Standards

According to Andrew Taylor, original inventor of SQL, *SQL* does *not* stand for *Structured Query Language* (or anything else for that matter). But for the rest of the world, it does now. As you would expect from the (non-) title, SQL represents a stricter and more general method of data storage than the previous standard of flat-file DBM-style databases.

SQL is a standard under both the American National Standards Institute (ANSI) and the Equipment Managers Council of America (ECMA); both are international standards-maintenance organizations. You can read the standards on payment of a fee to these organizations:

- `www.ansi.org`
- `www.ecma.org`

However, within the general guidelines of the standard there are considerable differences among the products of individual companies and open source database development organizations. The past few years, for instance, have seen the rapid growth of so-called *object-relational* databases, as well as of SQL products specifically slanted toward the web market.

The key to choosing a database is to be selfish, or at least supremely self-centered. You will see plenty of unusually virulent postings out there opining that a certain advanced database feature (like triggers or cross joins) is a "must," and any SQL installation without this feature hardly deserves the name. Take this stuff with a grain of salt. It's far better to make a blind shopping list of functions you need in order of importance and then go out looking for the product that best meets your requirements.

That said, a good deal of SQL really is pretty standardized. You will be using a few SQL statements over and over and over, no matter which specific product you choose to deploy.

The Workhorses of SQL

The basic logical structure of a SQL database is very simple. A given SQL installation can usually contain multiple *databases* — for instance, one for customer data and one for product data. (It's problematic that both the SQL server itself and the collections of tables within it are commonly referred to by the term *database* — but what can you do?) Each database contains a number of *tables*. Each table is made up of carefully defined *columns*, and every entry can be thought of as an added *record* or row. (It's not really a row, but this is a concept so stuck in our visualization that we may as well go with it.)

Four so-called *data manipulation statements* are supported by every SQL server and will constitute an extremely high percentage of all the things you'll want to do with a relational database. These four horsemen of the database are SELECT, INSERT, UPDATE, and DELETE. These commands are your friends and helpmates; get comfy with them, and they will serve you well.

The thing to remember about these four SQL statements is that they manipulate only database *values,* not the structure of the database itself. In other words, you can use these commands to add data but not to make a database; you can get rid of every piece of data in a database, but the shell will still be there — so, for instance, you wouldn't be able to name another database on the same server with the same name. If you want to add or get rid of columns, blow away entire databases as if they never existed, or make up new databases, you need to use other commands such as DROP, ALTER, and CREATE. We discuss these in the "Database Design" section later in this chapter.

TIP A note on SQL style*: Many SQL queries that you see are written in one long line of code — which becomes totally illegible once you're dealing with more than four or five fields. A very accomplished PL/SQL programmer of our acquaintance recommends that you break up every SQL statement into as many lines as you need for maximum legibility. He also does not shy away from using indentation in a SQL query with many variables. (SQL queries are usually quite whitespace insensitive.) He has years of experience working on big Oracle installations, and his recommendations actually are very helpful — so that is the style we try to use in this book.*

SELECT

SELECT is the main command you need to get information out of a SQL database. The basic syntax is extremely simple:

```
SELECT field1, field2, field3
FROM table
```

That's no harder than asking your coworker to get you last month's sales records from the file cabinet in the hallway.

In some cases, you'll want to ask for entire records instead of picking out individual pieces of information. This practice is generally frowned upon, but it is still widely used and, therefore, we need to mention it. A whole record is called for by using the wildcard (asterisk) symbol:

```
SELECT *
FROM mytable
```

Selecting Certain Records

The previous two examples show how to retrieve all rows from the table. It's not all that common to do this in the real world, which is where the WHERE clause comes in. The WHERE clause places a condition on the SELECT statement that causes only those rows matching the WHERE clause to be returned in the result set. For example:

```
SELECT *
FROM mytable
WHERE ID < 100;
```

This example retrieves all fields from the table `mytable` where the ID column value is less than the integer 100. `WHERE` clauses can get quite complex, and, frequently, multiple conditions are used together with the `AND` keyword.

Joins

Joins are one of the main useful features of SQL.

A `SELECT` statement on a single table without joins might be visualized as being something like a row in a spreadsheet. But an SQL database is by definition relational. To understand the philosophy behind the relational database concept, you have to think back to some occasion on which you were forced to fill out a whole bunch of forms — such as applying for a loan, visiting a doctor's office for the first time, or dealing with some kind of governmental formality. (If you've never had this experience, it's because you're young enough to have lived entirely in a world of relational databases.) As you were writing down your name, address, phone, and Social Security number for the 15[th] time, you probably thought, "Why can't I just write my address down once, and then they could just look it up on a need-to-know basis?" That's exactly the concept behind a relational database.

The way a relational database differs from paper forms is the main identifier. Humans do well with text and prefer to categorize by textual identifiers such as names. If a dentist's office or auto body shop stored its paper files in numerical order, it would be difficult for anyone to lay his hands on John Johnson's forms when John next required service. Frankly, most paper file users these days ask for your Social Security number as a backup — it works solely to differentiate you from other people in their files with exactly the same first, last, and middle names.

Databases, on the other hand, work well with integers. You'll frequently use integer values to create unique identifiers or IDs within a database table. This field or column is then called a primary key, which indicates that each value in that column will be unique and that the rows within that column will always have a value in the primary key field. Because primary keys are unique by nature, a database needs only one to identify a person, place, or thing uniquely — no matter how many tables refer to that piece of information.

So instead of needing to repeat information several times, like this:

```
Name:  John Johnson
SS#:  123-45-6789

Name:  John Johnson
Fears:  Cats, Friday the 13th, Flying

Name:  Jane Jones
SS#:  987-65-4321

Name:  Jane Jones
Fears:  Heights, Flying
```

with a relational database you can write down each piece of information just once and then relate it to each other piece using integers, as shown in Tables 13-1 to 13-3.

TABLE 13-1

People

PersonID	Name	SS#
1	John Johnson	998-00-9889
2	Jane Jones	987-65-4321
3	Aloysius Snuffleupagus	987-65-4329

TABLE 13-2

Fears

FearID	Fear
1	Black cats
2	Friday the 13th
3	Peanut butter sticking to the roof of your mouth
4	Heights
5	Flying

TABLE 13-3

Person_Fear

ID	PersonID	FearID
1	1	1
2	1	2
3	1	5
4	2	4
5	2	5

This is clearly a neater and faster (for a database) way to store this information. But when you need to pull out the data into a human-readable form, there's a problem: You have to get and correlate information from more than one database. That's the job of a join.

To find out what phobias were suffered by Ms. Jones, you could first look up her personal unique ID:

```
SELECT PersonID
FROM People
WHERE Name = 'Jane Jones';
```

that returns the unique integer 2. Then you can define another SELECT statement using that information:

```
SELECT FearID
FROM Person_Fear
WHERE PersonID = 2;
```

You get the values 4 and 5 back, which you can use in a third query:

```
SELECT Fear
FROM Fears
WHERE FearID = 4 OR FearID = 5;
```

This returns the values Heights and Flying. We should make it clear that there is nothing inherently incorrect about doing it this way, as long as any performance loss is within parameters acceptable to you.

Alternatively, you can perform a join, which returns the same information in a single SELECT statement:

```
SELECT Fears.Fear
FROM Fears, Person_Fear, People
WHERE Fears.FearID = Person_Fear.FearID
AND Person_Fear.PersonID = People.PersonID
AND People.Name = 'Jane Jones';
```

An alternate syntax for this join is:

```
SELECT Fears.Fear
FROM (Fears INNER JOIN Person_Fear ON FearID INNER JOIN People on
PersonID)
WHERE People.Name = 'Jane Jones';
```

As you can see, you need only know one single piece of information to be able to get all the data in the database about that subject using joins. In effect, a join makes two or more tables into one for purposes of searching for a particular piece of information.

Joins come in several different flavors. The one in the preceding example is called an *inner join*, which is the most common and restrictive type. Another common type is the *outer join*. This is used

to return a list of all fears even if they do not have people attached to them. In this example, we are using a *left outer join* (also known as a *natural join*):

```
SELECT Fear
FROM Fears LEFT JOIN People ON PersonID;
```

Fears that have people attached to them would appear in the data set multiple times, but fears without people would each appear once.

You can also get a list of all people even if they do not have fears attached to them, using a *right outer join*:

```
SELECT Name
FROM Fears RIGHT JOIN People ON PersonID;
```

Again, the fears that are actually attached to people appear multiple times, whereas the fears that are not suffered by any people still show up once in the data set. As you can see, left and right outer joins differ in which of the two tables you want the actual data set from: the first (left) or the second (right). Because you can switch them around at will, many people consistently use the left outer join for all outer joins.

CAUTION **Ask yourself whether you really need to be using outer joins. Because outer joins require less precision to format, inexperienced SQL users often perform an outer join and then filter the results in the code layer. This is wasteful and slow. Outer joins are all about the NULL values, which are not easily returned by inner joins. An example of a good use for an outer join is a report where you want to see which of your registered users had and had not downloaded your latest software product and how many times they had downloaded. If you are not in this situation, learn to use inner joins instead.**

Finally, there is something known as the *self-join*, which is a more advanced technique and won't really make a lot of sense with the example data set. It's often used with *denormalized data*, which means data that deliberately bends the rules of good SQL design (for example, never repeating any data point) for performance reasons (for example, to reduce the number of complex multitable joins).

If you need to make complex and frequent joins, this may constrain the brand of SQL database you can use, because not all of them support every type of join.

Subselects

Before we leave the realm of SELECT statements, we should mention the *subselect*. This is a statement such as:

```
SELECT phone_number
FROM table
WHERE name = (SELECT name FROM table2 WHERE ID = 1);
```

Subselects are more of a convenience than a necessity. They can be very handy if you're working with enormous batches of data, but you can get the same result with two simpler SELECTs. The subselect is faster if the subselect clause returns a large data set, but there are cases where two selects will not appreciably affect performance.

INSERT

Of course, no matter how many SELECT queries you write, all is for naught if you haven't put any information in the database to begin with. The command you need to put new data into a database is INSERT. The basic syntax is:

```
INSERT INTO table (col1, col2, col3) VALUES(val1, val2, val3);
```

Obviously, the columns and their values need to match up; if you mix up your array items, nothing good will happen. If some of the rows will not have values for some of the fields, you will need to use an empty, null, or auto-incremented value — and, at a deeper level, you may need to have ensured beforehand that fields can be nullable or auto-incrementable. If this is not possible, you should simply leave out any columns you wish to default to an empty value in an INSERT statement.

A twist on the basic INSERT is the INSERT INTO...SELECT. This just means you can INSERT the results of a SELECT statement:

```
INSERT INTO customer(birthmonth, birthflower, birthstone) SELECT * FROM
birthday_info WHERE birthmonth = $birthmonth;
```

Not every SQL server has this capability. Also, you need to be careful with this command because you can cause problems for yourself quite easily — for instance you can overwrite data or experience locking issues. In general, it's not a good idea to select from the same database you're inserting into.

UPDATE

UPDATE is used to edit information already in the database, without deleting any rows. In other words, you can selectively change some information without having to delete an entire old record and insert a new one. The syntax is:

```
UPDATE table
SET field1='val1', field2='val2', field3='val3'
WHERE condition;
```

The conditional statement is just like a SELECT condition, such as WHERE ID>15 AND ID<21 or WHERE gender='F'.

DELETE

DELETE is pretty self-explanatory: You use it to delete the contents of one or more fields permanently from the database. The syntax is:

```
DELETE datapoint
FROM table
WHERE condition;
```

The most important thing to remember is the condition — if you don't set one, you will delete every entry in the specified columns from the database, without a confirmation or a second chance in many cases!

CAUTION **Let us reemphasize: you *must* remember to use a condition *every single time* you** UPDATE **or** DELETE**! If you do not, every single row in the table will experience the same alteration or deletion. Even very experienced programmers have forgotten to include the condition, to their vast professional embarrassment. You should also give a good deal of thought to restricting database permissions so the minimum number of people can perform these potentially dangerous operations. I'll usually jump ahead and write the beginnings of the** WHERE **condition before filling out the rest of the** DELETE FROM **portion of the statement, just to make sure I don't inadvertently delete the entire table's worth of data. Another tip is to use the** limit **keyword within the** DELETE **statement so as only to delete the number of rows specified in the limit.**

Database Design

As should be obvious from the previous section, learning to use a SQL database isn't exactly rocket science — you can get a lot done with just a few simple commands. The hard part is designing the database in the first place and, of course, operating it in the real world over time. Not every web developer will be asked to design a schema in a professional context, but it never hurts to know how.

At the most fundamental level, database design can be broken down into the following mantra:

```
One to one,
One to many,
Many to many,
Many to one;
And always use a unique ID.
```

An example of *one-to-one data* for Americans is the Social Security number (other nations probably have similar identification cards with unique numbers). Each U.S. citizen has only one unique identifier; it is, in fact, a crime to use the Social Security number of another individual or apply for more than one number. Database designers seize upon truly unique identifiers such as this because almost every other piece of personal information is subject to change — which accounts for the large number of businesses who inappropriately use the Social Security number for identification purposes.

One-to-many data and *many-to-one* are the same, differing only in how the columns are placed in a database. An example of one-to-many data comes from the medical realm: patients to visits. Each patient will always be a discrete individual but may have any number of visits to the doctor. If you designed the table to represent visits to patients, it would instantly become many-to-one data.

Finally, *many-to-many data* is well represented by the relationship of authors to books. Not only can a given book have multiple authors, but each author may have written or coauthored many books.

This is not a matrix of relationships that would be easy to represent efficiently in a spreadsheet, but it is precisely this category of data at which relational databases most excel.

Every data relationship falls into one of these categories. As a database designer, it's your job to decide which one of these represents what you need to know in the way you need to know it.

This is not as trivial as it sounds. Imagine that you want to develop a database of movie information. One decision you might have to make is whether *movie* and *title* are in a one-to-one relationship with each other, or whether enough films have alternate titles to merit an *alternate title* field or even a one-to-many representation. There's no right answer here — the decision depends on exactly how the information will be used, how large the database will be, if the extra resources required to maintain a more precise data structure are worth the cost, and whether there's a better-than-even chance that today's tangential trivia will become tomorrow's crucial discovery. Some people may be surprised to learn that archiving information can be as much about ruthless excluding as about careful hoarding. As historians say, history is about forgetting as much as it is about remembering.

The simplest relationship is the one-to-one because you can group all these fields into a single table that can be searched more quickly. For instance, a table holding customer information might contain the following fields:

```
Customer ID
Customer name
Administrative contact
Technical contact
```

The hardest thing about the one-to-one relationship is definitively deciding that you will never need to make it into a one-to-many relationship. For instance, what if your biggest customer decides it wants to designate two technical contacts?

As soon as you have a one-to-many, many-to-one, or many-to-many relationship, you're looking at going from a single table to multiple tables: one each for the main variables and one stating the relationship. Tables 13-4 through 13-6 show a common example of a many-to-many relationship:

TABLE 13-4

Customer

Customer_id	Name
1	Acme Bread
2	Baker Construction
3	Coolee Dam

TABLE 13-5

Interactions

Interaction_id	Type
1	Phone-support incident
2	On-site incident
3	Written complaint
4	Phone complaint
5	Kudo

TABLE 13-6

Customer-Interaction

Customer-interaction_id	Customer_id	Interaction_id
1	1	1
2	3	5
3	2	4
4	2	3
5	1	2

After you've decided on a database design, the mechanical details of constructing the database are minimal. The main data structure statements of SQL are CREATE, ALTER, and DROP.

CREATE is used to make a completely new table. All the work is in defining the columns of each table. First, you declare the name of the table, and then you must detail the specific data types of that table's columns in what is called a *create definition*. A CREATE statement will take this form:

```
CREATE TABLE tablename (
  id_col INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  col1 TEXT NULL INDEX,
  col2 DATE NOT NULL
);
```

203

Different SQL Servers have slightly different data types and definition options, so the syntax of one may not transfer exactly to another. For instance, Oracle databases do not auto-increment; to get a new value, you must generally call a function.

DROP can be used to completely delete a table and all its associated data. It's not the most subtle command:

```
DROP TABLE tablename;
```

Obviously, you need to be very careful with this statement.

ALTER is the way to change a table's structure. You simply indicate which table you're changing and redefine its specs. Again, SQL products differ in functionality here. The ALTER statement usually takes this form:

```
ALTER TABLE table RENAME AS new_table;
ALTER TABLE new_table ADD COLUMN col3 VARCHAR(50);
ALTER TABLE new_table DROP COLUMN col2;
```

Privileges and Security

As we state in Chapter 28, security online is analogous to security in the real world. Any cop will tell you that you cannot make your home absolutely crime-proof. A more realistic goal is to increase the difficulty and risk to a level where a large percentage of intruders will choose to go to an easier target down the block.

Setting database permissions

The most fundamental rule of database use (of any computer security, really) is to give each user or group only the minimum permissions necessary to do what needs to be done. Besides the threat of malicious/experimental outsiders, setting the correct permissions can protect you from your coworkers and yourself. Insiders have been known to cause massive problems through disgruntlement, ignorance, momentary brain freeze, or a combination of motives. You do not want to have to cope with the consequences of a fired employee's parting shot or a new intern trying out the DROP database command just to see what happens.

A typical database permissions package might be something like:

- **Web visitor:** SELECT only
- **Contributor:** SELECT, INSERT, and maybe UPDATE
- **Editor:** SELECT, INSERT, UPDATE, and maybe DELETE and maybe GRANT
- **Database Administrator:** SELECT, INSERT, UPDATE, DELETE, GRANT, and DROP

`DROP` in particular is the nuclear bomb of SQL because it allows you to blow away an entire table or database with a single command. Someone's got to have the ability, but heavy lies the tiara of responsibility on the head of the root database user. Use the power wisely, grasshopper.

In many databases, including MySQL, passwords are encrypted using a different algorithm from system passwords (and, of course, they are typically stored in entirely different locations). Even if one is cracked, the other is not necessarily vulnerable. This assumes that you take the time to set permissions correctly, pick good passwords, and usually employ a special command to insert usernames and passwords correctly into the grant table (as opposed to inserting them like other data).

CAUTION   Database usernames and passwords should not be identical to system usernames and passwords.

CROSS-REF   Chapter 14 covers permissions for MySQL specifically.

Keep database passwords outside the web area

It's a good idea to separate passwords from the web pages that use them. With PHP's `include()/include_once()` and `require()/require_once()` functions, it's very easy to drop in text (such as database passwords) from another file at runtime. Remember that these included files do not have to be in a PHP or web server–enabled directory! Whenever possible, keep them somewhere outside your web area or the file hierarchy viewable to the public through the web server. A good example is a directory above or outside of your web document root or in a home directory.

Taking the database variables out of PHP files is also good for other reasons. If you have many PHP scripts using the same database, they can all use the same password file. When you suspect the password has been compromised, or when you change the password on a regular schedule, you need only alter one script for all the files to be updated.

The unavoidable downside of this technique is that the file must be *readable* by the user through which the web server runs, such as wwwuser, httpd, or Apache. This usually involves changing the ownership of the file with the database credentials to that of the Apache web server user, and, of course, making sure that the mode of the file doesn't allow it to be world-readable.

If you have a set of database variables you use infrequently — a configuration script or the like — you can keep it in a non-Apache-readable directory and change the permissions only on the rare occasions necessary. We infrequently have to go to the trouble to delete postings from our sites' forums. So it's not that much more work (and much more secure) to keep this file in a non-Apache-user-owned directory, once in awhile change the permissions just long enough to delete the offending post, and then immediately change everything back.

If for whatever reason, you decide to put your database username, password, hostname, and database name into a PHP script in plain text, this is what you can expect. If the web server is functioning normally, the database passwords should be as safe as any file on that server. But if the daemon goes down, there is some chance your raw PHP (including plain-text database variables) will be delivered in a human-readable form. You can reduce this risk by avoiding the use of the `.html` suffix for PHP files.

205

In some versions of PHP, if database connectivity went down and you hadn't specified silent mode, you would see something like the following:

```
Warning: MySQL Connection Failed: Access denied for user:
'someuser@localhost' (Using password: NO) in
/home/web/html/mysqltest.php3
on line 2
```

This constitutes a security breach, because it reveals your MySQL username and whether or not you use a password. From PHP4 forward, MySQL error messages are no longer displayed by default. Two functions, `mysql_errno()` and `mysql_error()`, allow you to opt for error codes or text warnings — but now you have to deliberately choose to ask for the information. Because, in most cases, you can opt for the more configurable `die()` instead or remove error messages after debugging, it's still not a good idea to use `mysql_error` on a public production server unless you scrupulously send messages to error logs using the `error_log()` function rather than to standard output.

Learn to make backups

And finally, the biggest part of database security may be backing up. Take an hour to learn the best way to back up data in your particular database (for example, via the `mysqldump` command in MySQL), and then schedule regular backups right away. Even better, with a little foresight you can also set up an automatic database backup schedule.

Summary

SQL is not rocket science. The four basic data-manipulation statements supported by essentially all SQL databases are `SELECT`, `INSERT`, `UPDATE`, and `DELETE`. `SELECT` gets data out of the database, `INSERT` puts in a new entry, `UPDATE` edits pieces of the entry in place, and `DELETE` gets rid of an entry.

Designing databases is where most of the difficulty lies. Not all web developers will be asked to do this. The designer must think long and hard about the best way to represent each piece of data and relationship for the intended use. Well-designed databases are a pleasure to program with, while poorly designed ones can leave you pulling your hair out while contemplating numerous connections and icky joins.

SQL databases are created by so-called data structure statements. The most important of these are `CREATE`, `ALTER`, and `DROP`. As one would expect, `CREATE TABLE` defines a new table within a database. `ALTER` changes the structure of a table. `DROP` is the nuclear bomb of SQL commands because it deletes entire tables or sometimes even whole databases.

Chapter 14

Learning Database Administration and Design

MySQL is one of the easiest databases to administer on all platforms, and because it's so lightweight, it can run on even low-powered PCs. Thus, PHP developers have long found it convenient to throw a copy of MySQL on client machines — even on laptops — for a complete local web development environment. Many developers learn to run their own MySQL installations so that they can work at home or on the road, using the OS of their choice. Work teams also sometimes prefer developers to each use a separate local MySQL installation, so that there is no single point of failure that could affect an entire development group. And many PHP-based open source projects assume complete familiarity with MySQL database administration for all developers.

Unlike some other databases, it should be well within the capability of any PHP developer to self-administer a MySQL database. There is a plethora of tools, both in MySQL itself and available from third parties, to make this job even easier. Many PHP-based application packages, both commercial and open source, also require familiarity with a MySQL database to install, run, and debug the web app. So even if you don't plan to write all your PHP code yourself, getting comfortable with MySQL administration will pay many dividends.

Basic MySQL Client Commands

It may surprise you to know that the binary named `mysql` in your `mysql/bin` directory is not the server, but the client (the server is `mysqld`). When you type `mysql` into a shell, you are using the MySQL command-line client to access some MySQL server.

To connect to the MySQL server using the command-line client, the basic command is:

```
mysql [-h hostname] [-P portnumber] -u username -p
```

You almost certainly need to pass the username; if you don't, the client will try the name of your shell user. If you don't pass the password flag, `mysql` will check whether a password is needed for the user you claim to be — and if so, it will reject you. If you're connecting to a local host, you don't need the hostname flag; if you're connecting to the default port (3306), you don't need the port number flag. There are a bunch of other options, but usually this is all you need the first time. Assuming that you use the username `root`, you will be prompted for the root password that you just set in the previous step.

At this point, you will need to select a database to use. The command for that is:

```
USE databasename;
```

The semicolon is optional for this command, but you need one for every other SQL command, so you might as well get used to using it. Until you create new databases, there are only two databases in a fresh install: `mysql` and `test`. If you just connected to MySQL as the root user, you have access to both; if you are connected as any other user, you have access only to `test`.

The command `SHOW TABLES;` will dump a list of all the tables in this database.

To quickly see the structure of a database table, use `SHOW COLUMNS FROM tablename;`. This displays all the columns with their types, sizes, default values, and other helpful information.

To see all the values in a table, just do a `SELECT` with unrestrictive conditions:

```
SELECT * FROM tablename;
```

Be careful though, since in live databases this kind of query can be huge and take up a lot of resources. If you have reason to suspect that the data set is more than a few rows, you should take steps to limit the query.

CROSS-REF See Chapter 13 for more information on how to write SQL statements such as `SELECT`, `INSERT`, **and so forth. Remember that one of the best ways of debugging problems with SQL statements in your PHP code is to try them out (with suitable fake data plugged into the variables) using the MySQL command-line client rather than the PHP client. See Chapter 19 for more information on debugging SQL in your PHP.**

Finally, to get out of the MySQL client session, use the command `quit;`. Again, the semicolon is optional for this command. This should drop you back into your normal shell.

MySQL User Administration

A big part of using MySQL safely and effectively is understanding its privilege system and learning how to use the tools provided for controlling user privileges.

MySQL allows you to grant quite fine-grained permissions to different users from different client locations. There are four descending levels of privileges: global, database, table, and column. So in theory, you could allow a particular user to write data only to certain columns of certain tables of certain databases on your MySQL server. Or you could just as easily give any database user connecting from anywhere the same powers as the root database user (although this is totally not recommended).

Of course, for security reasons it's generally a good rule of thumb to grant each user only the minimal permissions necessary to perform his or her function.

There are two different ways to add or edit user permissions in MySQL (assuming that you're the root database user): by direct SQL statements (for example, putting a Y by hand into every relevant field of every relevant grant table) or by use of the GRANT and REVOKE syntax. The latter is easier, and less dangerous if you make a small mistake, since in most cases your query will choke with a SQL error instead of just leaving a gaping security hole.

To add a new MySQL user, type the following:

```
GRANT priv_type [(column1, column2, column3)]
ON database.[table]
TO user@host IDENTIFIED BY 'new_password';
```

where columns and tables are optional and additional `priv_types` can be appended in a comma-separated list.

The types of privileges and their scope are shown in Table 14-1.

Obviously, there's no point in trying to give anyone the SHUTDOWN privilege at the table level. You will merely get an error message referring you to the manual. If you grant ALL to a column, table, or database, the user will get only the basket of privileges appropriate to that level.

You should be especially careful about giving users the following privileges, which are all dangerous: GRANT, ALTER, CREATE, DROP, FILE, SHUTDOWN, PROCESS. No normal database user, especially a PHP user, should need these permissions in production.

The syntax for revoking privileges is very similar, although simpler:

```
REVOKE priv_type [(column1, column2, column3)]
ON database[.table]
FROM user@host;
```

MySQL Privilege Scope for Selected Privileges

Privilege	Global	Database	Table	Column
ALL	✓			
ALTER	✓	✓	✓	
CREATE	✓	✓	✓	
CREATE TEMPORARY TABLE	✓	✓	✓	
DELETE	✓	✓	✓	
DROP	✓	✓	✓	
EXECUTE	✓	✓		
FILE	✓	✓		
INDEX	✓	✓	✓	
INSERT	✓	✓	✓	✓
LOCK TABLES	✓	✓		
PROCESS	✓	✓		
REFERENCES	✓	✓		
RELOAD	✓	✓		
REPLICATION CLIENT	✓			
REPLICATION SLAVE	✓			
SELECT	✓	✓	✓	✓
SHOW DATABASES	✓			
SHUTDOWN	✓			
SUPER	✓	✓		
UPDATE	✓	✓	✓	✓
USAGE	✓	✓		
GRANT OPTION	✓	✓	✓	

After you grant or revoke privileges to any user, you need to force the database to reload the new privilege data into memory. You do this by issuing the `FLUSH PRIVILEGES` command. You could also start and stop the server, but that's impractical in many circumstances.

This is all well and good, but by now you're probably thinking: But what actual permissions should I actually grant to my actual PHP user? Let's look at some common cases from the real world.

Local development

For purely local stuff, especially on a machine that isn't connected to the Internet all the time or is tucked securely behind a good firewall, almost anything goes. If you need to experiment with your schema, this is the place to do it — so it's appropriate to have permissions like `ALTER`, `CREATE`, `DELETE`, and `DROP` in addition to the normal `SELECT`, `INSERT`, `UPDATE`. A lot of people will find it convenient to just grant `ALL PRIVILEGES` on a certain database to a local user, like this:

```
GRANT ALL PRIVILEGES on database.*
TO username@localhost
IDENTIFIED BY 'password';
```

Standalone web site

A self-hosted database probably needs to accept connections from numerous web servers in the same domain. In production, all machines should be limited to `SELECT`, `INSERT`, `UPDATE`, and possibly `DELETE` — although many systems never actually delete data, and it's a little safer not to do so. Since there probably won't be multiple databases on a standalone web site's production database, global permissions are faster with not much more real security risk. So a possible grant statement is:

```
GRANT SELECT, INSERT, UPDATE ON *.*
TO phpdbuser@%.example.com
IDENTIFIED BY 'password';
```

However, this is the situation that is most likely to use master-slave replication. Often, these MySQL clusters are configured so that all writes go to the master, while the slaves do nothing but serve up very fast reads. In that case, you would give only `SELECT` privileges on each slave and only `INSERT` and `UPDATE` privileges on the master — possibly to two different database users.

Shared-hosting web site

If you are an Internet service provider (ISP) that offers shared hosting, or a customer hosting your web site with one, your primary concern should be security over performance. Under no circumstances do you want one user to be able to tamper with or delete data belonging to another user.

Unless each user has her own MySQL instance running on her own port, the ISP administrator should not allow users to create or drop globally. Obviously, though, there is no good way to deny table creates or drops, which implies that each user will also be able to drop his own database if he so desires. Yes, that's right: If your users can define new tables, as they almost certainly will have to

in this situation, there's no good way to prevent them from blowing away all their data with a single command! That's part of the easy come, easy go thrill of MySQL. The database administrator can and should, however, prevent users from being able to do this to other users on the same server.

Backups

Database backups can be made in two ways: by copying the data directory directly (either manually or by means of the `mysqlhotcopy` script on Unix) or by using the `mysqldump` tool to write out a SQL file that will replicate your database. The former is a little faster, but the latter is more flexible. With `mysqldump` you can choose to copy just the structure of the database, just the data, or both.

The most basic usage of `mysqldump` is:

```
mysqldump -u username -p databasename > dumpfilename.sql
```

This command will dump a text file that can be read into another database server, like this:

```
mysql -u root -p databasename < dumpfilename.sql
```

Instead of directing the output of `mysqldump` to a file, you can also pipe it directly to another server, like this:

```
mysqldump -u username -p databasename |
mysql -h remote-host -u remoteuser -p -C databasename
```

However, this can be less secure in some cases, since you have to tell the remote host to accept database-modifying connections from external clients.

This basic command is fine as far as it goes — meaning that it will result in a nice SQL file containing both the structure and data of the named database. But sometimes you will want something more specific than that: maybe just the structure or just the data or all the databases on that server or just some tables from your chosen database. MySQL allows you to both specify different combinations of databases and/or tables and to add option flags to your command.

If you want to select specific tables to dump from your chosen database, just list them after the database name:

```
mysqldump -u username -p databasename table1 table2
> dumpfilename.sql
```

If you want to dump some but not all databases on your server, use the `--databases` flag and then list the databases. However, in this case, you will not be able to specify tables.

```
mysqldump -u username -p --databases database1 database2 >
dumpfilename.sql
```

If you want to dump all databases, use the `--all-databases` flag:

```
mysqldump -u username -p --all-databases > dumpfilename.sql
```

You can specify any of these options before specifying the databases and tables. There are many `mysqldump` options, but Table 14-2 lists the most commonly used options.

TABLE 14-2

mysqldump Options

Option	Explanation
`--add-locks`	Adds table locking to SQL file for faster inserts on the target table. See also `--opt`.
`--add-drop-table`	Will overwrite each table definition. Be careful with this option, as you could delete data! If you don't use this option but a table of the same name already exists, you will get an error on the target database.
`-a, --all`	All options. Be careful!
`-c, --complete-insert`	Use more complete insert statements with column names, instead of simply reading in values.
`--help`	Displays help message with options.
`-l, --lock-tables`	Locks tables on the source machine before the dump.
`-n, --no-create-db`	Will not create databases of the specified names if they don't exist already. Default with the `--databases` and `--all-databases` options.
`-t, --no-create-info`	Will not create tables of the specified names if they don't exist already.
`-d, --no-data`	Just the structure of the specified database(s) or tables.
`--opt`	Equal to `--quick --add-drop-table --add-locks --extended-insert --lock-tables`. Fastest possible dump. Make sure that you want to drop existing tables if there's a conflict.
`-q, --quick`	No buffering.
`-r, --result-file=filename`	Dump result to file. In DOS, creates Unix-style line breaks.
`-w, --where='condition'`	Select results by the WHERE clause in single quotes.

Because `mysqldump` is so easy to use, you should have no excuse for not adhering to a regular backup schedule. This is why cronjobs were invented! If your data changes relatively infrequently, you might be able to get away with weekly or fortnightly backups; if you have a fairly high-traffic site, you'll want to schedule one every night.

Users of PHPMyAdmin have access to `mysqldump` through the Export tab. However, PHPMyAdmin currently offers only the most common options for your data dump. If you need more control over the format of your SQL file, you'll have to use `mysqldump` as previously described instead.

Replication

MySQL replication is based on a one-way single-master, single-or-multiple-slave model. The master database will handle all writes — meaning all `INSERT`s, `UPDATE`s, and `DELETE`s, as well as all schema changes. The slaves will periodically get these changes from the master and in the meantime will be available for highly optimized read-only data serving (meaning all `SELECT`s). The master does not know anything about slave databases. It simply makes its binary logs available, and the slaves do all the rest: scheduling updates, connecting to the master, getting the changes, applying the changes, and so on. Thus, slaves are aware of the identity of the master, but masters are not aware of the identities of slaves.

If the master database goes down for any reason, no replacement will be automatically elected. The entire system is likely to become unresponsive, as the slaves spend many resources trying in vain to connect to the master for updates, while PHP tries to perform writes without success. The database administrator will have to manually break the existing master-slave relationships and designate a new master by hand. Luckily, if something goes wrong with the master, there's no way the slaves will have gotten out of sync — so if a database administrator notices the problem and is available to deal with it, changing to a new master database should be relatively quick.

Because there have been many changes and upgrades to the replication function in recent versions of MySQL, many recent versions are incompatible with other recent versions in a replication setup. If you want to try replication, we recommend that you make sure all the database servers involved are using the same version of MySQL, and furthermore, that this version is 4.0.3+. If you are trying to replicate with disparate versions of MySQL between 3.23 and 4.0.3, it is very likely that things will not work properly.

In a nutshell, the operations that must be performed to establish MySQL replication are:

1. Grant permissions to slave user on master.
2. Take snapshot of master data; copy to slave machines.
3. Shut down MySQL servers.
4. Restart MySQL servers with correct `server-ids`.
5. Establish master-slave relationship from each slave.

Now we'll explain the process in more detail.

You will need to create an account on the master database for slaves to use, with the `REPLICATE SLAVE` privilege. You do not need to grant any other privileges to this account.

```
GRANT REPLICATE SLAVE ON *.*
TO replicant@'%' IDENTIFIED BY 'replpwd';
```

Next, lock the master server and take a snapshot of its state immediately before the replication. On the master server, log in to a MySQL client session as the root user and issue the commands:

```
FLUSH TABLES WITH READ LOCK;
SHOW MASTER STATUS;
```

This will prevent any changes from being made to the database until you are ready to bring up the cluster. You may also (depending on whether this server has been run with binary logging) see some data about the location of the binary log file and offset. If so, write it down; if not, use the default values `''` (empty string) and 4, respectively.

Next, copy the master database structure and data. There are two ways to do this. The first is to simply copy the `mysql/data` directory into a tarball or zip file by using one of these commands or a GUI procedure:

```
tar -cvf master_snapshot.tar data/
zip master_snapshot.zip data/
```

Alternatively, you can use `mysqldump` to make a backup as described in the next section. Copy this snapshot file to each slave server.

Now shut down all the master and slave servers. Quit any `mysql` client shell sessions, and issue the command:

```
mysqladmin -u root -p shutdown
```

on each server. The reason you are shutting the servers down is to give them unique `server-id` values. They will use these values to find each other when they establish the master-slave relationship. This value is set in each server's `my.cnf` file and will be read in on startup. On Windows, the `my.cnf` file is located in one of two places: `C:\my.cnf` or `C:\[Windows directory]\my.ini`. On Unix systems, the global `my.cnf` file is found in `/etc/my.cnf` and the server-specific file (which is probably the one you want to use) is found in `/path/to/mysql/data/my.cnf`.

First, set the `server-id` on the master machine. Find or create a file called `my.cnf` in the proper location for your platform, and make sure that it contains the lines:

```
[mysqld]
log-bin
server-id=1
```

Restart the master server:

```
bin/mysqld_safe --user=mysql
```

In each slave server's `my.cnf` files, you need only the `server-id`, not the `log-bin` line. The most important thing is that you are absolutely positive that all the `server-id` values in your cluster are unique! If they are not, bad things will happen. So the first slave's `my.cnf` file would contain this line:

```
[mysqld]
server-id=2
```

The second slave would set `server-id=3`, and so forth.

Now, before you bring up each slave server, you may need to do a little bit of housekeeping. If this MySQL server has been used as a slave before, you may want to delete the files `data/master.info` and `data/relay-log.info`. You may also want to delete the `.err` and `.pid` files in the `data` directory. Also, if you copied the master's data snapshot into a tarball or zipfile, now is the time to copy it to the slave with a command like one of these (from the `mysql` directory):

```
tar -xvf master_snapshot.tar
unzip master_snapshot.zip
```

If you used `mysqldump` instead, you have to wait until the server is back up.

Now bring up the slave:

```
bin/mysqld_safe --user=mysql --skip-slave-start --log-warnings
```

If you took your master data snapshot with `mysqldump`, now is the time to apply the SQL file to the slave:

```
mysql -u root -p databasename < master_snapshot.sql
```

Finally, you will establish the master-slave relationship. Log in to a `mysql` shell and then enter the following commands, substituting the values you wrote down at the beginning of the process:

```
CHANGE MASTER TO
    MASTER_HOST='masterhostname',
    MASTER_USER='replicant',
    MASTER_PASSWORD='replpwd',
    MASTER_LOG_FILE='',
    MASTER_LOG_POS=4;
START SLAVE;
```

If there are problems, they will appear in the slave machine's error log.

Recovery

Normally, MySQL does not require much attention. MySQL servers have happily puttered away for months if not years with minimal administration. However, bad things do happen to data: Hard disks melt down, hosting centers lose power suddenly, and human error is a constant and awful probability. If you have insufficient memory for all the applications you're running on a server, or insufficient disk space on a partition, you may also get an error that requires a recovery process. It must be admitted that MySQL seems to have minor database corruption events with greater frequency than heavier-weight databases — or perhaps it's just easier for the administrator to notice these events.

Luckily, MySQL is designed to make it amazingly easy to repair small flaws in your data and get back up quickly. Only once have we had to actually scrap an entire database after repeated attempts at recovery, and that disaster was caused by a total hard disk failure, which is something a developer can do nothing to plan for or recover gracefully from — except make frequent backups.

MySQL has long shipped with a command-line tool called `myisamchk` for checking and repairing tables. This was a fine script but it suffered from one flaw: It could be run effectively only when the database was shut down. That's fine when you're actually recovering from a disaster, since you're unlikely to be able to start your database anyway, but it's a significant barrier to trying to head off problems by regularly checking your data tables. Luckily, there is now a new tool that can be used during operation — `mysqlcheck`. You can continue to use `myisamchk` (used only for `myisam` tables) when the server is not running. Refer to the MySQL manual for more information on trouble-shooting table problems.

Both these tools basically can do three things: check a MyISAM table for errors, repair problems, and optimize the database. The syntax by which you use the scripts is different, however.

myisamchk

The `myisamchk` utility is invoked like this:

```
myisamchk [options] table_name
```

or

```
myisamchk [options] /path/to/mysql/data/database/table.MYI
```

You can wildcard both database directories and table names with an asterisk, which is more common than specifying a table, since you usually don't know exactly which table is causing the problems. Use the following commands to check all the tables of all the databases on a server:

```
myisamchk [options] /path/to/mysql/data/*/*.MYI
myisamchk [options] /path/to/mysql/data/*/*.MYD
```

.MYI extensions designate index files, and .MYD extensions designate data files — you need to check both.

With no option flags, `myisamchk` will simply check the designated table. If you pass the `-r` option flag, `myisamchk` will repair the designated tables. You can also check and repair any corrupted tables in a single operation:

```
myisamchk --silent --force --fast --update-state -O
key_buffer=64M -O sort_buffer=64M -O read_buffer=1 -O
write_buffer=1M /path/to/mysql/data/*/*.MYI
```

The command `myisamchk -r tablename` will also optimize a table that has been fragmented by deletes and updates.

mysqlcheck

The `mysqlcheck` tool has several handy advantages over `myisamchk`. As previously mentioned, it can be used while the server is running — even while serving up queries. It works on databases rather than tables, using the same syntax as the `mysqldump` tool. And instead of having to remember the meaning of a bunch of option flags, you can copy and rename the executable to get different behaviors.

The `mysqlcheck` tool is invoked in one of these ways:

```
mysqlcheck [options] databasename table1 table2 table3
mysqlcheck [options] --databases database1 database2
mysqlcheck [options] --all-databases
```

To repair, analyze, or optimize databases, you simply copy the `mysqlcheck` file and change its name to `mysqlrepair`, `mysqlanalyze`, or `mysqloptimize` — and then invoke it the same way. So, for instance, to repair all the databases on your server, you might give this command:

```
mysqlrepair -u root -p --all-databases
```

MySQL AB recommends that you set up a regular schedule of data file checking via cronjob, plus run one of these utilities every time you start up your MySQL server. This should help keep your data compact for fast reads, head off problems while they're still tiny, and minimize your chances of a database problem that is visible to your users.

Summary

MySQL is one of the easiest databases to administer, and learning to do so provides many benefits to PHP developers. MySQL installations have become easier of late on many platforms, and there are GUI as well as command-line tools available to help you view the structure of your database, manage database users, and make backups. More advanced MySQL administration tasks include disaster recovery and replication — both of which are probably as easy to accomplish on MySQL as they could possibly be made. However, even long-time MySQL users should consider the impact of recent changes to the MySQL-PHP relationship: licensing issues, client-version incompatibility, the new `mysql` extension, and transactions.

Chapter 15

Integrating PHP and MySQL

A fter you've installed and set up your MySQL database, you can begin to write PHP scripts that interact with it. Here, we will try to explain all the basic functions that enable you to pass data back and forth from web site to database.

> **NOTE** Information related to creating a MySQL database is at the end of this chapter, because it is a more advanced skill that builds on the fundamental MySQL skills discussed in the earlier parts of the chapter.

Connecting to MySQL

The basic command to initiate a MySQL connection is

```
mysql_connect($hostname, $user, $password);
```

if you're using variables, or

```
mysql_connect('localhost', 'root', 'sesame');
```

if you're using literal strings.

The password is optional, depending on whether this particular database user requires one (it's a good idea). If not, just leave that variable off. You can also specify a port and socket for the server ($hostname:port:socket), but unless you've specifically chosen a nonstandard port and socket, there's little to gain by doing so.

The corresponding `mysqli` function is `mysqli_connect`, which adds a fourth parameter allowing you to select a database in the same function you use to connect. The function `mysqli_select_db` exists, but you'll need it only if you want to use multiple databases on the same connection.

You do not need to establish a new connection each time you want to query the database in the same script. You will need to run this function again, however, for each script that interacts with the database in some fashion.

Next, you'll want to choose a database to work on:

```
mysql_select_db($database);
```

if you're using variables, or

```
mysql_select_db('phpbook');
```

if you're using a literal string.

> **TIP** You will sometimes see these two functions used with an @ prepended, such as `@mysql_ select_db($database)`. **This symbol denotes *silent mode*, meaning the function will not return any message on failure, as a security precaution. You should have** `display_errors` **set to** `off` **on production servers anyway.**

You must select a database each time you make a connection, which means at least once per page or every time you change databases. Otherwise, you'll get a Database not selected error. Even if you've created only one database per daemon, you must do this, because MySQL also comes with default databases (called mysql and test) you might not be taking into account.

You may find it convenient to group all your connection information into a custom connect function and put it someplace where you can access it from all your scripts, such as the php includes directory, or in the case of a virtual server, a site-specific include file. This function might look like the following:

```
// Connect to a single db
function qdbconn() {
  $dbUser = "myuser";
  $dbPass = "mypassword";
  $dbName = "mydatabase";
  $dbHost = "myhost";
  if (!($link=mysql_connect($dbHost, $dbUser, $dbPass))) {
    error_log(mysql_error(), 3, "/tmp/phplog.err");
  }
  if (!mysql_select_db($dbName, $link)) {
    error_log(mysql_error(), 3, "/tmp/phplog.err");
  }
}
```

If you like, you could extend this function by creating links (for example, $link1, $link2) to multiple databases on the same server. This code also records a MySQL error message in the PHP error log.

Now that you've established a connection to a specific database, you're ready to make a query.

Making MySQL Queries

A database query from PHP is basically a MySQL command wrapped up in a tiny PHP function called `mysql_query()`. This is where you use the basic SQL workhorses of `SELECT`, `INSERT`, `UPDATE`, and `DELETE` that we discussed in Chapter 13. The MySQL commands to `CREATE` or `DROP` a table can also be used with this PHP function if you do not wish to make your databases using the MySQL client.

You could write a query in the simplest possible way, as follows:

```
mysql_query("SELECT Surname FROM personal_info WHERE ID < 10");
```

PHP would dutifully try to execute it. However, there are very good reasons to split up this and similar commands into two lines with extra variables, like this:

```
$query = "SELECT Surname FROM personal_info WHERE ID < 10";
$result = mysql_query($query);
```

The main rationale is that the extra variable gives you a handle on an extremely valuable piece of information. Every MySQL query gives you a receipt whether you succeed or not — sort of like a cash machine when you try to withdraw money. If things go well, you hardly need or notice the receipt — you can throw it away without a qualm. But if a problem occurs, the receipt will give you a clue as to what might have gone wrong, similar to the "Is the machine not dispensing or is your account overdrawn?" type of message that might be printed on your ATM receipt.

Another advantage of assigning the query string to a variable is that you can more easily view the query if you run into an error. Of course, you would accomplish this by writing the variable out to an error log — never by dumping it out to the browser in production!

The function `mysql_query` takes as arguments the query string (which should not have a semicolon within the double quotation marks) and optionally a link identifier. Unless you have multiple connections, you don't need the link identifier. It returns a `TRUE` (nonzero) integer value if the query was executed successfully *even if no rows were affected.* It returns a `FALSE` integer if the query was illegal or not properly executed for some other reason.

For purposes of this chapter, we've left the link identifier off; however, if you need to use multiple databases in your script, you can use code like this:

```
$query = "SELECT Surname FROM personal_info WHERE ID < 10";
$result = mysql_query($query, $link_1);
$query = "SELECT * FROM orders WHERE date > 20030702";
$result = mysql_query($query, $link_2);
```

As expected, the MySQL improved analog for this function is `mysqli_query`. It is very similar to its counterpart; however, the `link` and `query` parameters change places, and a third parameter allows you to specify a result flag indicating how PHP should handle the result.

If your query was an `INSERT`, `UPDATE`, `DELETE`, `CREATE TABLE`, or `DROP TABLE` and returned `TRUE`, you can now use `mysql_affected_rows` to see how many rows were changed by the query. This function optionally takes a link identifier, which is only necessary if you are using multiple connections. It *does not* take the result handle as an argument! You call the function like this, without a result handle:

```
$affected_rows = mysql_affected_rows();
```

If your query was a `SELECT` statement, you can use `mysql_num_rows($result)` to find out how many rows were returned by a successful `SELECT`.

The `mysqli_affected_rows` and `mysqli_num_rows` behave exactly the same as their `mysql_` counterparts.

TIP The `mysql_num_rows` **function can be useful in paginating large data sets returned by MySQL queries.**

Fetching Data Sets

One thing that often seems to temporarily stymie new PHP users is the whole concept of fetching data from PHP. It would be logical to assume that the result of a query would be the desired data, but that is not correct. As we discussed in the previous section, the result of a PHP query is an integer representing the success or failure or identity of the query.

What actually happens is that a `mysql_query()` command pulls the data out of the database and sends a receipt back to PHP reporting on the status of the operation. At this point, the data exists in a purgatory that is immediately accessible from neither MySQL nor PHP — you can think of it as a staging area of sorts. The data is there, but it's waiting for the commanding officer to give the order to deploy. It requires one of the `mysql_fetch` functions to make the data fully available to PHP.

The fetching functions are as follows:

- `mysql_fetch_row`: Returns row as an enumerated array
- `mysql_fetch_object`: Returns row as an object
- `mysql_fetch_array`: Returns row as an associative array
- `mysql_result`: Returns one cell of data

CAUTION **In our humble opinion, the functions** `mysql_fetch_field` **and** `mysql_fetch_lengths` **are misleadingly named. They both provide information** *about* **database entries rather than the entry values themselves. For instance, one might expect a function named** `mysql_fetch_field` **to be a quick way to fetch a single-field result set (the ID associated with a particular username, for instance), but that is not the case at all. The actual purpose of these functions is explained in Table 15-2 at the end of the chapter — but for the moment, the point is not to be misled into thinking that these functions will return database values.**

The differences among the three main fetching functions is small. The most general one is `mysql_fetch_row`, which can be used something like this:

```
$query = "SELECT ID, LastName, FirstName
          FROM users WHERE Status = 1";
$result = mysql_query($query);
while ($name_row = mysql_fetch_row($result)) {
  print("{$name_row[0]} {$name_row[1]} {$name_row[2]}<BR>\n");
}
```

This code will output the specified rows from the database, each line containing one row or the information associated with a unique ID (if any).

CAUTION **In an enumerated array, the integers in brackets are called *field offsets*. Remember that they always begin with the integer zero. If you start counting at 1, you will miss the value of your first column.**

The function `mysql_fetch_object` performs much the same task, except the row is returned as an object rather than an array. Obviously, this is helpful for those among the PHP brethren who utilize the object-oriented notation:

```
$query = "SELECT ID, LastName, FirstName
          FROM users WHERE Status = 1";
$result = mysql_query($query);
while ($row = mysql_fetch_object($result)) {
  echo "{$row->ID}, {$row->LastName}, {$row->FirstName}<BR>\n";
}
```

The most useful fetching function, `mysql_fetch_array`, offers the choice of results as an associative or an enumerated array — or both, which is the default. This means you can refer to outputs by database field name rather than number:

```
$query = "SELECT ID, LastName, FirstName
          FROM users WHERE Status = 1";
$result = mysql_query($query);
while ($row = mysql_fetch_array($result)) {
  echo "{$row['ID']}, {$row['LastName']}, {$row['FirstName']}<BR>\n";
}
```

Remember that `mysql_fetch_array` can *also* be used exactly the same way as `mysql_fetch_row` — with numerical identifiers rather than field names. By using this function, you leave yourself the option. If you want to specify offset or field name rather than making both available, you can do it like this:

```
$offset_row = mysql_fetch_array($result, MYSQL_NUM);
or
$associative_row = mysql_fetch_array($result, MYSQL_ASSOC);
```

It's also possible to use MYSQL_BOTH as the second value, but because that's the default, it's redundant.

In early versions of PHP, mysql_fetch_row was considered to be significantly faster than mysql_fetch_object and mysql_fetch_array, but this is no longer an issue, as the speed differences have become imperceptible. The PHP junta now recommends use of mysql_fetch_array over mysql_fetch_row because it offers increased functionality and choice at little cost in terms of programming difficulty, performance loss, or maintainability.

Last and least of the fetching functions is mysql_result(). You should only even *consider* using this function in situations where you are positive you need only one piece of data to be returned from MySQL. An example of its usage is:

```
$query = "SELECT count(*) FROM personal_info";
$db_result = mysql_query($query);
$datapoint = mysql_result($db_result, 0, 0);
```

The mysql_result function takes three arguments: *result identifier*, *row identifier*, and (optionally) *field*. Field can take the value of the field offset as above or its name as in an associative array ("Surname") or its MySQL field-dot-table name ("personal_info.Surname"). Use the offset if at all possible, as it is substantially faster than the other two. Even better, don't use this function with any frequency. A well-formed query will almost always return a specific result more efficiently.

CAUTION   You should never use mysql_result() to return information that is available to you through a predefined PHP-MySQL function. The classic no-no is inserting a row and then selecting out its ID number (extra demerits if you select on MAX(ID)!). Wicked bad style — use mysql_insert_id() instead.

All of the PHP functions for fetching MySQL data have identical mysqli counterparts. They take the same parameters and return comparable results.

A special MySQL function can be used with any of the fetching functions to more specifically designate the row number desired. This is mysql_data_seek, which takes as arguments the result identifier and a row number and moves the internal row pointer to that row of the data set. The most common use of this function is to reiterate through a result set from the beginning by resetting the row number to zero, similar to an array reset. This obviates another expensive database call to get data you already have sitting around on the PHP side. Here's an example of using mysql_data_seek():

```
<?php
echo("<TABLE>\n<TR><TH>Titles</TH></TR>\n<TR>");
$query = "SELECT title, publisher FROM books";
$result = mysql_query($query);
while ($book_row = mysql_fetch_array($result)) {
  echo("<TD>$book_row[0]</TD>\n");
}
echo("</TR></TABLE><BR>\n");
echo("<TABLE>\n<TR><TH>Publishers</TH></TR>\n<TR>");
```

```
      mysql_data_seek($result, 0);
      while ($book_row = mysql_fetch_array($result)) {
        echo("<TD>{$book_row[1]}</TD>\n");
      }
      echo("</TR></TABLE><BR>\n");
      ?>
```

Without using `mysql_data_seek`, the second usage of the result set would turn back no 0 rows because it has already iterated through to the end of the dataset and the pointer stays there until you explicitly move it. This handy function helps greatly when you are formatting data in a way that does not place fields in columns and records in rows.

Getting Data about Data

You only need four PHP functions to put data into or get data out of a preexisting MySQL database: `mysql_connect`, `mysql_select_db`, `mysql_query`, and `mysql_fetch_array`. Most of the rest of the functions in this section are about getting information about the data you put into or took out of the database or about the construction of the database itself. PHP offers extensive built-in functions to help you learn the name of the table in which your data resides, the data type handled by a particular column, or the number of the row into which you have just inserted data. With these functions, you can effectively work with a database about which you know very little.

The MySQL metadata functions fall into two major categories:

■ Functions that return information about the previous operation only

■ Functions that return information about the database structure in general

A very commonly used example of the first type is `mysql_insert_id()`, which returns the auto-incremented ID assigned to a row of data you just inserted. A commonly used example of the second type is `mysql_field_type()`, which reveals whether a particular database field's data must be an integer, a varchar, text, or what have you. Observe however, that this function is also deceptively named. Rather than returning the MySQL type, it returns the PHP data type. For example, an ENUM-type field will return `'string'`. Use `mysql_field_flags` to return more specialized field information. This should be apparent when you consider that it works on a result rather than on an actual MySQL field. It would be useful to have a function that got the possible values for an ENUM field, but there isn't a canned version at this point. Instead, use a "describe table" query and parse the result using PHP's `regex` functions.

Most of the data-about-data functions are pretty self-explanatory. There are a couple of things to keep in mind when using them, though. First, most of these functions are only effective if used in the proper combination — don't try to use a `mysql_affected_rows` after a `SELECT` query and then wonder what went wrong. Second, be careful about security with the functions that return information about your database structure. Knowing the name and structure of each table is very valuable to a cracker. And finally, be aware that some of these functions are shopping baskets full of simpler

functions. If you need several pieces of information about a particular result set or database, it could be faster to use `mysql_fetch_field` than all the `mysql_field` functions one after the other.

All of the MySQL metadata functions are fairly easy to use. However, their efficacy is directly related to intelligent database design rather than a mere marker of the PHP's strengths. Good database practices will make these functions useful over the long haul. The `mysqli` equivalent functions are perfect analogues in each of these cases.

Multiple Connections

Unless you have a specific reason to require multiple connections, you only need to make one database connection per PHP page. Even if you escape into HTML many times within the page, your connection is still good (assuming that it was good in the first place). You do not want to make multiple connections if you don't have to, because that is one of the most costly and time-consuming parts of most database queries.

Conversely, there's no easy way to keep your connection open from page to page — because PHP and MySQL would never know for sure when to close it after visitors wander off. Therefore, your connection is closed at the end of each script unless you use persistent connections.

The main time that you need to use different connections is when you're querying two or more completely separate databases. The most common situation in which you might do this is when you're using MySQL in a replicated situation. MySQL replication is accomplished through a master-slave setup, where you typically get reads from a slave and make writes to the master.

To use multiple connections, you simply open connections to each database as needed and make sure to hang on to the right result sets. PHP will help you do this by utilizing the result identifiers discussed in the "Making MySQL Queries" section earlier in the chapter. You pass the identifiers along with each MySQL function as an optional argument. If you're completing all your queries on one connection before moving on to the next, you don't even need to do this; PHP will automatically use the last link opened.

In this example, we are using connections from three different databases on different servers:

```php
<?php
$link1 = mysql_connect('host1', 'me', 'sesame');
mysql_select_db('userdb', $link1);
$query1 = "SELECT ID FROM usertable
          WHERE username = '$username'";
$result1 = mysql_query($query1, $link1);
$array1 = mysql_fetch_array($result1);
$usercount = mysql_num_rows($result1);
mysql_close($link1);

$today = '2002-05-01';
$link2 = mysql_connect('host2', 'myself', 'benne');
mysql_select_db('inventorydb', $link2);
```

```
    $query2 = "SELECT sku FROM widgets
               WHERE ship_date = '$today'";
    $result2 = mysql_query($query2, $link2);
    $array2 = mysql_fetch_array($result2);
    $widgetcount = mysql_num_rows($result2);
    mysql_close($link2);

    if ($usercount > 0 && $widgetcount > 0) {
      $link3 = mysql_connect('host3', 'I', 'seed');
      mysql_select_db('salesdb', $link3);
      $query3 = "INSERT INTO saleslog (ID, date, userID, sku)
               VALUES (NULL, '$today', '$array1[0]', '$array2[0]')";
      $result3 = mysql_query($query3, $link3);
      $insertID = mysql_insert_id($link3);
      mysql_close($link3);
      if ($insertID >= 1) {
        print("Perfect entry");
      }
      else {
        print("Danger, danger, Will Robinson!");
      }
    } else {
      print("Not enough information");
    }
    ?>
```

In this example, we have deliberately kept the connections as discrete as possible for clarity's sake, even going to the trouble to close each link after we use it. Without the `mysql_close()` commands, we would be running multiple concurrent connections — which you may want to do. There's nothing stopping you from doing so. Just remember to pass the link value carefully from one function to the next, and you should be fine.

Building in Error Checking

This section could have been titled "Die, die, die!" because the main error-checking function is actually called `die()`. There was something about that title that failed to reinforce the warm, hospitable learning environment we cherish, so we went with the more prosaic subheading.

`die()` is not a MySQL-specific function — the PHP manual lists it in "Miscellaneous Functions." It simply terminates the script (or a delimited portion thereof) and returns a string of your choice.

```
    mysql_query("SELECT * FROM mutual_funds
               WHERE code = '$searchstring'")
    or die("Please check your query and try again.");
```

Notice the syntax: the word `or` (you could alternatively use `||`, but that isn't as much fun as saying `or die`) and only one semicolon per pair of alternatives.

Until quite recently, MySQL via PHP returned very insecure and unenlightening (except to crackers) error messages upon encountering a problem with a database query. die() was often used as a way to exert control over what the public would see on failure. Now that no error messages are returned at all, die() may be even more necessary — unless you want your visitors to be left wondering what happened.

Other built-in means of error-checking are error messages. These are particularly helpful during the development and debugging phase, and they can be easily commented out in the final edit before going live on a production server. As mentioned, MySQL error messages no longer appear by default. If you want them, you have to ask for them by using the functions mysql_errno() (which returns a code number for each error type) or mysql_error() (which returns the text message). Then you can send them to a custom error log by using the error_log() function:

```
if (!mysql_select_db($bad_db)) {
  print(mysql_error());
}
```

There's more to database error handling than judicious use of die(), however. Servers become unavailable, data sets get corrupted, and so forth. We've been fairly liberal in setting up connections and executing queries, but ideally, every interaction with the database should be nested inside a conditional that returns the desired result on success and a nice clean error page on failure. This is where die() drops the ball. Execution immediately stops for the entire script, leaving off, if nothing else, closing tags for your HTML page if they are defined in PHP. Additionally, there may be plenty more perfectly good scripting or HTML left to go on the page — code that is unaffected by a dropped database connection or a failed query. Finally, die() doesn't let you know anything went wrong. Do you really think that your users will tell you? Probably not. It's much more realistic that they will leave your site in disgust and never return. An example of good error checking is:

```
function printError($errorMesg) {
  printf("<B>%s </B><BR>\n", $errorMesg);
}
function notify($errorMesg) {
  mail(webmaster@example.com, "An Error has occurred at
example.com", $errorMesg)
}
if ($link = mysql_connect("host", "user", "pass")) {
  // Things to do if the connection is successful
} else {
  printError("Sorry for the inconvenience; but we are unable
to process your request at this time. Please check back
later");
  notify("Problem connecting to database in $SCRIPT_NAME at
line 12 on date('Y-m-D')");
}
```

Even better, if you really want to get your feet wet with PHP6's new object-oriented programming (OOP) features, try using exceptions, which are covered in Chapter 30.

Creating MySQL Databases with PHP

You can, if you wish, actually create your databases with PHP rather than using the MySQL client tool. This practice has potential advantages — you can use an attractive front end that may appeal to those who find the MySQL command-line client horribly plain or finicky to use — counterbalanced by one big disadvantage, which is security.

To create a database from PHP, the user of your scripts will need to have full `CREATE/DROP` privileges on MySQL. That means anyone who can get hold of your scripts can potentially blow away all your databases and their contents with the greatest of ease. This is not such a great idea from a security standpoint.

If you're even considering creating databases with PHP, do yourself a big favor and at least don't store the database username and password in a text file. Make yourself type your database username and password into a form and pass the variables to the inserting handler each and every time you use this script. This is one case where keeping the variables in an `include` file outside your web tree is not sufficient precaution. Better yet, run the scripts manually from the command line through SSH:

```
mysql -u <username> -p <databasename> < sql-script.sql
```

For those times when you need to create databases programmatically, the relevant functions are:

- `mysql_create_db()`: Creates a database on the designated host, with name specified in arguments
- `mysql_drop_db()`: Deletes the specified database
- `mysql_query()`: Passes table definitions and drops in this function

A bare-bones database-generation script might look like this:

```
<?php
$linkID = mysql_connect('localhost', 'root', 'sesame');
mysql_create_db('new_db', $linkID);
mysql_select_db('new_db');
$query = "CREATE TABLE new_table (
        id INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
      new_col VARCHAR(25)
        )";
$result = mysql_query($query);
$axe = mysql_drop_db('new_db');
?>
```

Several other GUI tools are available that are not database-specific but will probably work with MySQL. As MySQL has become more and more popular, a number of applications for both Windows and Linux have come into play that allow you to administer MySQL databases in the graphical fashion you may have become accustomed to. Like their web counterparts, these applications offer full administrative control, but without the headache of exposing yourself to the security

risk of a web-based interface. The list changes often as software comes and goes, so a listing here would probably very quickly go out of date. However, the MySQL web site keeps a pretty comprehensive list at `http://dev.mysql.com`.

MySQL data types

The actual PHP functions used to create MySQL databases are trivial compared to the MySQL data structure statements that are passed in those functions. The "Database Design" section of Chapter 13 has general rules on how to conceptualize a database schema and use the `CREATE`, `DROP`, and `ALTER` statements. To implement your abstract schema in MySQL, however, you also need to understand MySQL data types and how to use them.

The general rule is to use the smallest and most specific data type that will adequately meet the needs of this particular column in your database. MySQL is known for having compact types, such as `TINYINT` and `TINYTEXT`, that are good for things like 0/1 values or first names. It also has very large types that can store 4GB (or more) of data in one field.

There are three buckets of MySQL data types: numeric types, date and time types, and string (or character) types. For the most part, their use is fairly straightforward — in the sense that the average user is not going to know or care whether you used an `INT` or a `MEDIUMINT`. However, if you're the type of programmer who cares about doing everything in the absolutely tightest and fastest way possible, the MySQL manual gives subtle tips on maximizing efficiency — for instance, always use the `DECIMAL` type with money, or it takes 8 bytes to store a `DATETIME` but only 4 bytes to store a Unix `TIMESTAMP`, which PHP can convert to any date-time format you desire. Careful perusal of the "Column Types" section of the MySQL manual (at `www.mysql.com/doc/en/Column_types.html`) may yield hidden treasures of insight.

Table 15-1 shows the current MySQL data types and their possible values. `M` stands for the maximum number of digits displayed, and `D` stands for the maximum number of decimal places in a floating-point number. Both are optional.

TABLE 15-1

MySQL Data Types

Name and Aliases	Storage size	Usage
`TINYINT(M)`		
`BIT`, `BOOL`, `BOOLEAN` are synonyms for `TINYINT(1)`	1 byte	If unsigned, stores values from 0 to 255; otherwise, from -128 to 127. A new Boolean type will appear in future, but until now has been implemented as a `TINYINT(1)`.
`SMALLINT(M)`	2 bytes	If unsigned, stores values from 0 to 65535; otherwise, from -32768 to 32767.

Name and Aliases	Storage size	Usage
MEDIUMINT(M)	3 bytes	If unsigned, stores values from 0 to 16777215; otherwise, from -8388608 to 8388607.
INT(M)		
INTEGER(M)	4 bytes	If unsigned, stores values from 0 to 4294967295; otherwise, from -2147483648 to 2147483647.
BIGINT(M)	8 bytes	If unsigned, stores values from 0 to 18446744073709551615; otherwise, from -9223372036854775808 to 9223372036854775807. You may experience strangeness when performing arithmetic with unsigned integers of this size due to limitations in your operating system.
FLOAT(precision)	4 or 8 bytes	Where precision is an integer up to 53. If precision <= 24, converted to a FLOAT; if precision > 24 and <= 53, converted to a DOUBLE. Provided for Open DataBase Connectivity (ODBC) compatibility; in general, use the normal MySQL FLOAT and DOUBLE types.
FLOAT(M,D)	4 bytes	Single-precision floating-point number.
DOUBLE(M,D)		
DOUBLE PRECISION, REAL	8 bytes	Double-precision floating-point number.
DECIMAL(M,D)		
DEC, NUMERIC, FIXED	M+1 or M+2 bytes	An unpacked floating-point number that is stored like a CHAR. Used for small decimals, such as money.
DATE	3 bytes	Displayed in the format YYYY-MM-DD.
DATETIME	8 bytes	Displayed in the format YYYY-MM-DD HH:MM:SS.
TIMESTAMP	4 bytes	Since MySQL 4.1, can no longer set display size. Displayed in the same format as DATETIME.
TIME	3 bytes	Displayed in the format HHH:MM:SS where HHH is a value from -838 to 838. This allows a TIME value to represent an elapsed time between two events.
YEAR	1 byte	Displayed in the format YYYY, which is a value from 1901 to 2155. To use an earlier date, you should use a TINYINT type.

continued

231

TABLE 15-1   *(continued)*

Name and Aliases	Storage size	Usage
CHAR(M)	M bytes	Fixed in length. If your string is not long enough, it will be padded with spaces at the end. M must be <= 255.
VARCHAR(M)	Up to M bytes	Variable in length. M must be <= 255.
BINARY(M)	Up to M bytes	Stores byte strings.
VARBINARY(M)	Up to M bytes	Similar to VARCHAR. Stores byte strings.
TINYBLOB or TINYTEXT	Up to 255 bytes	TINYBLOB is case-sensitive for sorting and comparison; TINYTEXT is case-insensitive.
BLOB or TEXT	Up to 64KB	BLOB is case-sensitive for sorting and comparison; TEXT is case-insensitive.
MEDIUMBLOB or MEDIUMTEXT	Up to 16MB	MEDIUMBLOB is case-sensitive for sorting and comparison; MEDIUMTEXT is case-insensitive.
LONGBLOB or LONGTEXT	Up to 4GB	LONGBLOB is case-sensitive for sorting and comparison; LONGTEXT is case-insensitive.
ENUM(value1, ...valueN)	1 or 2 bytes	Up to 65535 distinct values.
SET(value1,... valueN)	Up to 8 bytes	Up to 64 distinct values.

MySQL Functions

Table 15-2 includes a recap of the MySQL functions. All arguments in brackets are optional.

TABLE 15-2

PHP-MySQL Functions

Function Name	Usage
mysql_affected_rows([link_id])	Use after a nonzero INSERT, UPDATE, or DELETE query to check number of rows changed.

Function Name	Usage
`mysql_change_user(user, `*`password[,`*` database]` `[, link_id])`	Changes MySQL user on an open link.
`mysql_close([link_id])`	Closes the identified link (usually unnecessary).
`mysql_connect([host][:port][:socket][,` `username][, password])`	Opens a link on the specified host, port, socket; as specified user with password. All arguments are optional.
`mysql_create_db(db_name[, link_id])`	Creates a new MySQL database on the host associated with the nearest open link.
`mysql_data_seek(result_id, `*`row_num)`*	Moves internal row pointer to specified row number. Use a fetching function to return data from that row.
`mysql_drop_db(db_name[, `*`link_id])`*	Drops specified MySQL database.
`mysql_errno([link_id])`	Returns ID of error.
`mysql_error([link_id])`	Returns text error message.
`mysql_fetch_array(result_id[, `*`result_type])`*	Fetches result set as associative array. Result type can be `MYSQL_ASSOC`, `MYSQL_NUM`, or `MYSQL_BOTH` (default).
`mysql_fetch_field(result_id[, `*`field_offset])`*	Returns information about a field as an object.
`mysql_fetch_lengths(result_id)`	Returns length of each field in a result set.
`mysql_fetch_object(result_id[, `*`result_type])`*	Fetches result set as an object. See `mysql_fetch_array` for result types.
`mysql_fetch_row(result_id)`	Fetches result set as an enumerated array.
`mysql_field_name(result_id, `*`field_index)`*	Returns name of enumerated field.
`mysql_field_seek(result_id, `*`field_offset)`*	Moves result pointer to specified field offset. Used with `mysql_fetch_field`.
`mysql_field_table(result_id, field_offset)`	Returns name of specified field's table.
`mysql_field_type(result_id, field_offset)`	Returns type of offset field (for example, `TINYINT`, `BLOB`, `VARCHAR`).
`mysql_field_flags(result_id, field_offset)`	Returns flags associated with enumerated field (for example, `NOT NULL`, `AUTO_ INCREMENT`, `BINARY`).

continued

233

TABLE 15-2	*(continued)*
Function Name	**Usage**
mysql_field_len(result_id, field_offset)	Returns length of enumerated field.
mysql_free_result(result_id)	Frees memory used by result set (usually unnecessary).
mysql_insert_id([link_id])	Returns AUTO_INCREMENTED ID of INSERT; or FALSE if insert failed or last query was not an insert.
mysql_list_fields(database, table[, link_id])	Returns result ID for use in mysql_field functions, without performing an actual query.
mysql_list_dbs([link_id])	Returns result pointer of databases on mysqld. Used with mysql_tablename.
mysql_list_tables(database[, link_id])	Returns result pointer of tables in database. Used with mysql_tablename.
mysql_num_fields(result_id)	Returns number of fields in a result set.
mysql_num_rows(result_id)	Returns number of rows in a result set.
mysql_pconnect([host][:port][:socket][, username][, password])	Opens persistent connection to database. All arguments are optional. Be careful — mysql_close and script termination will not close the connection.
mysql_query(query_string[, link_id])	Sends query to database. Remember to put the semicolon outside the double-quoted query string.
mysql_result(result_id, row_id, field_identifier)	Returns single-field result. Field identifier can be field offset (0), field name (FirstName) or table-dot name (myfield.mytable).
mysql_select_db(database[, link_id])	Selects database for queries.
mysql_tablename(result_id, table_id)	Used with any of the mysql_list functions to return the value referenced by a result pointer.

Summary

PHP's MySQL and MySQL Improved functions are easy to use, if sometimes named confusingly. Each instance of a PHP/MySQL interaction must have a connection, a database select, and a query or command that returns a result identifier. The result identifier is like an ATM receipt that reports on the success or failure of an operation.

If data is returned after a `SELECT` statement, one of the PHP/MySQL fetching functions must also be employed. Data pulled from a MySQL database exists in a kind of limbo until one of the fetching functions is applied to the result set. If you wish to loop through the result set again, you can use `mysql_data_seek()` to reset the row pointer to zero.

PHP also has a large number of functions that return data about the database itself or about a particular operation. Two of the most common are `mysql_num_rows()`, which returns the number of rows in a result set, and `mysql_insert_id()`, which returns the ID of the proximate `INSERT` operation.

PHP handles much of the MySQL connectivity for you without requiring specific link identifiers or result pointers. The exception comes when you need multiple database connections on the same web page. In this case, you use exactly the same functions and syntax but simply pass the correct link identifier with most commands.

We do not personally recommend creating MySQL databases with PHP front ends.

Chapter 16

Performing Database Queries

Much of the point of PHP is to help you translate between a backend database and its frontend presentation on the web. Data can be viewed, added, removed, and tweaked as a result of your web user's keystrokes and mouse clicks.

For most of this chapter, we restrict ourselves to ways to use PHP to look at the contents of a database without altering it, using only the `SELECT` statement from SQL and displaying the results in HTML tables. We use a single database example to show different strategies, including some handy reusable functions. Finally, we look at code to create the sample data shown in the display examples, using the `INSERT` statement.

The two big productivity points from this chapter are:

- *Reuse functions* in simple cases. The problem of database table display shows up over and over in database-enabled site design. If the display is not complicated, you should be able to throw the same simple function at the problem rather than reinventing the wheel with each PHP page you write.

- *Choose between techniques* in complex cases. You may find yourself wanting to pull out a complex combination of information from different tables (which, of course, is part of the point of using a relational database to begin with). You may not be able to map this onto a simple reusable function, but there aren't that many novel solutions either — get to know the alternatives, and you can decide how to trade off efficiency, readability, and your own effort.

> **NOTE** This chapter uses the MySQL database and functions exclusively, but the display strategies should be directly transferable to almost any SQL-compliant database supported by PHP.

237

HTML Tables and Database Tables

First of all, some terminology — unfortunately, both relational databases and HTML scripting use the term *table*, but the term means very different things in the two cases. A *database table* persistently stores information in columns, which have predefined names and types so that the information in them can be recovered later. An *HTML table* is a construct that tells the browser to lay out arbitrary HTML contents in a rectangular array in the browser window. We'll try to always make it clear which kind of table we are talking about.

One-to-one mapping

HTML tables are really constructed out of rows (the `<TR></TR>` construct), and columns have no independent existence — each row has some number of table datum items (the `<TD></TD>` construct), which will produce a nice rectangular array only if there are the same number of `TD`s for every `TR`. (There is no corresponding `<TC>` construct that lets you display by column first.) By contrast, *fields* (aka *columns*) in database tables are the more primary entity — defining a table means defining the fields, and then you can add as many rows as you like. In this chapter, we will focus on printing out tables and queries in such a way that each database field prints in its own HTML column, simply because there are usually more database rows than database fields, and people are more used to up-and-down scrolling than left-to-right scrolling. If you find yourself wanting to map database fields to HTML rows, it is a simple inversion exercise.

The simplest case of displaying a table is the one in which the structure of a database table or query *does* correspond to the structure of the HTML table we want to display — the database entity has *m* columns and *n* rows, and we'd like to display an *m*-by-*n* rectangular grid in the user's browser window, with all the cells filled in appropriately.

Example: A single-table displayer

So let's write a simple translator that queries the database for the contents of a single table and displays the results onscreen. Here's the top-down outline of how the code will get the job done:

1. Establish a database connection.
2. Construct a query to send to the database.
3. Send the query and hold on to the result identifier that is returned.
4. Using the result identifier, find out how many columns (fields) there are in each row.
5. Start an HTML table.
6. Loop through the database result rows, printing a `<TR></TR>` pair to make a corresponding HTML table row.
7. In each row, retrieve the successive fields and display them wrapped in a `<TD></TD>` pair.
8. Close off the HTML table.
9. Close the database connection.

Finally, we'd like to wrap all the preceding steps up into a handy function that we can use whenever we want to. Also, for reasons of efficiency, we don't want to include the first and last steps of creating and closing the database connection in the function — we may want to use such a function more than once per page, and it wouldn't make sense to open and close the connection each time. Instead, we'll assume that we have a connection already and pass the connection to the function along with the table name.

Such a function is shown in Listing 16-1, embedded in a complete PHP page that uses the function to display the contents of a couple of tables.

A table displayer

```php
<?php
include("/home/phpbook/phpbook-vars.inc");
$global_dbh = mysql_connect($hostname, $username, $password);
mysql_select_db($db, $global_dbh);

function display_db_table($tablename, $connection)
{
  $query_string = "SELECT * FROM $tablename";
  $result_id = mysql_query($query_string, $connection);
  $column_count = mysql_num_fields($result_id);

  print("<TABLE BORDER=1>\n");
  while ($row = mysql_fetch_row($result_id))
    {
      print("<TR ALIGN=LEFT VALIGN=TOP>");
      for ($column_num = 0;
           $column_num < $column_count;
           $column_num++)
        print("<TD>$row[$column_num]</TD>\n");
      print("</TR>\n");
    }
  print("</TABLE>\n");
}
?>

<HTML>
<HEAD>
<TITLE>Cities and countries</TITLE>
</HEAD>
<BODY>

<TABLE><TR><TD>
<?php display_db_table("country", $global_dbh); ?>
</TD><TD>
```

```
<?php display_db_table("city", $global_dbh); ?>
</TD></TR></TABLE></BODY></HTML>
```

Some things to notice about this script:

- Although the script refers to specific database tables, the `display_db_table()` function itself is general. You could put the function definition in an `include` file and then use it anywhere on your site.

- The first thing the script does is load in an `include` file that contains variable assignments for the database name, database username, and database password. It then uses those variables to connect to MySQL and then to choose the desired database. (The fact that this file is located outside the publicly available web hierarchy makes it *slightly* more secure than just including that information in your code.)

- In the function itself, we chose to use a `while` loop for printing rows and a `for` loop to print the individual items. We could as easily have used a bounded `for` loop for both and recovered the number of rows with `mysql_num_rows()`.

- The main `while` loop reflects a very common idiom, which exploits the fact that the value of a PHP assignment statement is the value assigned. The variable `$row` is assigned to the result of the function `mysql_fetch_row()`, which will be either an array of values from that row or a false value if there are no more rows. If we're out of rows, `$row` is false, which means that the value of the whole expression is false, which means that the `while` loop terminates.

- We put line breaks (\n) at the end of selected lines, so that the HTML source would have a readable structure when printed or viewed as source from the browser. Notice that these breaks are not HTML line breaks (<BR>) and do not affect the look of the resulting web page. (In fact, if you want to make it annoying for someone else to scrutinize the HTML you generate, don't put breaks in at all!)

The sample tables

To see the Listing 16-1 script in action, see Figure 16-1, which shows the displayed contents of the Country and City sample tables. These tables have the following structure:

```
Country:
    ID int (auto-incremented primary key)
    continent varchar(50)
    countryname varchar(50)
City:
    ID int (auto-incremented primary key)
    countryID int
    cityname varchar(50)
```

FIGURE 16-1

A simple database table display

Think of these tables as a rough draft of the database for an eventual online almanac. They employ our usual convention of always having one field per table called `ID`, which is a primary key and has successive integers assigned to it automatically for each new row. Although you can't tell for sure from the preceding description, the tables have one "relation" embodied in their structure — the `countryID` field of the `City` table is matched up with the `ID` field of the `Country` table, representing which country the city belongs to. (If you were designing a real almanac database, you would want to take this one step further and break the `Country` table into a relational pair of `Country` and `Continent` tables.)

CROSS-REF To see how we created these tables and populated them with sample data, see the "Creating the Sample Tables" section at the end of this chapter.

Improving the displayer

Our first version of this function has some limitations: It works with a single table only, does no error-checking and is very bare-bones in its presentation. We'll address these problems one by one and then fix them in one fell revision. (If you want to look ahead, the new-and-improved version of the function is in Listing 16-2.)

Displaying column headers

Our first version of a database table displayer simply displays all the table cells, without any labeling of what the different fields are. It's conventional in HTML to use the `<TH>` element for column and/or row headers — in most browsers and styles, this displays as a bold table cell. One improvement we can make is to optionally display column headers that are based on the names of the table fields themselves. To actually retrieve those names, we can use the function `mysql_field_name()`.

Error checking

Our original version of the code assumes that we have written it correctly and also that our database server is up and functioning normally — if either of these is not the case, we will run into puzzling errors. We can partially address this by appending a call to `die()` to the actual database queries — if they fail, an informative message will be printed. This is a reasonable approach for such a small example, but as projects get larger it is better to use the exception-handling capability introduced back in PHP5.

Cosmetic issues

Another source of dissatisfaction with our simple table-displayer is that it always has the same look. It would be nice, at a minimum, to control whether table borders are displayed. The simple solution we will use in our new function is just to permit passing in a string of arguments that will be spliced into the HTML table definition. This is a pretty crude form of style control that style sheet proponents would discourage, but it will permit us to directly specify some elements of the table's look without writing an entirely new function.

Displaying arbitrary queries

Finally, it would be nice to be able to exploit our relational database and display the results of complex queries rather than just single tables. Actually, our single-table displayer has an arbitrary query embedded in it — it just happens that it is hardcoded as `select * from table`, where `table` is the supplied table name. So let us transform our simple table displayer into a query displayer and then recreate the table displayer as a simple wrapper around the query displayer. These two functions, complete with the cosmetic improvements and better error checking, are shown in Listing 16-2.

LISTING 16-2

A query displayer

```php
<?php
include("/home/phpbook/phpbook-vars.inc");
$global_dbh = mysql_connect($hostname, $username, $password)
              or die("Could not connect to database");
```

```
mysql_select_db($db, $global_dbh)
  or die("Could not select database");

function display_db_query($query_string, $connection,
                          $header_bool, $table_params)
{

  // perform the database query
  $result_id = mysql_query($query_string, $connection)
               or die("display_db_query:" . mysql_error());

  // find out the number of columns in result
  $column_count = mysql_num_fields($result_id)
                  or die("display_db_query:" . mysql_error());

  // TABLE form includes optional HTML arguments passed
  //   into function
  print("<TABLE $table_params >\n");

  // optionally print a bold header at top of table
  if ($header_bool)
  {
    print("<TR>");
    for ($column_num = 0;
         $column_num < $column_count;
         $column_num++)
    {
      $field_name =
        mysql_field_name($result_id, $column_num);
      print("<TH>$field_name</TH>");
    }
    print("</TR>\n");
  }
  // print the body of the table
  while ($row = mysql_fetch_row($result_id))
  {
    print("<TR ALIGN=LEFT VALIGN=TOP>");
    for ($column_num = 0;
         $column_num < $column_count;
         $column_num++)
      {
        print("<TD>$row[$column_num]</TD>\n");
      }
    print("</TR>\n");
  }
  print("</TABLE>\n");
}

function display_db_table($tablename, $connection,
                          $header_bool, $table_params)
```

243

```
{
  $query_string = "SELECT * FROM $tablename";
  display_db_query($query_string, $connection,
                   $header_bool, $table_params);
}
?>

<HTML><HEAD><TITLE>Countries and cities</TITLE></HEAD>
<BODY>
<TABLE><TR><TD>
<?php display_db_table("country", $global_dbh,
                       TRUE, "BORDER=2"); ?>
</TD><TD>
<?php display_db_table("city", $global_dbh,
                       TRUE, "BORDER=2"); ?>
</TD></TR></TABLE></BODY></HTML>
```

The result of using this code on the same database contents is shown in Figure 16-2. The only visible difference is the column header. Splitting the functions apart means that we also have a new function in our bag of tricks — we could do the same kind of display with an arbitrary query string that joins data from different tables.

FIGURE 16-2

Using the query displayer

244

Complex Mappings

So far in this chapter, we've enjoyed a very nice and simple-minded correspondence between query resultsets and HTML tables — every row in the resultset corresponds to a row in the table, and the structure of the code is simply two nested loops. Unfortunately, life isn't often this simple, and sometimes the structure of the HTML table we want to display has a complex relationship to the relational structure of the database tables.

Views and Stored Procedures

Our query displayer assumes a particular division of labor between the PHP code and the database system itself — the PHP code sends off an arbitrary query string, which the database responds to by setting up a resultset. In particular, this means that the database system has to parse that query and then figure out the best way to go about retrieving the results. This is part of what can make querying a database a mildly expensive operation.

In cases where your code may construct novel queries on the fly, this is the best you can hope for. However, some databases offer ways to set up queries in advance, which gives the database system a chance to preoptimize how it handles the query. One such construct is called a *view* under MS SQL Server and some other RDBMSs — after you have set up a query as a named view, it can be treated just like a real table. A related idea is the *stored procedure*, which is like a view that also accepts runtime arguments that are spliced into the query. In general, if you realize that you are suffering from slow query performance, you may want to investigate what similar optimizations your particular RDBMS makes available.

Multiple queries versus complex printing

Let's say that, rather than displaying our sample `City` and `Country` tables individually, we want to match them up in a tabular display.

We can easily write a `SELECT` statement that joins these tables appropriately:

```
SELECT country.continent, country.countryname,
       city.cityname
FROM country, city
WHERE city.countryID = country.ID
ORDER BY continent, countryname, cityname
```

Now, this would be a handy place to use our query-displayer function — all we have to do is send it the preceding statement as a string, and it will print out a table of cities matched up with their continents and countries. However, if we do this, we will see an individual HTML table row for each city, and the continent and country will print each time — for example, we'll see `North America` printed several times. Instead, what if we want one name matched with many titles? This is a case where the structure of what we print differs from the structure of the most convenient query.

If we want to do a more complex mapping, we have a choice: We can throw database queries at the problem, or we can write more complex display code. Let's look at each option in turn. (For each of these examples, we'll be moving away from the reusable generality of the functions we wrote earlier toward functions that address a particular display problem.)

A multiple-query example

If we want to print just one HTML row per country, we can make a query for the countries and then make another query for the relevant cities in each trip through a `country` row. A function written using this strategy is shown in Listing 16-3.

LISTING 16-3

A display with multiple queries

```php
<?php
include("/home/phpbook/phpbook-vars.inc");
/* open database connection */
$global_dbh = mysql_connect($hostname, $username, $password)
              or die("Could not connect to database");
mysql_select_db($db, $global_dbh)
  or die("Could not select database");

function display_cities($db_connection)
{
  /* Displays table of cities and countries */
  $country_query = "SELECT id, continent, countryname
                    FROM country
                    ORDER BY continent, countryname";
  $country_result =
     mysql_query($country_query, $db_connection);

  /* begin table, print hard-coded table header */
  print("<TABLE BORDER=1>\n");
  print("<TR><TH>Continent</TH><TH>Country</TH>
           <TH>Cities</TH></TR>");

  /* loop through countries */
  while ($country_row = mysql_fetch_row($country_result))
    {
      /* set up country info */
      $country_id = $country_row[0];
      $continent = $country_row[1];
      $country_name = $country_row[2];

      print("<TR ALIGN=LEFT VALIGN=TOP>");
      print("<TD>$continent</TD>");
```

```
      print("<TD>$country_name</TD>");

      /* begin table cell for city list */
      print("<TD>");
      $city_query = "select cityname from city
                     where countryID = $country_id
                     order by cityname";
      $city_result =
         mysql_query($city_query, $db_connection)
           OR die(mysql_error());
      /* loop through cities */
      while ($city_row = mysql_fetch_row($city_result))
        {
          $city_name = $city_row[0];
          print("$city_name<BR>");
        }
      /* close city cell and country row */
      print("</TD></TR>");
    }
  print("</TABLE>\n");
}
?>

<HTML>
<HEAD>
<TITLE>Cities by Country</TITLE>
</HEAD>
<BODY>
<?php
  display_cities($global_dbh);
?>
</BODY>
</HTML>
```

The strategy is appealingly simple: There is an outer loop that uses one query to proceed through all the countries, saving the country's name and the primary ID field of each country row. Then for each country, the ID field is used to look up all the cities belonging to that country. Notice the trick of embedding the $countryid variable in the inner query — the query string sent is actually different on each iteration through the country loop.

Simple? Yes. Efficient? Probably not. This code makes a *separate* city query for each country. If there are 500 countries in the database, this function will make 501 separate database queries (the extra one being the enclosing country query).

Your mileage will vary according to how efficient your particular database is in parsing queries and planning query retrieval, but the sum of these queries will certainly take more time than the simple query we started this section with.

A complex printing example

Now let's solve exactly the same problem, but using a different strategy. Instead of making multiple queries, we will make a single query and print the resulting rows selectively, so that each HTML table row corresponds to more than one database row (see Listing 16-4). The resulting browser display is exactly the same as in the previous example.

LISTING 16-4

A complex display with a single query

```php
<?php
include("/home/phpbook/phpbook-vars.inc");
/* open a single DB connection for this page */
$global_dbh = mysql_connect($hostname, $username, $password)
              or die("Could not connect to database");
mysql_select_db($db, $global_dbh)
              or die("Could not select database");

function display_cities($db_connection)
{
  /*  print table of countries and their cities,
      selectively printing only one HTML table row
      per country */
  $query = "SELECT country.id,
                   country.continent, country.countryname,
                   city.cityname
                   FROM country, city
                   WHERE country.id = city.countryID
                   ORDER BY country.continent,
                            country.countryname,
                            city.cityname";
  $result_id =
     mysql_query($query, $db_connection)
       OR die(mysql_error($query));

  /* begin table, print hard-coded table header */
  print("<TABLE BORDER=1\n");
  print("<TH>Continent</TH><TH>Country</TH>
           <TH>Cities</TH></TR>");

  /* Initialize the ID for the "previous" country.
     We will rely on the fact that Country.ID is
     numbered beginning with 1, so a previous ID
     value of zero means that the current country
     is the first */
  $old_country_id = 0;
```

```
   /* loop through result rows (one per city) */
   while ($row_array = mysql_fetch_row($result_id))
     {
       $country_id = $row_array[0];
       /* if we have a new country */
       if ($country_id != $old_country_id)
         {
           /* set up country info */
           $continent = $row_array[1];
           $country_name = $row_array[2];

           /* if there was a previous country
              close the city datum and country row */
           if ($old_country_id != 0)
             print("</TD></TR>\n");

           /* start a row for the new country,
              and begin the city table datum */
           print("<TR ALIGN=LEFT VALIGN=TOP>");
           print("<TD>$continent</TD>");
           print("<TD>$country_name</TD><TD>");

           /* the new country is no longer new */
           $old_country_id = $country_id;
         }
       /* the only thing that is printed for every result
          row is the name of a city */
       $city_name = $row_array[3];
       print("$city_name<BR>");
   }
  /* close off final country and table */
  print("</TD></TR></TABLE>");
}
?>
<HTML><HEAD><TITLE>Cities by Country</TITLE></HEAD>
<BODY>
<?php display_cities($global_dbh);
 ?>
</BODY></HTML>
```

This code is somewhat tricky — although it goes through the result rows in order, and everything it prints is grabbed from the current row, it prints countries only when their values have changed. (Continents are still printed redundantly.)

The change in a country is detected by monitoring the ID field of the country row. A country change is also a signal to print out the HTML necessary to close off the preceding table row and start a new one. Finally, the code must handle printing the HTML necessary to start the first row and end the last one.

249

Creating the Sample Tables

Now we will show you the PHP/MySQL code we actually used to create the sample tables. (Such data might more normally be created by interacting only with MySQL, but we decided to respect our book's title by doing it from PHP.) The code (shown in Listing 16-5) is a special-purpose, one-time hack, not a model of style, but it has useful examples of using the SQL INSERT statement.

LISTING 16-5

Creating the sample tables

```php
<?php
include("/home/phpbook/phpbook-vars.inc");
$global_dbh = mysql_connect($hostname, $username, $password)
             or die("Could not connect to database");
mysql_select_db($db, $global_dbh)
             or die ("Could not select databased");

function add_new_country($dbh, $continent, $countryname,
                         $city_array)

{
  $country_query =
    "INSERT INTO country (continent, countryname)
     VALUES ('$continent', '$countryname')";
  $result_id =  mysql_query($country_query)
               OR die($country_query . mysql_error());
  if ($result_id)
    {
      $countryID = mysql_insert_id($dbh);
      for ($city = current($city_array);
           $city;
           $city = next($city_array))
        {
          $city_query =
            "INSERT INTO city (countryID, cityname)
                  VALUES ($countryID, '$city')";
          mysql_query($city_query, $dbh)
            OR die($city_query . mysql_error());
        }
    }
}

function populate_cities_db($dbh)
{
  /* drop tables if they exist—permits function to be
     tried more than once */
```

```
  mysql_query("DROP TABLE city", $dbh);
  mysql_query("DROP TABLE country", $dbh);

  /* create the tables */
  mysql_query("CREATE TABLE country
                (ID int not null auto_increment primary key,
                 continent varchar(50),
                 countryname varchar(50))",
               $dbh)
               OR die(mysql_error());
  mysql_query("create table city
                (ID int not null auto_increment primary key,
                 countryID int not null,
                 cityname varchar(50))",
               $dbh)
               OR die(mysql_error());

  /* store data in the tables */
  add_new_country($dbh, 'Africa', 'Kenya',
          array('Nairobi','Mombasa','Meru'));
  add_new_country($dbh, 'South America', 'Brazil',
          array('Rio de Janeiro', 'Sao Paulo',
                  'Salvador', 'Belo Horizonte'));
  add_new_country($dbh, 'North America', 'USA',
          array('Chicago', 'New York', 'Houston', 'Miami'));
  add_new_country($dbh, 'North America', 'Canada',
          array('Montreal','Windsor','Winnipeg'));

  print("Sample database created<BR>");
}
?>
<HTML><HEAD><TITLE>Creating a sample database</TITLE></HEAD>
<BODY>
<?php populate_cities_db($global_dbh); ?>
</BODY></HTML>
```

You should be able to use this code to recreate the sample database on your development machine, assuming that you have PHP and MySQL configured, and an appropriately located file called php-book-vars.inc containing username, password, and database-name strings.

Just as in the display examples, this code sends off query strings (with embedded variables), but this time the queries are INSERT statements, which create new table rows. For the most part, the data inserted is just string data passed in to the function, although we chose to pass in an arbitrary number of cities per country by using an array.

The only tricky thing in creating these sample tables is setting up the relational structure. We want each city row to have an appropriate countryID, which should be equal to the actual ID of the appropriate row from the country table. However, these countryIDs are automatically assigned

251

in sequence by MySQL and are not under our control. How can we know the right `countryID` to assign? The answer is in the incredibly handy function `mysql_insert_id()`, which recovers the `ID` associated with the last `INSERT` query made via the given database connection. We insert the new country, recover the `ID` of the newly created row, and then use that `ID` in our city insertion queries.

Summary

Database interaction is one of the areas where PHP really shines. One very common use for database-enabled web code is simply to display database contents attractively. One approach to this kind of display is to map the contents of database tables, or `SELECT` statements, to corresponding HTML table elements.

When the mapping is simple enough, you can employ reusable functions that take arbitrary table names, or `SELECT` statements, and display them as a grid. When you need a more complicated combination of information from relational tables, you probably need a special-purpose function, but certain tricks recur there as well. One such trick is to craft a SQL statement that returns all the information you need, in the order you want, and selectively print only the nonredundant portions.

Near the end of this chapter, you saw a quick example of populating a set of database tables using `INSERT` statements. Aside from that, all the techniques in this chapter were read-only and do not modify the contents of databases at all. In Chapter 17, you'll see how you can get a more intimate connection to your database by combining SQL queries with HTML forms.

Chapter 17

Integrating Web Forms and Databases

F orm handling is one of PHP's very best features. The combination of HTML to construct a data-input form, PHP to handle the data, and a database server to store the data lies at the heart of all kinds of supremely useful web tasks.

HTML Forms

You already know most of what you need to make good forms to be handled by PHP and a database. There are a few PHP-specific points to brush up on:

- You must use extra caution when using any data that comes from a visitor's web browser. It may seem like common sense, but there are still too many PHP programs that don't escape incoming data from a web form or from a web browser (or anywhere). Never use unfiltered data in a database query.

- Always, always, *always* use a NAME for every data entry element (INPUT, SELECT, TEXTAREA, and so on). These NAME attributes will become PHP variable names — you will not be able to access your values if you do not use a NAME attribute for each one. If your WYSIWYG editor doesn't allow you to do this, you'll need to remember to add these NAME attributes by hand.

- A form field NAME does not need to be the same as the corresponding database field name.

- The VALUE can be set to data you wish to display in the form.

- Remember that you can pass hidden variables from form to form (or page), using the HIDDEN data entry elements. This practice has

negative security implications, so don't use it to store sensitive data and always validate the data you receive in a HIDDEN element; never trust it to be what you expect.

CROSS-REF **See Chapter 6 for more information on how to format an HTML form for use with PHP.**

Basic Form Submission to a Database

Submitting data to a database via an HTML form is straightforward if the form and form handler are two separate pages. Listing 17-1, `newsletter_signup.html`, is a simple form with only one input field.

LISTING 17-1

A simple form (newsletter_signup.html)

```
<HTML>
<HEAD>
<STYLE TYPE="text/css">
<!--
BODY, P      {color: black; font-family: verdana;
font-size: 10 pt}
H1        {color: black; font-family: arial; font-size: 12 pt}
-->
</STYLE>
</HEAD>

<BODY>
<TABLE BORDER=0 CELLPADDING=10 WIDTH=100%>
<TR>
<TD BGCOLOR="#F0F8FF" ALIGN=CENTER VALIGN=TOP WIDTH=17%>
</TD>
<TD BGCOLOR="#FFFFFF" ALIGN=LEFT VALIGN=TOP WIDTH=83%>
<H1>Newsletter sign-up form</H1>
<P>Enter your email address and we will send you our
weekly newsletter.</P>
<FORM METHOD="post" ACTION="formhandler.php">
<INPUT TYPE="text" SIZE=25 NAME="email">
<BR><BR>
<INPUT TYPE="submit" NAME="submit" VALUE="Submit">
</FORM>
</TD>
</TR>
</TABLE>

</BODY>
</HTML>
```

Figure 17-1 shows the result of the preceding code sample, a basic form to insert data into a database.

FIGURE 17-1

A form to insert data into a database

You enter the data in the database and acknowledge receipt in the form handler in Listing 17-2, which (with great originality) we are calling `formhandler.php`.

LISTING 17-2

Form handler for newsletter_signup.html (formhandler.php)

```
<HTML>
<HEAD>
<STYLE TYPE="text/css">
<!--
BODY, P      {color: black; font-family: verdana;
font-size: 10 pt}
H1        {color: black; font-family: arial; font-size: 12 pt}
-->
</STYLE>
</HEAD>

<BODY>
<TABLE BORDER=0 CELLPADDING=10 WIDTH=100%>
<TR>
<TD BGCOLOR="#F0F8FF" ALIGN=CENTER VALIGN=TOP WIDTH=17%>
</TD>
<TD BGCOLOR="#FFFFFF" ALIGN=LEFT VALIGN=TOP WIDTH=83%>
<H1>Newsletter sign-up form</H1>
```

255

```php
<?php

if (!$_POST['email'] || $_POST['email'] == "" ||
    strlen(isset($_POST['email'])) && $_POST['email'] > 30) {
  echo '<P>Is your e-mail address really that long?</P>';
} else {
  // Open connection to the database
  mysql_connect("localhost", "phpuser", "sesame")
or die("Failure to communicate with database");
  mysql_select_db("test");

  // Insert email address
  $as_email = mysql_real_escape_string($_POST['email']);
  $tr_email = trim($as_email);
  $query = "INSERT INTO mailinglist (ID, Email, Source)
            VALUES(NULL, '$tr_email',
            'www.example.com/newsletter_signup.html')
            ";
  $result = mysql_query($query);
  if (mysql_affected_rows() == 1) {
    echo '<P>Your information has been recorded.</P>';
  } else {
    error_log(mysql_error());
    echo '<P>Something went wrong with your signup
attempt.</P>';
  }
}
?>
</TD>
</TR>
</TABLE>
</BODY>
</HTML>
```

Having a separate form and form handler is a very clean design that can potentially be easier to maintain. However, there are quite a few things that you might want to do that you can't do easily with this model, caused by the difficulty of going back to the form from the form handler and the fact that variables are not available to both at the same time.

For one thing, if something goes wrong with the submission, it's very difficult to redisplay the form with the values you just filled in. This is particularly important with something like a user registration form, where you might want to check for unique e-mail addresses or matching passwords and reject the entire registration with an error message if it doesn't pass the tests. People are going to be very annoyed if one little typo causes them to lose all the data that they just filled in — and after one or two go-rounds, they will simply stop trying to register.

The first step to solving all these problems is to combine form and handler into one self-submitting PHP script.

Self-Submission

Self-submission refers to the process of combining one or more forms and form handlers in a single script, using the HTML FORM standard to submit data to the script one or more times.

Another situation in which self-submission is a win occurs when you need to submit the same form more than once. Say that you are applying for auto insurance online, and you need to give the particulars of three or four different cars. It's extra work for the user to submit the form, get a success message, and then have to click a button to go back to the form for car #2. This kind of navigation problem has no perfect solution, but in situations where there's a high probability of multiple submissions, self-submission causes fewer clickthroughs for your web users.

Finally, the separate form and form handler make it difficult to pull data from the database, edit it, and submit it — repeating the process however many times it takes for the user to be satisfied. A common example of this usage is a form to allow users to change their personal information, such as photos and bios, which people often like to fiddle with until they look exactly the way that the users want. If you want to make five small incremental edits to your user profile, you aren't going to want to go back and forth between the form and form handler 10 times.

Self-submission is accomplished by the simplest of means: specifying the same script name as the ACTION target in the FORM element, like this:

```
<FORM METHOD="POST" ACTION="myself.php">
```

The single most important thing to remember about self-submitting forms is: *The logic comes before the display.* If you're used to writing separate forms and handlers, this may seem a little counterintuitive at first — but think of it this way: Because your form will look different or display variables based on interactions with the database, obviously these interactions must happen before the HTML for the page is output to the browser. After you construct a few self-submitting forms, logic-before-display will seem totally natural and painless.

CAUTION **To use self-submission with controls, you will need to employ a more programmatic PHP-writing style — what we term the *maximum* or *medium* style. Beginners may find this somewhat more difficult than a clear division between the functions of HTML (form display) and PHP (form handling). This can be mitigated somewhat by using the *heredoc* syntax, as we do in many of our examples.**

If you're a think-ahead type, by now you're wondering: "*But* if the logic comes before the display, won't my script try to do the database operations before showing me the HTML form in the first place?" Good question — and an indication that we need some way to tell the script either "We want to see the form now" or "We want to insert data into the database now." This "What am I supposed to be doing now?" bit is called a *stage variable*. It lets you keep track of how many times the form has submitted values to itself and, therefore, which stage of a multistep process you have reached.

The cheapest stage variable to test for is the Submit button. You can name your Submit button and give it a value, which will be set as a PHP value only after the form is submitted at least once. The easiest way to demonstrate what we're talking about is by rewriting the previous form and form handler as one self-submitting form, as we do in Listing 17-3.

Unified form and form handler (newsletter_signup.php)

```php
<?php

if (isset($_POST['submit']) && $_POST['submit'] == 'Submit') {
  if (!isset($_POST['email']) || $_POST['email'] == "" ||
strlen($_POST['email'] > 30)) {
    $message = '<P>There is a problem.  Did you enter an email
address?</P>';
  } else {
    // Open connection to the database
    mysql_connect("localhost", "phpuser", "sesame")
or die("Failure to communicate with database");
    mysql_select_db("test");

    // Insert email address
    $as_email = mysql_real_escape_string($_POST['email']);
    $tr_email = trim($as_email);
    $query = "INSERT INTO mailinglist (ID, Email, Source)
             VALUES(NULL, '$tr_email',
             'www.example.com/newsletter_signup.html')
             ";
    $result = mysql_query($query);
    if (mysql_affected_rows() == 1) {
      $message = '<P>Your information has been recorded.</P>';

      $noform_var = 1;
    } else {
      error_log(mysql_error());
      $message = '<P>Something went wrong with your signup
attempt.</P>';
    }
  }

  // Show the form in every case except successful submission
  if (!isset($noform_var)) {
    $thisfile = "newsletter_signup.php";
    $message .= <<< EOMSG
<P>Enter your email address and we will send you our weekly
newsletter.</P>
<FORM METHOD="post" ACTION="$thisfile">
<INPUT TYPE="text" SIZE=25 NAME="email">
<BR><BR>
<INPUT TYPE="submit" NAME="submit" VALUE="Submit">
</FORM>
```

```
EOMSG;
   }
}
?>

<HTML>
<HEAD>
<STYLE TYPE="text/css">
<!--
BODY, P     {color: black; font-family: verdana;
font-size: 10 pt}
H1       {color: black; font-family: arial; font-size: 12 pt}
-->
</STYLE>
</HEAD>

<BODY>
<TABLE BORDER=0 CELLPADDING=10 WIDTH=100%>
<TR>
<TD BGCOLOR="#F0F8FF" ALIGN=CENTER VALIGN=TOP WIDTH=17%>
</TD>
<TD BGCOLOR="#FFFFFF" ALIGN=LEFT VALIGN=TOP WIDTH=83%>
<H1>Newsletter sign-up form</H1>
<?php echo $message; ?>
</TD>
</TR>
</TABLE>

</BODY>
</HTML>
```

The first time you load up this page, you should see a normal HTML form exactly like the one in Figure 17-1. If you submit it without any data or with a string that's too long (often a sign of a cracking attempt), you'll see an error message and the form again. If something goes wrong with the database INSERT, you'll see an error message and the form again. Only if the INSERT completes successfully will you not see the form again — which is the navigation we want because we don't want people to sign up for the newsletter more than once.

In the preceding example, we need to check only for two states of the form (unsubmitted or submitted), so we can use the Submit button as our stage variable. But what if you want to check for more than one state? You need a variable that is capable of taking more than one value. You could either give your Submit button different values, which would show up as different labels in the button itself, or you could set a hidden variable that is capable of taking more than one value, depending on the state. We demonstrate the technique in Listing 17-4, which collects some information and then allows you to rate your boss anonymously.

259

LISTING 17-4

A three-part form (rate_boss.php)

```php
<?php

// First set the form strings, which will be displayed
//in various cases below
$thisfile = "rate_boss.php"; //Have to set this for heredoc

$reg_form = <<< EOREGFORM
<P>We must ask for your name and email address to ensure that no
one votes more than once, but we do not associate your personal
information with your rating.</P>
<FORM METHOD="post" ACTION="$thisfile">
Name: <INPUT TYPE="text" SIZE=25 NAME="name"><BR><BR>
Email: <INPUT TYPE="text" SIZE=25 NAME="email">
<INPUT TYPE="hidden" NAME="stage" VALUE="register">
<BR><BR>
<INPUT TYPE="submit" NAME="submit" VALUE="Submit">
</FORM>
EOREGFORM;

$rate_form = <<< EORATEFORM
<P>My boss is:</P>
<FORM METHOD="post" ACTION="$thisfile">
<INPUT TYPE="radio" NAME="rating" VALUE=1>
Driving me to look for a new job.<BR>
<INPUT TYPE="radio" NAME="rating" VALUE=2>
Not the worst, but pretty bad.<BR>
<INPUT TYPE="radio" NAME="rating" VALUE=3>
Just so-so.<BR>
<INPUT TYPE="radio" NAME="rating" VALUE=4>
Pretty good.<BR>
<INPUT TYPE="radio" NAME="rating" VALUE=5>
A pleasure to work with.<BR><BR>
Boss's name: <INPUT TYPE="text" SIZE=25 NAME="boss"><BR>
<INPUT TYPE="hidden" NAME="stage" VALUE="rate">
<BR><BR>
<INPUT TYPE="submit" NAME="submit" VALUE="Submit">
</FORM>
EORATEFORM;

if (!isset($_POST['submit'])) {
  // First time, just show the registration form
  $message = $reg_form;
```

```
} elseif (isset($_POST['submit']) && $_POST['submit'] == 'Submit' && $_
POST['stage'] ==
'register') {
  // Second time, show the registration form again on error,
  // rating form on successful INSERT

  if (!isset($_POST['name']) || $_POST['name'] == "" ||
strlen($_POST['name'] > 30) || !$_POST['email'] ||
$_POST['email'] == "" || strlen($_POST['email'] > 30)) {
    $message = '<P>There is a problem.  Did you enter a name and
email address?</P>';
    $message .= $reg_form;
  } else {
    // Open connection to the database
    mysql_connect("localhost", "phpuser", "sesame")
or die("Failure to communicate with database");
    mysql_select_db("test");

    // Check to see this name and email have not appeared before
    $as_name = mysql_real_escape_string($_POST['name']);
    $tr_name = trim($as_name);
    $as_email = mysql_real_escape_string($_POST['email']);
    $tr_email = trim($as_email);
    $query = "SELECT sub_id FROM raters
              WHERE Name = '$tr_name'
              AND Email = '$tr_email'
             ";
    $result = mysql_query($query);
    if (mysql_num_rows($result) > 0) {
      error_log(mysql_error());
      $message = 'Someone with this name and password has
already rated .  If you think a mistake was made, please email
help@example.com.';
    } else {
      // Insert name and email address
      $query = "INSERT INTO raters (ID, Name, Email)
                VALUES(NULL, '$tr_name', '$tr_email')
               ";
      $result = mysql_query($query);
      if (mysql_affected_rows() == 1) {
        $message = $rate_form;
      } else {
        error_log(mysql_error());
        $message = '<P>Something went wrong with your signup
attempt.</P>';
        $message .= $reg_form;
      }
    }
```

```
  }

} elseif (isset($_POST['submit']) && $_POST['submit'] == 'Submit' && $_
POST['stage'] ==
'rate') {
  // Third time, store the rating and boss's name

  // Open connection to the database
  mysql_connect("localhost", "phpuser", "sesame")
or die("Failure to communicate with database");
  mysql_select_db("test");

  // Insert rating and boss's name
  $as_boss = mysql_real_escape_string($_POST['boss']);
  $tr_boss = trim($as_boss);
  $rating = mysql_real_escape_string($_POST['rating']);
  $query = "INSERT INTO ratings (ID, Rating, Boss)
            VALUES(NULL, '$rating', '$tr_boss')
           ";
  $result = mysql_query($query);
  if (mysql_affected_rows() == 1) {
    $message = '<P>Your rating has been submitted.</P>';
  } else {
    error_log(mysql_error());
    $message = '<P>Something went wrong with your rating
attempt.  Try again.</P>';
    $message .= $rate_form;
  }
}
?>

<HTML>
<HEAD>
<STYLE TYPE="text/css">
<!--
BODY, P     {color: black; font-family: verdana;
font-size: 10 pt}
H1      {color: black; font-family: arial; font-size: 12 pt}
-->
</STYLE>
</HEAD>

<BODY>
<TABLE BORDER=0 CELLPADDING=10 WIDTH=100%>
<TR>
<TD BGCOLOR="#F0F8FF" ALIGN=CENTER VALIGN=TOP WIDTH=17%>
</TD>
<TD BGCOLOR="#FFFFFF" ALIGN=LEFT VALIGN=TOP WIDTH=83%>
<H1>Rate your boss anonymously</H1>
<?php echo $message; ?>
```
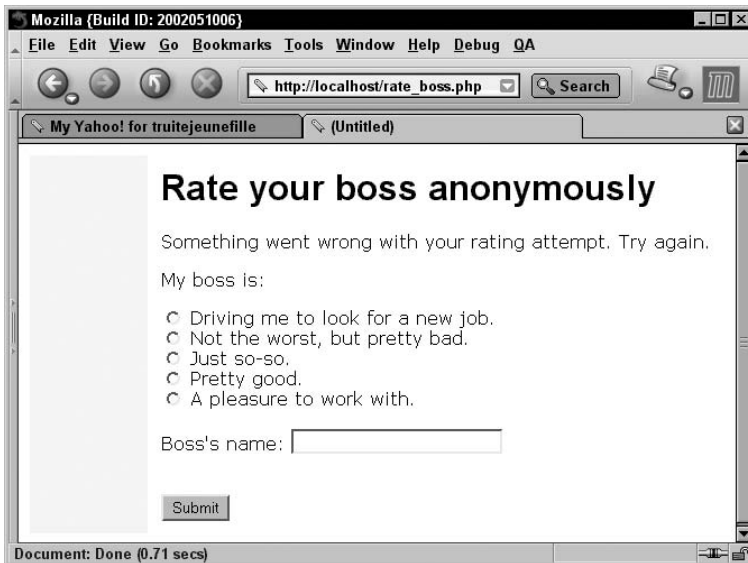
262

```
</TD>
</TR>
</TABLE>

</BODY>
</HTML>
```

Figure 17-2 shows the rating form after an error has occurred.

A multiple self-submitting form

Some of you might be thinking, "Hey, wait! You said logic always comes before display — but then you started this script with a bunch of HTML." Very observant — but not quite right. Look closely, and you will realize that we are merely setting a bunch of text to a couple of variable strings (`$reg_form` and `$rate_form`). In the entire PHP section, we actually don't display anything. We merely construct a string, `$message`, which will be plugged in to the HTML at the bottom. If we took away the HTML, you would see a blank page in the browser. So it's okay to assemble the text you're going to want to display in the logic part; just don't echo it out to the browser until the end.

Another issue with self-submitted forms is navigation. With the traditional HTML form, navigation is strictly one-way: form to handler to whatever navigational device (if any) the designer decrees.

263

Self-submitted forms need not conform to this rule, however. In each individual instance, you need to decide:

- Whether the form can be resubmitted multiple times by the user, in whole or in part
- Whether the user decides when to move on by clicking a link or the form moves users along automatically
- Whether you need to pass variables on to the next page, hidden or in plain view
- Whether you want to control where the user can go next or if you want to give users multiple choices

The answers to these questions will determine whether you need a control, another form, a simple link or button, or multiple links.

> **TIP** Whatever you decide about navigation, remember to provide plenty of text that clearly explains what's going to happen at every step. Because PHP gives you so much flexibility with forms, new users' default expectations may be crossed up, and they could end up uncertain whether they accomplished their mission with your form.

Editing Data with an HTML Form

PHP is brilliant at putting variables into a database, but it really shines when taking data from a database, displaying it in a form to be edited, and then putting it back in the database. Its HTML-embeddedness, easy variable passing, and slick database connectivity are at their best in this kind of job. These techniques are extremely useful, because you will find a million occasions to edit data you're storing in a database.

Let's look at the specific kinds of HTML `FORM` data elements and how they are handled.

TEXT and TEXTAREA

`TEXT` and `TEXTAREA` are the most straightforward types because they enjoy an unambiguous one-to-one relationship between identifier and content. In other words, there is only one possible `VALUE` per `NAME`. You just pull the data field from the database and display it in the form by referencing the appropriate array value, as shown in Figure 17-3.

Listing 17-5, `comment_edit.php`, takes a comment out of the database and allows you to edit it.

> **TIP** You may need to use the `stripslashes` **function when displaying** `TEXTAREA` **and** `TEXT` **if there's any chance the values might have single quotation marks or apostrophes. Watch out for people with apostrophe'd names like O'Malley or D'Nesh!**

FIGURE 17-3

Displaying text for editing

LISTING 17-5

Editing data from database (comment_edit.php)

```php
<?php

// Open connection to the database
mysql_connect("localhost", "phpuser", "sesame")
or die("Failure to communicate with database");
mysql_select_db("test");

if (isset($_POST['submit']) && $_POST['submit'] == 'Submit') {
  // Format the data
  $comment_id = mysql_real_escape_string($_POST['comment_id']);
  $comment_header = mysql_real_escape_string($_POST['comment_header']);
  $as_comment_header = mysql_real_escape_string($comment_header);
  $comment = mysql_real_escape_string($_POST['comment']);
  $as_comment = mysql_real_escape_string($_POST['comment']);

  // Update values
```

```
   $query = "UPDATE comments
             SET comment_header = '$as_comment_header',
             comment = '$as_comment'
             WHERE ID = $comment_id";
   $result = mysql_query($query);
   if (mysql_affected_rows() == 1) {
     $success_msg = '<P>Your comment has been updated.</P>';
   } else {
     error_log(mysql_error());
     $success_msg = '<P>Something went wrong.</P>';
   }
} else {
   // Get the comment header and comment
   $comment_id = mysql_real_escape_string($_GET['comment_id']);
   $query = "SELECT comment_header, comment
             FROM comments
             WHERE ID = $comment_id";
   $result = mysql_query($query);
   $comment_arr = mysql_fetch_array($result);
   $comment_header = stripslashes($comment_arr[0]);
   $comment = stripslashes($comment_arr[1]);
}


$thispage = "comment_edit.php"; //Have to do this for heredoc

$form_page = <<< EOFORMPAGE
<STYLE TYPE="text/css">
<!--
BODY, P     {color: black; font-family: verdana;
font-size: 10 pt}
H1      {color: black; font-family: arial; font-size: 12 pt}
-->
</STYLE>
</HEAD>

<BODY>
<TABLE BORDER=0 CELLPADDING=10 WIDTH=100%>
<TR>
<TD BGCOLOR="#F0F8FF" ALIGN=CENTER VALIGN=TOP WIDTH=17%>
</TD>
<TD BGCOLOR="#FFFFFF" ALIGN=LEFT VALIGN=TOP WIDTH=83%>
<H1>Comment edit</H1>

$success_msg
<FORM METHOD="post" ACTION="$thispage">
<INPUT TYPE="text" SIZE="40" NAME="comment_header"
VALUE="$comment_header"><BR><BR>
```

```
<TEXTAREA NAME="comment" ROWS=10 COLS=50>$comment</TEXTAREA>
<BR><BR>
<INPUT TYPE="hidden" NAME="comment_id" VALUE="$comment_id">
<INPUT TYPE="submit" NAME="submit" VALUE="Submit">
</FORM>

</TD></TR></TABLE>
</BODY>
</HTML>
EOFORMPAGE;
echo $form_page;
?>
```

> **TIP** Remember that in an HTML form integers and doubles must use the `TEXT` or `TEXTAREA` type, as there is no specifically numeric HTML form field type.

CHECKBOX

The `CHECKBOX` type has only one possible value per input: off (unchecked) or on (checked). The database field that records this information is almost always going to be a small integer or bit type with values 0 and 1 corresponding to unchecked or checked check boxes. Figure 17-4 shows a common type of check box being edited.

Listing 17-6 demonstrates how to use a check box to display and change a Boolean value.

FIGURE 17-4

A prepopulated check box

LISTING 17-6

Check box displaying boolean data from database (optout.php)

```php
<?php

// Open connection to the database
mysql_connect("localhost", "phpuser", "sesame")
or die("Failure to communicate with database");
mysql_select_db("test");

// If the form has been submitted, record the preference and
// redisplay
if (isset($_POST['submit'] && $_POST['submit'] == 'Submit') {
  $email = $_POST['email'];
  $as_email = mysql_real_escape_string($_POST['email']);
  if (isSet($_POST['OptOut'] && $_POST['OptOut'] == 1) {
    $optout = 1;
  } else {
    $optout = 0;
  }

  // Update value
  $query = "UPDATE checkbox
            SET BoxValue = $optout
            WHERE BoxName = 'OptOut'
            AND email = '$as_email'";
  $result = mysql_query($query);
  if (mysql_error() == "") {
    $success_msg = '<P>Your preference has been updated.</P>';
  } else {
    error_log(mysql_error());
    $success_msg = '<P>Something went wrong.</P>';
  }
  // Get the value
  $query = "SELECT BoxValue FROM checkbox
            WHERE BoxName = 'OptOut' AND email = '$as_email'";
  $result = mysql_query($query);
  $optout = mysql_result($result, 0, 0);

  if ($optout == 0) {

    $checked = "";
  } elseif ($optout == 1) {
    $checked = 'CHECKED';
  }
}


// Now display the page
$thispage = "optout.php"; //Have to do this for heredoc
```

```
$form_page = <<< EOFORMPAGE
<HTML>
<HEAD>
<TITLE>Semi-sleazy opt-in form</TITLE>
</HEAD>

<BODY>
$success_msg
<FORM METHOD=POST ACTION="$thispage">
Email address:
<INPUT TYPE="text" NAME="email" SIZE=25 VALUE="$email">
<BR><BR>
<FONT SIZE=+4>Please send me lots of e-mail bulletins!</FONT>
<BR>
<FONT SIZE=-2>opt out by clicking this tiny checkbox</FONT>
<INPUT TYPE="checkbox"  NAME="OptOut" VALUE=1 $checked><BR><BR>
<INPUT TYPE="submit" NAME="submit" VALUE="Submit">
</FORM>

</BODY>
</HTML>
EOFORMPAGE;
echo $form_page;

?>
```

Although each check box is capable of expressing only a fixed chunk of data, check boxes are often used in bunches to convey more complex aggregate meanings. Look at the check box grouping in Figure 17-5.

RADIO

RADIO data elements allow for a one-to-many relationship between identifier and value. In other words, they have multiple possible values, but only one can be predisplayed or selected. They are best for small sets of options, generally between two and ten, which need more than a word or two of text to identify themselves.

Unfortunately, it's somewhat more difficult to represent stored data in a radio button than in a check box or text field. This is because there is only one possible value for text or a textarea and only two possible values for a check box — but radio buttons can have more than two possible values. Therefore, you will have to output part of the actual form with PHP. This looks a little bit less neat than the styles we employed previously, so you have to go to a little more trouble to have an easily readable script. Again, the user interface experience allowed by radio buttons is worth the extra trouble it gives to the web developer.

In the example in Figure 17-6 and accompanying code, we are assembling a series of radio buttons that display preference data from the database.

269

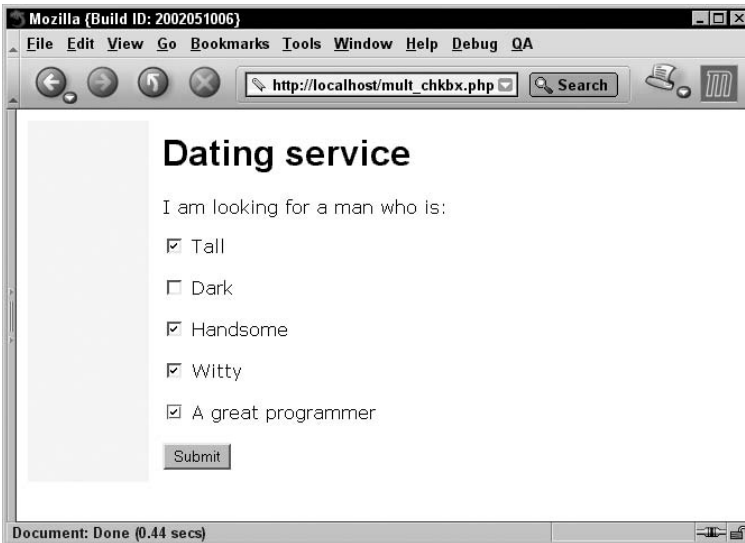**FIGURE 17-5**

A cluster of check boxes

FIGURE 17-6

Prepopulated radio buttons

Listing 17-7 shows the code for Figure 17-6, which shows how to edit forms with radio buttons.

LISTING 17-7

Radio buttons displaying boolean data from database (date_prefs.php)

```php
<?php

// Subscriber ID is stored in a cookie on the user's browser
if (isset($_COOKIE['userID'])) {
    $sub_id = mysql_real_escape_string($_COOKIE['userID']);
}

if (!isset($sub_id)) {
  die("Cookie Not Found.");
}

// Open connection to the database
mysql_connect("localhost", "mysqluser", "sesame")
or die("Failure to communicate with database");
mysql_select_db("test");

// If the form has been submitted, record the preferences
if (isset($_POST['submit']) && $_POST['submit'] == 'Submit') {
  $height = mysql_real_escape_string($_POST['height']);
  $haircolor = mysql_real_escape_string($_POST['haircolor']);
  $edu = mysql_real_escape_string($_POST['edu']);

  // Update value
  $query = "UPDATE qualities
          SET height = $height, haircolor = $haircolor,
              edu = $edu
          WHERE subscriber = $sub_id";
  $result = mysql_query($query);
  if (mysql_affected_rows() == 1) {
    $success_msg = '<P>Your preferences have been updated.</P>';
  } else {
    error_log(mysql_error());
    $success_msg = '<P>Something went wrong.</P>';
  }

}

// Get the values
$query = "SELECT height, haircolor, edu FROM qualities
        WHERE subscriber = $sub_id";
$result = mysql_query($query);
$pref_arr = mysql_fetch_array($result);
$height = $pref_arr[0];
```

271

```
$haircolor = $pref_arr[1];
$edu = $pref_arr[2];

// Assemble the radio button part of the form
if ($height == 1) {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"height\" VALUE=1
checked> Short<BR>\n";
} else {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"height\" VALUE=1> Short<BR>\n";
}
if ($height == 2) {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"height\" VALUE=2
checked> Average height<BR>\n";
} else {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"height\" VALUE=2>
Average height<BR>\n";
}
if ($height == 3) {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"height\" VALUE=3
checked> Tall<BR>\n";
} else {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"height\" VALUE=3>
Tall<BR>\n";
}
if ($height == 0) {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"height\" VALUE=0
checked> Doesn't matter<BR><BR>\n";
} else {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"height\" VALUE=0>
Doesn't matter<BR><BR>\n";
}

if ($haircolor == 1) {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"haircolor\" VALUE=1
checked> Blonde<BR>\n";
} else {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"haircolor\" VALUE=1>
Blonde<BR>\n";
}
if ($haircolor == 2) {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"haircolor\" VALUE=2
checked> Brunette<BR>\n";
} else {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"haircolor\" VALUE=2>
Brunette<BR>\n";
}
if ($haircolor == 3) {
  $radio_str .= "<INPUT TYPE=RADIO NAME=\"haircolor\" VALUE=3
checked> Redhead<BR>\n";
} else {
```

```
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"haircolor\" VALUE=3>
Redhead<BR>\n";
}
if ($haircolor == 0) {
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"haircolor\" VALUE=0
checked> Doesn't matter<BR><BR>\n";
} else {
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"haircolor\" VALUE=0>
Doesn't matter<BR><BR>\n";
}

if ($edu == 1) {
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"edu\" VALUE=1 checked>
High school graduate<BR>\n";
} else {
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"edu\" VALUE=1> High
school graduate<BR>\n";
}
if ($edu == 2) {
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"edu\" VALUE=2 checked>
College graduate<BR>\n";
} else {
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"edu\" VALUE=2> College
graduate<BR>\n";
}
if ($edu == 3) {
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"edu\" VALUE=3 checked>
Advanced degree holder<BR>\n";
} else {
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"edu\" VALUE=3>
Advanced degree holder<BR>\n";
}
if ($edu == 0) {
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"edu\" VALUE=0 checked>
Doesn't matter<BR><BR>\n";
} else {
    $radio_str .= "<INPUT TYPE=RADIO NAME=\"edu\" VALUE=0> Doesn't
matter<BR><BR>\n";
}


// Now display the page
$thispage = "date_prefs.php"; //Have to do this for heredoc

$form_page = <<< EOFORMPAGE
<HTML>
<HEAD>
<STYLE TYPE="text/css">
<!--
BODY, P      {color: black; font-family: verdana;
```

273

```
font-size: 10 pt}
H1         {color: black; font-family: arial; font-size: 12 pt}
-->
</STYLE>
</HEAD>

<BODY>
<TABLE BORDER=0 CELLPADDING=10 WIDTH=100%>
<TR>
<TD BGCOLOR="#F0F8FF" ALIGN=CENTER VALIGN=TOP WIDTH=17%>
</TD>
<TD BGCOLOR="#FFFFFF" ALIGN=LEFT VALIGN=TOP WIDTH=83%>
<H1>Dating service</H1>
$success_msg
<P>I am looking for a girl who is:</P>
<FORM METHOD=POST ACTION="$thispage">
$radio_str
<INPUT TYPE=SUBMIT NAME="submit" VALUE="Submit">
</FORM>

</TD>
</TR>
</TABLE>
</BODY>
</HTML>
EOFORMPAGE;
echo $form_page;

?>
```
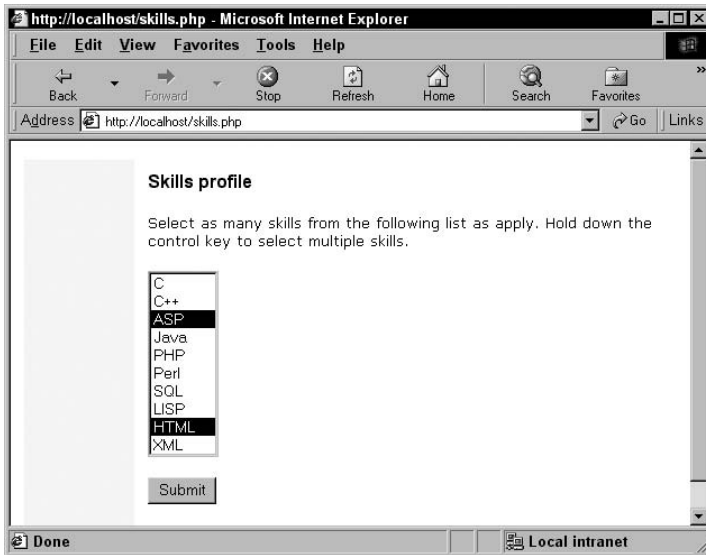
SELECT

The SELECT field type is perhaps the most interesting of all. It can handle the largest number of options, and it also allows the user to select multiple options that can be passed back to the database using arrays.

CROSS-REF **See Chapter 39 for ideas about using JavaScript to make even more interesting** SELECT **forms.**

In Figure 17-7, we are using the SELECT form element with multiple options. In PHP, this is done by creating an array of the multiple selected option values to pass to the form handler. You set up the array in the HTML form by declaring the MULTIPLE attribute of the SELECT element and by naming the SELECT element something like $val[] — in other words, appending a set of square brackets to the variable name. This will indicate to PHP that it's dealing with an array rather than a single variable, and it will construct the array appropriately with the multiple selected values. When the array gets to the form handler, you will need to deal with the values as you would any array's values — by dereferencing, or by listing out the contents of the array.

A prepopulated select with multiple choices

Listing 17-8 shows the code for Figure 17-7, which demonstrates how to display and edit a select list with multiple options.

Select list displaying database values (skills_profile.php)

```php
<?php

if (isset($_COOKIE['user_id'])) {
    $user_id = mysql_real_escape_string($_COOKIE['user_id']);
}


if (!isset($user_id)) {
  die("Cookie Not Found.");
}

// Open connection to the database
mysql_connect("localhost", "mysqluser", "sesame")
or die("Database error!");
mysql_select_db("test");
```

```php
if (isset($_POST['submit'] && $_POST['submit'] == 'Submit') {

  // Delete this user's skills
  $query2 = "DELETE FROM user_skill
             WHERE user_id = $user_id";
  $result2 = mysql_query($query2);

  foreach ($_POST['skills'] as $val) {
    $cleanVal = mysql_real_escape_string($val);
    $query = "INSERT INTO user_skill (ID, user_id, skill_id)
              VALUES (NULL, $user_id, $cleanVal)";
    $result = mysql_query($query);
    if (mysql_affected_rows() == 1) {
      continue;
    } else {
      error_log(mysql_error());
      $error_msg = '<P>Something went wrong</P>';
      break;
    }
  }
}

// Get all the results
$query = "SELECT * FROM skills";
$result = mysql_query($query);

// Download this user's skills
$query1 = "SELECT skill_id
           FROM user_skill
           WHERE user_id = $user_id";
$result1 = mysql_query($query1);
while ($user_skill = mysql_fetch_array($result1)) {
  $skill_id = $user_skill[0];
  $user_skill_arr[$skill_id] = $skill_id;
}


while ($skills = mysql_fetch_array($result)) {
  $key = $skills[0];
  if ($key == $user_skill_arr[$key]) {
    $select_str .= "<OPTION VALUE=\"$key\"
SELECTED>$skills[1]\n";
  } else {
    $select_str .= "<OPTION VALUE=\"$key\">$skills[1]\n";
  }
}


$thispage = "skills_profile.php"; //Have to do this for heredoc
```

```
$form_str = <<< EOFORMSTR
<HTML>
<HEAD>
<STYLE TYPE="text/css">
<!--
BODY, P      {color: black; font-family: verdana;
font-size: 10 pt}
H1        {color: black; font-family: arial; font-size: 12 pt}
-->
</STYLE>
</HEAD>

<BODY>
<TABLE BORDER=0 CELLPADDING=10 WIDTH=100%>
<TR>
<TD BGCOLOR="#F0F8FF" ALIGN=CENTER VALIGN=TOP WIDTH=17%>
</TD>
<TD BGCOLOR="#FFFFFF" ALIGN=LEFT VALIGN=TOP WIDTH=83%>
<H1>Skills profile</H1>
<P>Select as many skills from the following list as apply.  Hold
down the control key to select multiple skills.</P>
$error_msg

<FORM METHOD=POST ACTION="$thispage">
<SELECT NAME="skills[]" SIZE=10 MULTIPLE>
$select_str
</SELECT>
<BR><BR>
<INPUT TYPE="submit" NAME="submit" VALUE="Submit">
</FORM>

</TD></TR></TABLE>
</BODY></HTML>
EOFORMSTR;
echo $form_str;

?>
```

Summary

PHP is an extremely powerful form-handling tool, especially in conjunction with a database. You can use PHP to display database-stored data as form values, and of course, you can also store form-generated data in the database.

To prepare your HTML forms to work smoothly with PHP, you need to follow a few simple rules. First and foremost, never use data that comes from the user directly in a database call or query. This means using the `mysql_real_escape_string()` function on any `$_POST`, `$_GET`, and `$_COOKIE` values. Also, remember always to name every single form element — the HTML standard itself doesn't require this, but PHP does because the element names will become variable names in the form handler. One method that is sometimes helpful is to match the form element name to the corresponding database field name so that they are easy to remember, perhaps prefixing form variables with `frm` or something similar to help distinguish them from their database counterparts in code. PHP also allows you to make clever use of hidden form inputs and of multiple `SELECT` options, which should be delineated with square brackets (denoting an array) after the element name.

You have the choice with PHP to have separate HTML forms and PHP form handlers or to combine the two in a PHP script. The latter option is arguably the more powerful, but it can also be more difficult to work with and maintain. You will need to set a variable within the form to indicate whether the entries have been submitted; the PHP logic should be placed before the HTML display. You can even have multiple forms on one page that are handled by the same PHP script.

Chapter 18

Improving Database Efficiency

This quick chapter is for people making database-enabled PHP web sites who suspect that they are doing things awkwardly or inefficiently. Maybe you are new to databases, or maybe you know there must be a way to speed things up just because your pages are loading unacceptably slowly.

We offer some tips and tricks for making things run faster, and we show you some common ways that database systems can save you from writing unnecessary PHP code. As usual, some of our code examples will use MySQL functions, although the lessons are mostly general and independent of particular database implementations.

Connections — Reduce, Reuse, Recycle

One important thing to realize is that establishing an initial connection with a database is never a cheap operation in terms of resource usage and time. Unless your PHP script is doing some unusually computationally intensive work, the overall database interaction will be the most time- and resource-intensive part of your code, and it is frequently true that the establishment of a connection is the most expensive (in terms of resource usage) part of code that interacts with a database, even if the connection is only established once in serving the page.

You have two potentially competing goals here. On one hand, you want to minimize the number of times your code makes the time-consuming call to open an entirely new database connection. This argues for leaving connections open during the course of page execution, rather than closing and reopening. On the other hand, there are sometimes hard limits on the number of simultaneous connections that a database program can support. This might argue for closing connections whenever possible in hopes that less connected time per script might allow more scripts to execute simultaneously.

In our experience, however, most web scripts are evanescent enough that it is never worth the overhead to close and reopen a database connection within one page's execution. If you want to minimize total time connected, open the connection immediately before the first call to the database, and close it immediately after the last one.

A bad example: one connection per statement

The first bad example seems stylistically reasonable in one sense because it uses a function to eliminate repetitive code.

```php
<?php
function box_query ($query, $user, $pass, $db)
{
  $my_connection =
    mysql_connect('localhost', $user, $pass)
    or die("Couldn't connect to database");
  mysql_select_db($db, $my_connection)
    or die("Couldn't select database");
  $result_id = mysql_query($query, $my_connection)
    or die(mysql_error());
  print("<H3>Results for query:  $query</H3>");
  print("<TABLE>");
  while ($row = mysql_fetch_row($result_id))
  {
    print("<TR>");
    $row_length = mysql_num_fields($result_id);
    for ($x = 0; $x < $row_length; $x++)
    {
      $entry = $row[$x];
      print("<TD>$entry</TD>");
    }
    print("</TR>\n");
  }
  print("</TABLE>");
  mysql_close($my_connection);
}
/* code that uses box_query() */
?>
```

The idea is that we take a function that packages up an arbitrary MySQL query and displays the returned data in an attractive HTML table. The main virtue of this function as defined is that it is very self-contained — it opens its own database connection for its own purposes, and then it disposes of that connection when the function is done.

The preceding code is fine if we expect to display only one such table per page. If we use this function more than once per page, however, we will find ourselves opening and closing connections every time the function is invoked, which is bound to be less efficient than leaving the connection open. One approach is to leave a single connection open for as long as it is needed in the execution of a single page's script. Applying this rule to the preceding function would mean rewriting it so that it takes a connection as argument (or implicitly uses a connection opened at the beginning of the script) and then opening a single connection per page.

Multiple results don't need multiple connections

One thing that surprised us the very first time we saw web-database scripting was that, with many database programs, it is possible to retain the results from more than one query at one time, even though only one connection has been opened. For example, with a MySQL database you can do something like this:

```
mysql_connect('localhost', $user, $pass); //opens connection
mysql_select_db('scienceguide');
$author_result = mysql_query("SELECT ID FROM author")
                 or die(mysql_error());
while ($author_row = mysql_fetch_row($author_result))
{
  $book_result =
    mysql_query("SELECT title FROM book
               WHERE authorID = {$author_row[0]}")
    or die(mysql_error());
while ($book_row = mysql_fetch_row($book_result))
  {
    $title = $book_row[0];
    print("$title<BR>");
  }
}
```

This would print titles of books after retrieving them from the book table, using IDs from rows retrieved from the author table. If we assume there is not more than one author per book, then this is an extremely inefficient way to retrieve the data (see the section "Making the Database Work for You" later in this chapter), but it illustrates that two different result sets (identified by the variables $author_result and $book_result) can be actively used at the same time, after having been retrieved over a single connection.

Persistent connections

Finally, if you become convinced that the sheer overhead of opening new database connections is killing the performance of your application, you might want to investigate opening *persistent* connections. Unlike regular database connections, these connections are not automatically killed when your page exits (or even when `mysql_close()` is called) but are saved in a pool for future use. The first time one of your scripts opens such a connection, it is opened in the same resource-intensive way as with a regular database connection. The next script that executes, however, might get that very same connection in response to its request, which saves the cost of reopening a fresh connection. (The previous connection will be reused only if the parameters of the new request are identical.)

> **NOTE**   **Persistent database connections work only in the module installation of PHP. If you ask for a persistent connection in the CGI version, you will simply get a regular connection.**

The PHP function to request such a persistent connection for MySQL is `mysql_pconnect()`, which is used in exactly the same way as `mysql_connect()`. This naming convention seems to be stable across PHP functions for the different databases — if you use a particular DB `connect` function, you should consult the documentation to see if a `pconnect` version exists.

> **NOTE**   **Other than offering a particular kind of increased efficiency, persistent database connections do not provide any functionality beyond that of regular database connections. In particular, you should not expect persistent connections to have any memory of previous queries or of variables from previous page executions.**

Indexing and Table Design

MySQL is a pretty fast database, even absent any serious design considerations. In a lot of installations and applications, the database-design part of your job may be no more difficult than creating a single basic table with four or five fields in anticipation of holding no more than a few hundred records. However, as your database needs grow, your database itself will doubtless grow as well — in both size and complexity. That's no sweat for a good RDBMS: MySQL and other products in this class excel at handling these needs. Still, careful choice of both indexes and field types when designing tables can be crucial for performance as your tables get larger.

Indexing

Probably the first thing to investigate when `SELECT` statements are slow is whether you have defined appropriate indexes.

What is an index?

Wikipedia defines an index in the following manner: "A database index is a data structure that improves the speed of operations in a table" (`http://en.wikipedia.org/wiki/Index_(database)`). An

index on a table field is an indication by a database designer to the database system that any searches made on that field should be fast. Usually, this is implemented by the RDBMS as a side table that maintains all the values for the field in order, and maps them to rows in the original table. Whenever a SELECT statement has a WHERE condition that mentions the indexed field, the side table is consulted to locate the rows that have the desired values for the field. The ordering of the side table means that the database system can do fast lookups (for example, using binary search).

Indexing tradeoffs

There are two mantras to keep in mind when thinking about creating indexes:

- SELECT statements that filter on unindexed fields may require full table scans.
- While indexes speed up SELECT statements, they slow down INSERTs, UPDATEs, and DELETEs.

To see why both these statements are true, imagine that we gave you a large telephone book (sorted by last name) and asked you to find us everyone in the book with a first name of 'Zachary'. Unfortunately, it's difficult to see how to accomplish this without looking through the entire book.

A database system trying to execute a statement like:

```
SELECT lastname FROM phonebook WHERE firstname = 'Zachary'
```

is in exactly the same situation, if there is no index on the field 'firstname'. In database parlance, the system must resort to a *full table scan*, meaning that every row in the table is inspected.

If your job were to do this phonebook lookup frequently, you might find it worth your while to commission an extra *index* (in the book-publishing sense) that listed all the first names in order, along with the page numbers and associated last names. Once the newly indexed phone book arrived, your job would become a lot easier.

The bad news is that as soon as the new phone book arrived, we decided to promote you. Congratulations! Your new job is to keep the phone book up to date (including, of course, any associated indexes). Here is a list of 10,000 new customers, 8,000 people who have moved away, and 45 people who have had name changes. Now the firstname index is a burden rather than a benefit. Again, it's the same with the database system — the indexes that make lookups faster are a maintenance burden when the data must be modified.

The general lesson is that you should consider indexes on fields that you use frequently in the WHERE clauses of SELECT statements, especially when the data-modifying statements (INSERT, UPDATE, DELETE) will be used rarely. If modification is much more common than lookup, indexes make less sense.

Now we move on to the specifics of using indexes in MySQL, beginning with the most common usage: a single index that uniquely identifies each table row.

Primary keys

Simply put, a primary key is a field in a table that uniquely identifies each record in that table. A good primary key choice needs to meet a few criteria:

- A primary key should be of an integer type. These may vary some from one database tool to the next, but in MySQL, they are TINYINT, SMALLINT, MEDIUMINT, INT, and BIGINT. Refer to the MySQL online documentation for the current ranges and other properties of these types.

- A primary key should not return a null value. Your column definition should contain the SQL keyword NOT NULL. In fact, many databases, MySQL included, will not let you designate a primary key that is capable of returning a null value.

- A primary key MUST be unique. That's the point, isn't it? And because a primary key must be unique, it should also have an auto-increment feature set. Most databases offer this, and most call it the same thing.

CAUTION **Auto-increment and its use are often debated. In your Internet travels, you'll come across those who don't like auto-increment and variously describe it as an accident waiting to happen or a cop out. To be honest, there are some meritorious arguments in this vein. However, we believe the benefits significantly outweigh the concerns. The alternatives are either expensive database calls to determine what key values are available or to generate an ID programmatically and then insert it with your SQL statement. Neither of these is as reliable nor worry free as auto-increment.**

If you've already forged ahead and created some database tables of your own without a primary key, consider the fields you have already created. Does one of these meet the tests described previously? It may be that you have wisely foreseen or intuited this need and created something like it already. If this field exists, but lacks one or more of the components, you can alter it with a SQL statement like the following:

```
ALTER TABLE 'my_table' CHANGE 'existing_field' 'my_key' SMALLINT
NOT NULL AUTO_INCREMENT PRIMARY KEY
```

Or if your field already has all the necessary characteristics, you can simply make it the primary key like this:

```
ALTER TABLE 'my_table' ADD PRIMARY KEY ('my_key')
```

In the first statement, we indicate that we are altering a table and indicate which table we want to operate on. CHANGE further indicates that we are changing a field's properties and indicating which field with its quoted existing name. We can then specify a name that may indicate more specifically what sort of field it is and set the relevant properties in one fell swoop.

If you don't already have an appropriate field choice, the syntax doesn't change much:

```
ALTER TABLE 'my_table' ADD 'my_key' SMALLINT NOT NULL
AUTO_INCREMENT PRIMARY KEY
```

Finally, you may just be creating your table for the first time. If that's the case, you simply need to include the following field definition in your table `create` statement:

```
ID SMALLINT UNSIGNED AUTO_INREMENT NOT NULL PRIMARY KEY
```

where ID is the name you've assigned to your primary key. There's nothing magical about this name; you can call it Fido if you want, but ID is a good, meaningful self-descriptive name.

So now you've got a primary key. What's it good for? Well, it helps define the master record in a one-to-many relationship. Its other properties enforce an unambiguous identity for each record, such that the SQL statement `delete from 'my_table' where id = 12` can have only one possible result. Phew, and you thought you just blew that whole table away.

Creation of the primary key also has the net effect of speeding up queries that join tables on this unique ID because in the process of making it a primary key, we made it an index as well. An index is stored separately by MySQL and operates transparently to the end user.

When you are defining a relationship in your SQL, the child table — the *many* side of the one-to-many relationship — will also store a copy of the master table's primary key value. But it will store it once for every record that is a child of the parent record, making it unsuitable for use as a primary key. You may still wish to define a primary key for each record in the child table — in fact, it's a good idea to do so, but you won't be able to define a primary key on this particular field because values may not be unique to this column. On the other hand, you still want to improve the process MySQL uses to locate related records for queries that perform joins. That works out alright, because MySQL can still index a field without making it a primary key:

```
ALTER TABLE `child table' ADD INDEX MyIndex (child_id)
```

This will work great for an existing field, but as before, you may need to create a suitable field for this purpose:

```
ALTER TABLE 'child_table' ADD 'child_id' SMALLINT NOT NULL
```

Then make the field an index:

```
ALTER TABLE 'child_table' ADD INDEX ('child_id')
```

Everything including the kitchen sink

Indexes are almost a requirement for speedy, efficient joins. Even those most ardently concerned about things like disk space will rarely find room to argue about the merits of an index that speeds up the definition of relationships. More debatable, however, may be indexes that do not specifically operate on joins.

You can index virtually anything. Sure, binary data presents some problems and is almost always an ill-advised choice for indexing, but strings, the larger text fields, and numbers (including floats and decimals) are all fair game. Aside from defining a relationship, the only other overriding

qualification for index candidacy is that it should be something you're likely to use in the `WHERE` clause of your SQL statement.

Let's say you want to create a membership directory for your local Linux Users Group and you want members to be able to find other members in the same part of town so that they can easily get together for a drink or a movie. If you're like us, you're probably thinking Zip code. Excellent choice. A universally used (at least in the U.S.), well-documented, predictable and fairly stable search criterion. Of course, you don't have to index this field:

```
SELECT name, phone from members where zip = '32223'
```

will get you an answer, the same answer in fact, with or without an index. On a table with 100 or so records, you'll get your answer instantaneously — again, with or without an index.

But maybe you have several hundred, perhaps even thousands of members. An index may just speed up this search. Add one and try your search again:

```
ALTER TABLE 'members' ADD INDEX ('zip')
```

Perhaps do it while watching the output of Linux's `ps` or `top` commands. Perhaps you'll see user discernible improvement; perhaps you'll need a professional diagnostic tool of some kind to measure what just happened; perhaps your performance improvement will be measured in nanoseconds. The point is, at some number of records, you almost certainly will see an improvement at each of these levels. It will be up to you as the designer to determine whether the benefits justify the tradeoffs.

What are the tradeoffs? Disk space, for one. Depending on the number of records and the size of the field, an index can increase storage requirements by nearly as much as the table size itself. If you've got 80GBs of storage, you probably don't care. If you're on a 50MB shared hosting plan, you probably care very much. Another tradeoff is that although `SELECT` operations benefit, `INSERT`, `UPDATE`, and `DELETE` operations actually take longer because the indexes must be updated each time one of these is performed. The good thing about an index is that it's not irreversible. Try an index on anything you think might be useful, measure the performance improvement, and weigh it against what you may or may not be giving up to get that improvement.

Other types of indexes

There are a couple other types of indexes, or more appropriately, parameters to indexing functions, that specify how indexes work. Using them may have the net effect of making an index work better or worse. Again, consider each type, experiment and measure your results. It's a small effort to make with potentially huge dividends.

UNIQUE

Isn't that a primary key? Maybe. In MySQL at least, a primary key is by definition nothing more or less complicated than a `UNIQUE INDEX` with the name `PRIMARY`. If you find yourself defining a unique index, consider whether what you've got is really a primary key candidate. Social Security numbers, if your users are consistently willing to provide them, may work well in this regard. This

choice certainly meets the criteria and offers some additional advantages such as knowing what the primary key will be before you insert anything, enabling you to create master and child records without the intermediate call to `mysql_insert_id()`.

A phone number, on the other hand, may not be such a good choice. Sure, it's unique. It also is, or can be defined as, an integer. But you may wish to store phone numbers as a string to avoid some post-formatting for creating a readable display, such as parenthesizing an area code or inserting the traditional, if somewhat meaningless hyphen. But even if you are willing to forgo the aesthetic concerns, as an integer, a phone number is almost certainly larger than necessary. The largest possible phone number will store as 9,999,999,999. Yeah, that's what we said. This integer would require a field type of at least `INT`. You probably aren't going to store more than nine billion records. `SMALLINT` or `MEDIUMINT` would be better choices for a storage and searchable volume savings of 2^{18} or 2^9 bytes, respectively.

All that said, you can still use `UNIQUE` without having it as a primary key, and that is precisely why it exists. A `UNIQUE` attribute on a phone number field can still serve as a data integrity check, once again relieving you of the responsibility of performing the check programmatically (of course, you will still probably have to respond to the problem).

A unique index can be specified in MySQL like this:

```
ALTER TABLE 'members' ADD UNIQUE my_index ('phone')
```

Table design

In Chapter 14, we discussed table design pretty extensively; we're not going to recap all that information here. However, we do want to reiterate some points about field types because choice of table fields can have significant performance impact.

There are two interrelated concerns when choosing field types for a table: speed and size in memory. Your field definitions should anticipate the largest possible value that they may be asked to store, while not overanticipating and therefore creating unnecessarily huge tables with lots of unused space, both on disk and in memory. Appropriate field choices also come into play when choosing indexes for your table. Indexes are of the greatest benefit when they are set on a field type that is optimized for the type of data it is expected to hold. If, for example, you want an indexed number field where the count will never be more than 65,000 or so records, that index will perform more efficiently on the `SMALLINT` field type than it will on the `MEDIUMINT` field type, which allocates more space and therefore must search that extra space when attempting to isolate a specific value.

A similar principle holds true for the string types. Although there's some debate whether or not it's even advisable to index on a string column, that index will certainly perform more efficiently on a field that is defined precisely to the specifications of the data you will wish to store on it.

Earlier in this book, we pointed out that sometimes concerns about performance are so inflated that they border on the ridiculous. That's still the way we feel. It should not, however, appear inconsistent that we stress performance concerns now. This section and those that follow offer easily implemented design considerations that will collectively improve the performance of your databases.

Making the Database Work for You

Just as when you write code in a programming language, writing code that interacts with a database is an exercise in appropriate division of labor. People who write programming languages and databases have agreed to automate, standardize, and optimize certain tasks that come up over and over again in programming, so that programmers don't have to constantly reinvent the wheel when making their individual applications. The very general rule is that, unless you're willing to spend a lot of energy in optimizing code for your special case, you are better off using a database-provided facility than trying to invent your own solution for the same task.

It's probably faster than you are

Database programs are judged partly on their speed, so database programmers devote a large portion of their effort toward ensuring that queries execute as quickly as possible. In particular, any searching or sorting of the contents of a database is best done within that database (if possible) rather than by your own code.

A bad example: looping, not restricting

For example, take the following code fragment (and please don't laugh — we have actually seen code like this):

```
function print_first_name_bad ($lastname, $dbconnection)
{
  $query = "SELECT firstname, lastname FROM author";
  $result_id = mysql_query($query, $dbconnection)
    or die(mysql_error());
  while ($row = mysql_fetch_array($result_id))
  {
    if ($row['lastname'] == $lastname)
        print("The first name is " . $row['firstname']);
  }
}
```

When this code is handed a last name string and a database connection, it will print out associated first names, if any, in the "author table" of the database. For example, a call to `print_first_name_bad('Sagan', $dbconnection)` might produce the output:

```
The first name is Carl
```

If there were multiple authors in that table with the same last name, then multiple lines would be printed.

The problem here is that we don't need to grab all the data in this table, pull it through the narrow pipe of a connection, and then pick and choose from it on our side of the pipe. Instead, we should restrict the query with a `WHERE` clause:

```
function print_first_name_better ($lastname, $dbconnection)
{
  $query = "SELECT firstname, lastname FROM author
            WHERE lastname = '$lastname'";
  $result_id = mysql_query($query, $dbconnection)
     or die(mysql_error());
  while ($row = mysql_fetch_array($result_id))
  {
      print("The first name is " . $row['firstname']);
  }
}
```

The WHERE clause ensures that only the rows we care about are selected in the first place. Not only does this cut down on the data passed over the SQL connection, but the code used to locate the correct rows on the database side is almost certainly quicker than your PHP code.

Sorting and aggregating

Exactly the same argument applies if you find yourself writing code to sort results that have been returned from your database, or to count, average, or otherwise aggregate those results. In general, the ORDER BY syntax in SQL will allow you to presort your retrieved rows by any prioritized list of columns in the query, and that sort will probably be more efficient than either homegrown code or the PHP array-sorting functions. Similarly, rather than looping through DB rows to count, sum, or average a value, investigate whether the syntax of your particular DB's flavor of SQL supports the GROUP BY construct and in-query functions such as count(), sum(), and average(). In general, executing a query like:

```
$query = "SELECT count(ID) FROM author";
```

will be a radically more efficient approach to counting table rows than selecting them and iterating through them with a PHP looping construct.

Where possible, use MIN or MAX rather than sorting

Although it's good to let the database system do your sorting for you, it's even better to not have to sort at all. One task that is often addressed by unnecessary sorting is finding the minimum or maximum value in a set of result rows. You may see code like this:

```
$query = "SELECT ID FROM author ORDER BY ID limit 1;
// inefficient
```

This query will return a single ID from the author table after having sorted it in ascending order — in other words, the minimum ID. It does have the virtue that the actual result set returned is small, so it is a better approach for finding the minimum than using the same query without the limit clause and picking off the desired value from the top of that large result set. But if all we are interested in is the minimum (or maximum) value, there is no need to require the DB to figure out the rank order of all the other IDs that we are not interested in. A better solution is:

```
$query = select min(ID) from author; // efficient
```

289

The difference between these approaches will be imperceptible when your tables have only tens or hundreds of rows in them but will begin to matter as your tables grow to thousands or tens of thousands of rows in size.

Creating date and time fields

It is very common to want to associate a date and/or time with a row's worth of data. For instance, your table rows might represent requests made by your web site users, and the associated date/time is the time that that request hit your database.

Now, one way to insert or update date fields is to include a string that represents the desired date in a format parsable by your database. For example, if you want to set the `mydate` datetime field of all rows of `mytable` to a particular date, you might set up a query like this one:

```
$query = "UPDATE mytable SET mydate = '2003-11-24'";
```

and then send that query off for evaluation. (Unfortunately, the exact standards of readable date formats vary quite widely from one SQL database system to another. This particular date string means November 24, 2003, as far as MySQL is concerned.)

The preceding approach is fine, as long as you take care that the particular date string you send is, in fact, readable as a date by your DB. Things get more complicated if you need to construct such a string on the fly to represent a date that depends on the value of variables in your script.

The main thing to remember is that, with most database systems, there is no need to go through such contortions to set a field to the current date or time. Many have a current-date function that can be embedded directly in your query. For example, a MySQL version of the preceding query that sets the relevant date/time field to the current instant looks like this:

```
$query = "UPDATE mytable SET mydate = now()";
```

Note that the call to `now()` is not enclosed in single quotation marks, because it's a call to database function rather than a string to be interpreted by the database as data. The analogous query for Microsoft SQL Server looks like this:

```
$query = "UPDATE mytable SET mydate = getdate()";
```

Finally, even if the time you want stored is not that of the instant of execution, there may still be better alternatives than constructing readable date strings in your script. In addition to functions returning the current date, many versions of SQL offer functions for performing date arithmetic — start with a particular date/time, and then add or subtract years, months, or hours. In MySQL, these functions are:

- `date_add(date, date-interval)`
- `date_sub(date, date-interval)`

Here `date-interval` is a string that includes a number of time units and the type of unit. A MySQL query to set all rows to a time a week from now might look like this:

```
$query = "UPDATE mytable SET mydate = date_add(now(), INTERVAL 7
DAY)";
```

MySQL has a plethora of date and time related functions. See the MySQL documentation at: `http://dev.mysql.com/doc/refman/6.0/en/date-and-time-functions.html` for more information on all of the functions.

Finding the last inserted row

Another surprisingly helpful capability offered by some database systems is finding the ID of the last row inserted. This problem arises when you are trying to create a new database entry that is distributed across several database tables, each of which has an automatically incremented primary key. As an example, take the tables created by the following MySQL statements:

```
CREATE TABLE author (ID int primary key auto_increment,
                     lastname varchar(75),
                     firstname varchar (75));
CREATE TABLE book (ID int primary key auto_increment,
                   authorID int,
                    title varchar(100));
```

One intent of these statements is that the `book` table is linked to the `author` table by joining them so that `book.authorID` = `author.ID`. Another intent is that we don't have to worry about assigning unique ID fields for either table — the database will automatically assign them. Unfortunately, the combined intent leads to a problem. How do we write code that will gracefully insert a linked book-author pair, when both the author and the book are new to the database? If we insert a new author, the `ID` field of the inserted row will be automatically created by the database and so will not be a part of our SQL insert statement. How can we give the correct `authorID` to our new `book` row?

One possible strategy is to do something like the following (in MySQL):

```
$author_lastname = 'Feynman';
$author_firstname = 'Richard';
$book_title = 'The Character of Physical Law';
$author_insert = "INSERT INTO author (lastname, firstname)
    VALUES ('$author_lastname','$author_firstname')";
mysql_query($author_insert) OR die(mysql_error());
$author_id_query =
  "SELECT ID FROM author
   WHERE lastname = '$author_lastname'
   AND firstname = '$author_firstname'";
$author_id_result =
  mysql_query($author_id_query) OR die(mysql_error());
if (mysql_num_rows($author_id_result) <= 0)
  die("Inserted author not found!");
```

```
        else
          $author_row = mysql_fetch_row($author_id_result);
        $authorID = $author_row[0];
        $book_insert = "INSERT INTO book (authorID, title)
                       VALUES ($authorID, $book_title)";
        mysql_query($book_insert) OR die(mysql_error());
```

In this code, we create a new author row, use the last name and first name of the author to select the row we have just created, pull out the unique ID of that newly created row, and then incorporate that ID in a statement inserting a new row into the book table. This code would probably work in this particular instance, if we assume that the author's last name and first name are sufficient for unique identification. But for many databases, we will not be able to make such an assumption, which is, of course, why the convention of unique IDs developed in the first place.

A similar approach that is sometimes used is to insert a row (for example, into the author table) and then select the maximum ID from that table, on the theory that the highest row ID will be the one most recently inserted. If the most recently inserted row is, in fact, the one we just inserted, this will work like a charm. Unfortunately, this is exactly the kind of approach that appears to work when tested by a solitary user/programmer and then breaks when used with a real database server that is dealing with requests from multiple connections at the same time. The problem is that an insertion from another connection might well arrive in between our own insertion and the statement we send to retrieve the maximum ID to date, with the result that our second insertion is matched with an inappropriate ID.

The best solution, when it is available, is to have the database itself keep track of the last inserted ID in a retrievable way, and do this tracking on a *per-connection* basis, so that there are no worries about the synchronization issues in the previous paragraph. For MySQL users, PHP offers the function mysql_insert_id(), which takes a connection ID as argument and returns the auto-incremented ID of the last inserted row. We can use it to rewrite our previous code example:

```
        $author_lastname = 'Feynman';
        $author_firstname = 'Richard';
        $book_title = 'The Character of Physical Law';
        $author_insert = "INSERT INTO author (lastname, firstname)
                VALUES($author_insert) OR die(mysql_error());
        $authorID = mysql_insert_id();
        $book_insert = "INSERT INTO book (authorID, title)
                       VALUES ($authorID, '$book_title')";
        mysql_query($book_insert) OR die (mysql_error());
```

As with many PHP/MySQL functions, the connection argument to mysql_insert_id() is actually optional and defaults to the most recently opened connection.

In some other database systems, the ID of the most recent auto-increment is available (per session) as a "special" variable that can be embedded in the next query. In Microsoft SQL Server, for example, the variable is %%identity, which can be embedded in a query as follows to retrieve the last insert ID:

```
        $query = "SELECT @@identity";
```

Summary

Because database-related functionality is among the most resource-intensive things that PHP can do, you can become a hero by giving just a little thought to efficient coding practices. Particularly if your data-driven PHP scripts are sluggish, you want to learn to work with the database instead of against it.

The basic principles of database-intensive coding are simple. It costs a lot to open a connection to a database, so don't turn the tap on and off unnecessarily. Remember the pipe is narrow — you want to transport the bare minimum of data you need for each page. And take the time to learn all the functionality your particular database can offer you. SQL is really good at indexing, sorting, filtering, restricting, numbering, and grouping — use these powers rather than doing it less well and more slowly with PHP.

In Chapter 19, we move from these tips and stylistic concerns to problems and gotchas that can actually break your database code or give you unintended results.

Chapter 19

MySQL Gotchas

This chapter details some of the common difficulties that arise with using PHP and databases. The goal is to help you diagnose and solve problems more quickly and with less frustration. As usual, our specific code and function references are to MySQL (with one exception), although the set of gotchas is fairly independent across different databases.

CROSS-REF This chapter is about diagnosing and fixing PHP/database code that is genuinely broken — that is, it is not successfully retrieving data, or it is producing error messages. If your scripts are working, but too slowly, see Chapter 18.

No Connection

If you have a database call in your PHP script and the connection can't be opened, you will see a version of one of these two warning screens (depending on how high your error reporting levels are cranked up, and, to some extent, the precise cause of the problem).

The first possibility is the No Connection warning, as shown in Figure 19-1.

This option indicates a problem either with the MySQL server itself or with the path to `mysqld`. In its own special way, PHP is telling you that it knows about MySQL but can't hook up to it. This is the error you will see on a working PHP-MySQL installation if the database server crashes.

If the problem is on the PHP side, your error screen will look more like the one shown in Figure 19-2.

A No Connection warning
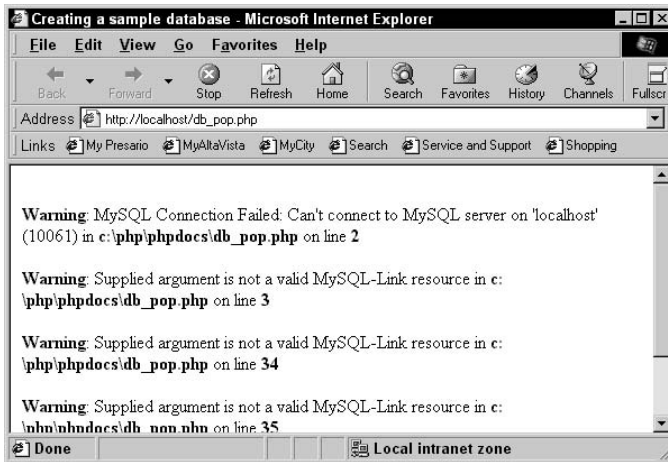
An undefined function fatal error
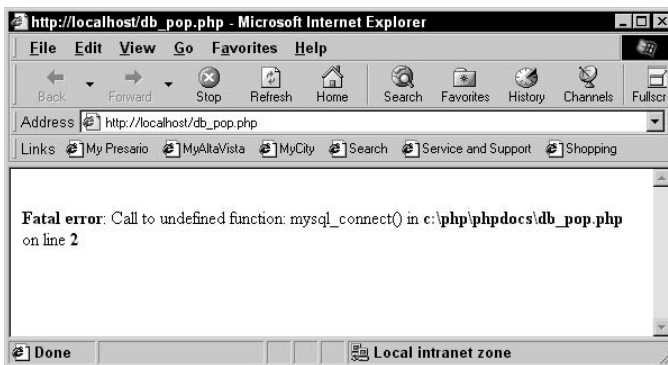
This means PHP doesn't know about MySQL at all.

Of the two, the fatal error is much more straightforward to fix. Clearly, if you're running into an undefined function that is supposed to be in the PHP function set, you can be pretty sure that you simply forgot to build that module into your installation. So on the Unix side, you will need to recompile the code with the `--with-mysql` option. On the Windows side, MySQL should be

precompiled into the binary for you and immediately available. In the case of any other supported database or a version of PHP older than 4.1, you merely need to uncomment the `extension=php_[database].dll` line in your `php.ini` file to be ready to go, unless you put your MySQL executable in a very, very strange place (which you shouldn't do unless you're prepared to handle the consequences, including fatal errors).

The innocuous-looking No Connection error is actually a little harder to diagnose because there are several possible causes. They fall into two main categories:

- The MySQL daemon isn't running.
- The MySQL socket isn't where PHP is looking for it.

It's easy to check whether `mysqld` is running, so you may as well do that first. Just use whatever method you prefer to check running processes. On Windows, this means it's time for the old Ctrl+Alt+Delete action to bring up the Task Manager. On Linux you can check the system processes by means of `ps`.

If `mysqld` is not running, perhaps you have merely forgotten to (re)start it. (Don't laugh. It happens.) If it's been running continuously for 143 days before suddenly quitting in the middle of an operation, your problem is beyond the scope of this book. We can only direct you to the MySQL web site (at `www.mysql.com`) with our deepest sympathies and most fervent hopes that you've maintained a good backup schedule.

The socket problem usually arises the first time you fire up MySQL on a new server. It's rather uncommon for this problem to occur in a long-running site, although it does happen. For instance, we recently had a web host move our MySQL daemon to another server on short notice, at which point all our scripts that used the hostname `localhost` immediately crashed.

The solution to your database connection problems is generally to be found in the `php.ini` file. There's a section of MySQL variables that you must carefully check against whatever hostname, port, and socket you're specifying in your PHP scripts. You want to ensure that you're not inadvertently directing PHP to look for MySQL on an odd port or at the wrong default host. On Linux, you can also check the `/etc/services` file for a different socket address, and the `/etc/hosts` file for an unexpected server alias. In general, you should leave these variables open unless you have a specific reason to set them.

Problems with Privileges

Error messages caused by privilege problems look a lot like the connection errors described previously. You will see a No Connection error that looks like Figure 19-3.

The key differentiator is that little piece about the user and password.

CAUTION **Because of the security issues caused by these failure messages, which include the database username and host and whether you're using a password or not, it's best to use silent mode on a production site. You do this by putting the character @ in front of the functions** `mysql_connect` **and** `mysql_select_db` **or by setting** `display_errors` **to** `off` **in the** `php.ini` **file.**

These errors are many in number but fall into pretty clear major types:

- Employing a database username that lacks the necessary permissions for the task.
- Mistyping usernames/passwords.
- Failing to use a necessary password.
- Trying to use a nonexistent password.
- Trying to use your system's username/password instead of the MySQL username/password.
- Logging in from a location or client that the MySQL database does not allow for a particular user.
- PHP's being unable to open the database-password `include` file because of incorrect file permissions. (It must be a file readable by your web.)
- The database root user's having deliberately changed permissions on you.

FIGURE 19-3

Privilege problems
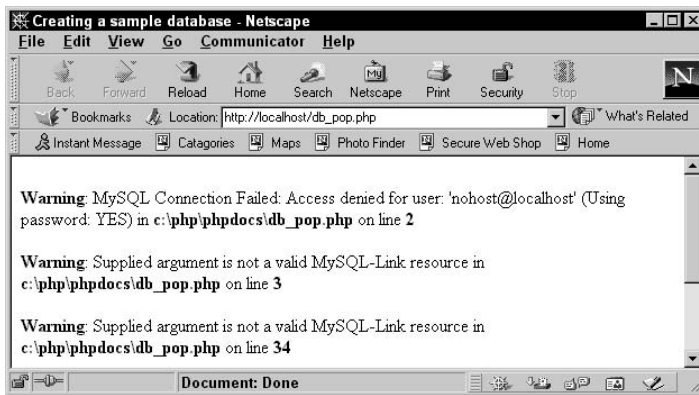
These are not structural problems but usually just simple slips of memory that result in miscues or misrecollections. They are very common. We aren't too proud to confess that we've fallen victim to all of them — and not just once but over and over. They should be trivial to fix in the vast majority of situations. If you are confident your username and password combination is correct, you try using MySQL's `FLUSH PRIVILEGES` command to ensure that the most current changes are loaded.

Unescaped Quotes

Quotes can cause many small but annoying buglets between PHP and MySQL. The crux of the issue is that PHP evaluates within double quotation marks and largely ignores single quotation marks, whereas MySQL evaluates within single quotation marks and largely ignores double quotation marks. This can lead to situations where you have to think hard about the purpose of each quotation mark. An example is:

```
mysql_query("INSERT INTO book (ID, title, year, ISBN)
            VALUES(NULL, '$title', '$year', '$ISBN')");
```

In most of PHP, variables within single quotation marks are not expanded, whereas variables in double quotation marks or unquoted variables are — so this query looks a bit strange. But if you think about it, the statement is valid in both languages. The single quotation marks exist within double quotation marks, so PHP takes them as literal characters, and the variables are actually within double quotation marks, so PHP replaces them with their values. You can think of the division of labor this way: In a database query, PHP does its thing on the stuff between double quotation marks (treating single quotation marks literally), and MySQL later deals with the stuff left over within single quotation marks.

Obviously, you'll need to exercise some care when writing these statements. This is one of the reasons why it's preferable to break up your MySQL queries into two parts, a query string and a `mysql_query()` function, like this:

```
$query = "INSERT INTO book (ID, title, year, ISBN)
          VALUES(NULL, '$title', '$year', '$ISBN')";
$result = mysql_query($query);
```

This style also eliminates the double parentheses that account for common PHP errors.

Even greater issues arise with strings that use single quotation marks and double quotation marks within the text. Remember that apostrophes and single quotation marks are the same thing for PHP and MySQL — they have no smart-quoting feature (not that most smart quotation marks are all that smart anyway). So this insertion query will break as follows if any of your lastname entries ever has an apostrophe in it (for example, O'Hara, D'Souza, and M'Naughten):

```
$query = "INSERT INTO employee (ID, lastname, firstname)
          VALUES(NULL, '$lastname', '$firstname')";
$result = mysql_query($query);
```

Other very common problems are caused by names of businesses with apostrophes in them, such as Rosalita's Bar and Grill or Yoshi's Hair Salon, and by any string that might have a contraction or possessive in it (such as can't, what's, or Mike's).

The parallel issue on the PHP side is a string with a double quotation mark in it. This construction will definitely not work as intended:

```
$string = "He said, "I'm not angry," but I knew he was.";
```

299

```
$statement = mysql_query("INSERT INTO diary (ID, entry)
                         VALUES(NULL, '$string')";
```

CAUTION   **In very long text entries, a quotation mark problem may present as a partial string being inserted, or it may appear as a complete failure, or it may seem as though only short entries are being accepted while longer entries fail.**

If you're using an HTML form with values, and only the first word of your string is being inserted, the problem is likely to be that you forgot to quote the form value properly. In other words, your form field says `<INPUT TYPE="text" VALUE=quoted string>` rather than `<INPUT TYPE="text" VALUE="quoted string">`.

The following list reviews the three ways of dealing with quoting issues:

- In cases where the string is directly stated within the code, you can escape the necessary characters with a backslash.

```
$query = "INSERT INTO employee (ID, lastname, firstname)
         VALUES(NULL, 'O\'Donnell', 'Sean')";
```

- In cases where the string is represented by a variable, you can use `addslashes()`, which will automatically add any necessary backslashes.

```
$string =
mysql_real_escape_string("He said, 'I'm not angry,' but I knew he
was.");
$statement = mysql_query("INSERT INTO diary (ID, entry)
             VALUES(NULL, '$string')");
```

For some murky psychological reason, many PHP users seem exceedingly averse to using `addslashes()` and its partner, `stripslashes()`. People will tie themselves in knots using single quotation marks when they really shouldn't, just so they don't have to escape double quotation marks. This practice is bad style at any time but is especially dangerous when using a database.

You need to add slashes when inserting values into a database; conversely, you'll need to strip out the slashes when pulling strings from a database (unless you have magic quotation marks turned on).

```
$query = "SELECT passphrase FROM userinfo
         WHERE username='$username'";
$result = mysql_query($query);
$query_row = mysql_fetch_array($result);
$passphrase = stripslashes($query_row[0]);
```

If you fail to do this, more and more slashes will be added each time you reenter the data into MySQL! This is an issue that is very frequently encountered with editable Web forms that redisplay values pulled from a database, as shown in Figure 19-4.

However, the preferred solution is to use `mysql_real_escape_string` to escape characters prior to sending them to the database.

FIGURE 19-4

Unstripped slashes in a form

Broken SQL Statements

In addition to quoting problems, there are a number of easy ways to send a *bad* query to the database. That query might be syntactically malformed, have the right syntax but refer to tables that do not exist, or have any of a number of problems that make the database unable to handle it properly. A typical error message is shown in Figure 19-5.

FIGURE 19-5

A bad SQL statement error
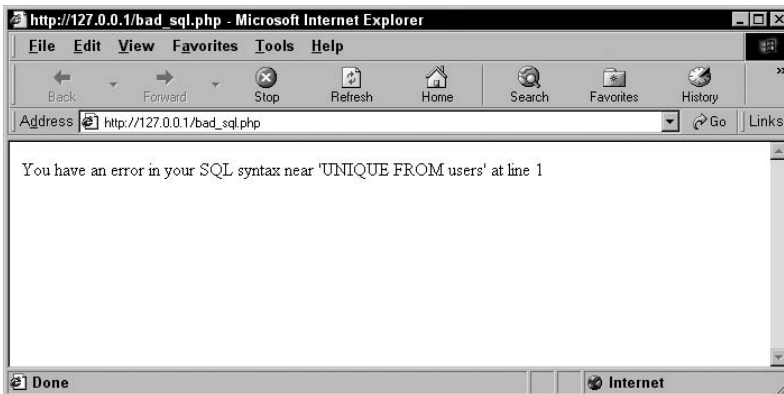
CAUTION   A MySQL error (such as the one shown in Figure 19-5) is different from a connection or link error, which looks something like Figure 19-1. A MySQL error is the error returned from the database when you try to do something that it doesn't like. It is not automatically echoed to the screen; you need to call `mysql_error()` to see any output. A connection error is a message that PHP is sending to you when an expected connection or link is not present. It is automatically echoed to the screen if you're using `display_errors` and must be silenced by being prepended with an `@`.

Older versions of PHP used to automatically `echo` an error statement in these circumstances. Now, if you wish to find out what the problem is, you must manually call `mysql_error()` (as we've done in the preceding example) or `mysql_errno()`. The safest way to capture these errors is to send them to a log file by using `error_log()`.

NOTE   A broken or invalid SQL query is not the same thing as a query that returns no rows. You can write a perfectly fine SQL query like the following:

```
$query = "select ID from cust where name = 'nonexistent'";
```

You send it to your DB and get back a perfectly valid result set, which happens to contain exactly 0 rows. Among other things, this means that error trapping that catches query failures will not help you detect the case of zero rows. For MySQLers, a helpful function is `mysql_num_rows()`, which is called on the query result ID and returns an integer.

Exactly how a bad SQL problem will present itself in your browser depends on your PHP version, your database version, your error settings, and how much error-checking code you have incorporated in your script. Just as with other kinds of malignancy, early detection of a failed query is key.

Your new best friend for making MySQL queries looks like this:

```
$result = mysql_query($query) or error_log(mysql_error());
```

Because `mysql_query()` will return a false value if it fails, the `error_log()` portion will be executed only if a failure occurs. The low operator precedence of the `or` operator ensures that the `error_log()` call also plays no role in the assignment statement — if the assignment succeeds, it is as if the `error()` portion did not exist. Failure leads to the script exiting just as soon as it has printed the most informative error message that the MySQL designers could concoct. If your particular database lacks such an error variable in PHP, you might want to simply call `error_log($query)`. Often, the problem is obvious after you see the query that is actually being sent.

If you have not incorporated error checking into your query calls, you will get the first bad news when you try to use the query result ID in subsequent database code. The typical pattern is:

```
$my_result = mysql_query($bad_query);
$row = mysql_fetch_row($my_result); // error shows up here
```

The typical error message for MySQL is `0 is not a mysql result identifier in [some row]`. This is because, rather than detecting the `0` value that `mysql_query()` returns when it fails, you have tried to use that value as if it were a valid identifier for a result set.

> **TIP** **Although a bad query is by far the most common way of producing the** `0 is not a valid result identifier` **message, it is not the only way. You would also get that message if you misspelled the name of the result identifier variable (and it was, therefore, unbound) or if the query statement had never actually executed (with the same result). Again, it is much easier to distinguish these problems if you trap the errors early on.**

If you suspect a broken query is causing your script to fail, liberal use of `print` and `var_dump` to output the actual SQL being executed is almost always helpful. I've found that seeing the SQL being executed will show a blatant error with quoting or an even subtler error. Taking the SQL and running it manually through the MySQL CLI can also help to reveal errors. I will say more on this, later.

Misspelled names

The sad truth is that for every bug that plumbs the depths of programming esoterica, there are a gazillion cheap mistakes that seem obvious once you've discovered them. The former may break your brain, but afterward you feel a certain exhilaration at testing your skills against a really hard nut. The latter just leave you feeling empty and regretful at the time you wasted on something so trivial.

So let us start with the single most common error: simple misspelling of table, column, and value names. It doesn't help that PHP and MySQL are both case-sensitive in Linux environments (but not on Windows). No force on earth can prevent you from using the wrong case once in a while, and the error messages will be uninformative at best. What can we say? Remember that even the most experienced programmers do it, too.

Comma faults

Remember to put the comma *outside* the single quotation marks within a SQL statement. This will not work:

```
$query = "UPDATE book SET title='$title,' subtitle='$subtitle,'
ISBN='$ISBN'";
```

But this will:

```
$query = "UPDATE book SET title='$title', subtitle='$subtitle',
ISBN='$ISBN'";
```

Think of the single quotation marks as part of the variable itself rather than following common American typographical practice, which puts a comma inside the ending quotation mark.

Unquoted string arguments

Any values that should be treated by the database as string data types typically need to be single-quoted within a SQL statement. For example, this query has the correct syntax:

```
$query = "SELECT * FROM author WHERE firstname = 'Daniel'";
```

By contrast, if we make a `mysql_query()` call using the following query, we should expect an error:

```
$query = "SELECT * FROM author WHERE firstname = Daniel";
```

The actual error returned by the database may be deceptive, though — quite likely the complaint will be about an unknown column named `'Daniel'`. This is because unquoted strings are assumed to name columns, as in:

```
$query = "select * from author where firstname = lastname";
```

This would be a perfectly acceptable way to search our database for Humbert Humbert and Lisa Lisa, but it won't work for people with more ordinary names.

Unbound variables

One of the sneakier ways to break a SQL statement is to interpolate an unbound variable into the middle of it.

When it works, the automatic splicing of variables into double-quoted strings is a perfect match for a SQL-based dialog with your database. Your code can determine values, for example, that are used to restrict the scope of a query made to the DB, as in this snippet:

```
$customerID = find_customer_id(); //returns int
$result_id = mysql_query("SELECT * FROM customers
                    WHERE ID = $customer_ID"); //BUG
$row = mysql_fetch_row($result_id);  //CRASH
```

Because this code makes no attempt to trap query errors, you will again see a complaint about the fact that `0 is not a valid MySQL result identifier`. It's possible (for us anyway) to stare at code like this for quite a while without seeing anything wrong (although the good PHP coders who habitually crank error reporting up to `E_ALL` will be rewarded with the cause of the error in a warning message). The problem, of course, is that we assigned one variable (`$customerID`) and then embedded a different one (`$customer_ID`) in our SQL statement. The latter variable is unbound and so behaves like an empty string when interpreted by the double-quote parsing. The result is that the database sees the following query, which is not valid SQL:

```
SELECT * FROM customers WHERE ID =
```

This kind of problem is one reason why it is often a good idea to construct your query and assign it to a variable in a separate statement, like this:

```
$my_query = "SELECT * FROM customers WHERE ID = $customer_ID";
```

Then make a distinct subsequent call to `mysql_query($my_query)`. If you do this, it is very easy to add printing or logging statements that show you the actual query you are sending.

Too Little Data, Too Much Data

Finally, you may find that your PHP/database script is working apparently without error but is displaying no data from the database or far more than you expected. As a vague and general rule, if your query function is returning successfully (and your code checks that), your suspicions might rightly turn to the SQL itself. Recheck the logic, particularly of WHERE clauses. It is easy, for example, to write a query like:

```
"SELECT * FROM families WHERE kidcount = 1 AND kidcount = 2";
```

In this query, you are really intending an or rather than an and, with the result that zero rows will be returned regardless of the contents of your database.

If your script is iterating through database rows and displaying them and you find that you have far, far too many of those rows, the problem is very often a SQL join that has too few restrictions. As a general rule, the number of restrictions in a WHERE clause should not be fewer than the number of tables joined minus one. For example, the following query has three tables but only one joining restriction:

```
"SELECT book.title FROM book, author, country
WHERE author.countryID = country.ID"
```

It is likely to return every possible book/author pair, without reference to whether the author wrote the book, which is probably not what was intended.

Specific SQL Functions

A few specific functions seem to cause a higher than normal number of problems, especially in the learning phase. These functions can send even the experienced PHP developer running to the online manual to check the arguments and returned data types time and time again.

mysql_affected_rows() versus mysql_num_rows()

Both of these functions tell you how many rows of data your last SQL statement touched. However, mysql_num_rows() works only on SELECT statements, while mysql_affected_rows() works only on INSERT, UPDATE, and DELETE statements. The way to think about it is that SELECTs do not affect (meaning *change*) any data that exists in the database.

Furthermore, mysql_affected_rows() takes an optional link identifier as the argument, whereas mysql_num_rows() takes a nonoptional result resource. This means that you can only get a valid result from mysql_affected_rows() until the moment you call another INSERT, UPDATE, or DELETE. In contrast, if you use different variable names for your result resources, you can use mysql_num_rows() anytime in the script. This code will help clarify the differences:

```
$link_id = mysql_connect($host, $user, $pass);
```

```
mysql_select_db($database, $link_id);

$query = "INSERT INTO mytable VALUES(NULL, '$myval')";
$result_resource = mysql_query($query);
$test_insert = mysql_affected_rows();
// This should work and return 1

$query1 = "SELECT * FROM mytable";
$result_resource1 = mysql_query($query1);
$test_select = mysql_num_rows($result_resource1);

$query2 = "DELETE FROM mytable";
$result_resource2 = mysql_query($query2);
$test_select2 = mysql_num_rows($result_resource2);
//Will not work
$test_delete = mysql_affected_rows();
//This will return the number of rows in the table; at this
// point you can no longer get the old result of 1
$test_select_again = mysql_num_rows($result_resource1);
//Should be the same as $test_select
```

mysql_result()

This function, which returns one value at a time from the database, is now used rather rarely. Unlike `mysql_fetch_row()` and `mysql_fetch_array()`, with `mysql_result` you need to specify the row and field of the value you're fetching as well as the result resource. Thus, you cannot do this:

```
// This won't work
while (mysql_result($result_resource)) {
  // Some loop
}

// This will
$firstname = mysql_result($result_resource1, 0, 'firstname');
```

You should really use this function only when you know you'll be fetching one or two pieces of data (a user's first name, for instance). Otherwise, the others are much faster.

OCI_Fetch()

When users of MySQL or SQL Server switch over to Oracle, they often have trouble with the OCI fetching functions — particularly this one. Unlike most other database row-fetching functions, you don't immediately access the result of `oci_fetch()` via `echo` or some other PHP function. This function fetches the result of a SQL statement into a result buffer — where it can be accessed via `OCIResult()`.

```
$query = "SELECT * from mytable";
$stmt = oci_parse($conn, $query);
$exec_result = oci_execute($stmt, OCI_DEFAULT);
```

```
$row2buffer = oci_fetch($stmt);
$myval = oci_result($stmt, "MYCOLUMN");
echo $myval;
```

This function should probably be thought of as analogous to `mysql_result()` rather than `mysql_fetch_row()`, or at best occupying a middle ground between the two. Similarly, it should only be used when you are sure you will be fetching very small data sets. Otherwise, use `oci_fetch_array()` which returns an array.

Debugging and Sanity Checking

If you are nearing your wit's end in trying to debug query-related errors and misbehavior, you may find it extremely useful to compare the results of your PHP-embedded queries with the same queries made directly to the database. If your technical setup permits actually running a SQL client directly (for example, the `mysql` or Oracle command-line clients), as well as cross-program cutting and pasting, try this two-step process:

1. Insert a debugging statement in your PHP script that prints the query itself immediately before it is actually used in a DB query call (for example, `echo $query`).

2. Directly paste that query from your browser output (or the HTML source) into your SQL client.

CAUTION   Obviously, this advice applies only to code under development, not to code you are running in production. It might be okay to echo errors to the browser while you're developing something for the first time, but when it's ready to go into production, you should make sure all your `echo()` statements are replaced with `error_log()` functions.

If the query looks reasonable to you, but it breaks both in the SQL program and in PHP, then there is some syntax or naming error in that SQL statement itself that you are missing, and your PHP code is not to blame (unless, of course, your code constructed that query in the first place). Similarly, with a dearth or overabundance of rows — if the behavior is the same in both places — the query is to blame. If, on the other hand, the behavior in the SQL interpreter looks like what you wanted, then the query is fine, and your suspicion should turn to your PHP code that actually sends that query and processes the results.

One final and general tip is to study any error messages very carefully, paying attention to phrases like `link identifier` and `result resource identifier`. In MySQL, the former means an identifier of a database *connection*, and the latter identifies the set of rows returned by a particular *query*. It is easy to confuse the two, as in the following code:

```
$my_connection = mysql_connect('localhost', $myname, $mypass);
mysql_select_db('MyDB');
$result = mysql_query($my_query, $my_connection);
while ($row = mysql_fetch_row($my_connection)) {
  // LOOP
}
```

This code will probably yield an error that contains the words `not a valid result identifier`. The problem is that we are using the connection ID where the result ID should be. The resulting error message is justified yet opaque.

Summary

PHP/database bugs are often not very deep or subtle but can still be difficult to diagnose. In general, the earlier in a script you can detect trouble, the easier the diagnosis will be. Especially when you are debugging, every statement that interacts with the database should have an associated `error_log()` clause, containing an informative error message.

By far, the most common cause of database-connection problems is incorrect arguments to the connection function (hostname, username, password). The most common causes of failed queries are quote faults, unbound variables, and simple misspellings.

If you have repeated failures with database queries that seem like they should be working, have your code print out the query that it is sending to the DB; if possible, try making that very query to the database directly. If the problem persists when PHP is out of the loop, you can safely restrict your attention to database design and your understanding of SQL queries.

Part III

More PHP

Chapter 20

Introducing Object-Oriented PHP

There are many possible audiences for this chapter, including people who know basic PHP but nothing about object-oriented programming (OOP), and people who know all about OOP and nothing about PHP. As usual, we aim to please everyone all at once, but be warned that you may want to pick and choose from the sections.

We start with a quick and very general introduction to object-oriented programming for those who are completely unfamiliar with it. If you are already comfortable with OOP from another language, please skip this section — it will not enlighten you (and might well enrage you). The section "PHP Constructs for OOP" gets into the meat of the basic syntax and behavior of PHP objects. Later in the chapter, we delve into more extended examples and cover some of the more obscure issues and gotchas around objects in PHP. Along the way, we offer a couple of sidebar meta-discussions, about the merits of object-oriented PHP and the extent to which PHP should be considered to be OOP.

NOTE   **In general in this chapter, we discuss OOP programming constructs as they are implemented in PHP5, which uses the new and significantly improved Zend Engine 2 as its parser.**

What Is Object-Oriented Programming?

So what is object-oriented programming (OOP) all about anyway? OOP turns out to be a very simple idea, which (when taken seriously and built into the structure of programming languages) leads to all sorts of more complicated elaborations.

The simple idea

The simple idea is this: Rather than creating data structures on the one hand and code on the other, suppose that we reorganize everything so that associated pieces of code and data are bundled together?

The procedural approach

For example, imagine a conventional procedural (non-object-oriented) program for manipulating personal calendars, with the capability to display, update, and edit calendars. Somewhere in the code for such a program, we would find the actual data definitions for representing someone's appointments for a particular month; somewhere else we would find code that did the right things to manipulate that data. Typically, the only connection between the data type definitions and the manipulation code is that a clever programmer has made sure that they get matched up appropriately.

Now imagine combining our calendar program with a recipe program (say that we want to plan our meals in detail for an entire year). Again, there will be data structures somewhere that represent the contents of the calendar, and other data structures that represent the contents of the recipes. The data structures will use the basic data types of the programming language; for all we know, the top-level type of a calendar might be an array, and the top-level type of a recipe might also be an array. Somewhere else in the program there is code for digging into the data structures that represent calendars and recipes and doing the right things with them. What is the connection between the data structures and the code? Only that a careful programmer has made sure that the arrays that represent calendars and the arrays that represent recipes get fed to the appropriate manipulation code. (Otherwise, we might find ourselves trying to schedule an appointment in Beef Stroganoff rather than in March 2006.)

If we think of procedural code as outlined like a book, the outline for the code we're talking about might look like:

- Data definitions
  - Data definitions for calendars
  - Data definitions for recipes
- Data manipulation code
  - Code for calendars
  - Code for recipes

The object-oriented approach

The most basic version of OOP reorganizes the procedural approach by grouping associated pieces of code and data together into conceptual units. This means that we replace the outline in the preceding subsection with:

- Calendars
  - Data definitions
  - Manipulation code
- Recipes
  - Data definitions
  - Manipulation code

This organizational inversion is the heart of object-oriented programming.

But so what (we can hear you say)? If we're just talking about a way to organize code, we could do that without any special terminology or programming languages. In normal procedural code, we can organize function definitions and data type definitions in any order we want to. For example, we could put all the data type definition code into one directory and all the manipulation code into another (a procedural organization), or we could put all the calendar code into one directory and all the recipe code into another (an object-oriented organization).

Object-oriented programming begins to be interestingly different from procedural programming, however, once the programming language itself is set up to make it easy to organize things in an object-oriented way. (See the sidebar "Do Web-Scripting Languages Really Need OOP?" for a discussion of how useful this organization is in languages like PHP.) The most basic form this takes is that data objects can be built out of local functions as well as local data. For example, as we build a data structure that represents a calendar, we can include the data members that are needed (structures to represent days, months, years, appointments) but also the functions that will be needed (`new_appointment()`, `calendar_display()`, and so forth). These functions are (in some sense) stored locally in the object definition itself. A calendar doesn't have an ingredient list, and a recipe doesn't have 31 days; similarly, a calendar object doesn't have a `print_ingredients()` function, and a recipe doesn't have a `new_appointment()` function. Finally, of course, the data members in an object may themselves be objects of a different type.

Bundling code and data together into units is the basic idea, and OOP languages always offer some support for this kind of bundling. However, most OOP languages take things further and offer one or more of the following elaborations that give OOP even more leverage. (See the sidebar "How OO Is PHP?" for a discussion of the extent to which PHP itself has these features.)

Elaboration: objects as data types

In addition to allowing us to store functions in our data, a good OO programming language lets us define these combinations as genuinely new data types that the language supports like any other type.

Do Web-Scripting Languages Really Need OOP?

The object-oriented revolution has not been without controversy. Although many programmers embraced OOP quickly, others preferred the procedural approach they were used to and wondered aloud if the extra machinery needed to support OOP wasn't more trouble than it was worth. Still, there's no doubt that the revolution has largely succeeded. Most of the popular programming languages in use today are either fully object oriented or have object-oriented extensions. Also, at least some of the promises about improved productivity and increased code reuse seem to have been realized, as design methodologies like the Unified Modeling Language (UML) and patterns gain greater influence, and as people get more used to subclassing as a standard way to reuse and extend vendor-supplied libraries.

We feel that the benefits of OOP for "major" (that is, compiled) programming languages like Java and C++ are clear. On the other hand, we feel that the benefits of OOP for scripting languages (like Perl and PHP) are less obvious and are most debatable in the case of web-scripting (PHP).

How is web scripting different from other kinds of programming tasks? The most obvious difference is simply that web scripts typically execute quickly and then go away. In other programming situations, you may have RAM-resident objects that live for hours or days and undergo complex evolutions of state that affect their behavior. A typical web script, on the other hand, might execute for half a second, as it serves up a particular page, and then dies happy. You may knit these scripts together to provide a more extended user experience (using databases, sessions, cookies), but often such efforts are all about making the experience outlive any PHP objects that may be created. More generally, scripting languages like PHP and Perl typically have a less thoroughgoing implementation of OOP than languages like Java, C++, and Smalltalk, and the limitations of implementation make these OOP extensions less attractive. (For more detail, see the sidebar "How OO is PHP?" later in this chapter.)

This is not to say that there aren't still benefits of OOP in PHP. In addition to the conceptual benefits that may result from structuring code in an object-centered way, there are two good reasons to use PHP objects: 1) It's a good way to distribute third-party code for reuse; 2) Many programmers who are used to OO syntax from other languages won't feel comfortable unless they can use the same idioms in PHP.

But our main point is that use of PHP constructs for OOP is a very "tradeoffy" and pragmatic decision, which we have often seen made more on the basis of religion or fashion. If you are comfy with OO, this kind of syntax is there for you, and if you work in a group that has decided to write in that style, you may want to let the majority rule. If you decide not to go OO, however, be strong — we urge you not to be swayed by the moral-superiority arguments you may hear from people who disdain your five-line procedural script in favor of their ten-line OO script that does exactly the same task.

After such a type is defined, we can create as many such objects as we like, just as we can create as many integers as we like given the integer type. In object-oriented terminology, the term *class* is used to refer to the general type definition, which specifies the data members and member functions that each instance of that class should have.

The term *object* (or *instance*) refers to any individual instance of the type. For example, after we define a class called `Calendar` (which specifies the different kinds of data and functions that every self-respecting calendar should have), we can make any number of `Calendar` objects (which might be associated with individual people).

Elaboration: Inheritance

After we've written a program that uses the class `Calendar`, we might want to make a more specific version of the program for a particular purpose. What we would really like to do is copy most of the code from the `Calendar` class but change it in just a few places, so that it prints differently or has a culturally appropriate set of holidays or allows us to schedule appointments to the second rather than to the hour.

This desire is common enough that OOP offers a mechanism to support it called *inheritance*. The basic idea is that you can define a class in terms of another class and then specify only the things that you want to be different in your own class. If you view the original class as the parent, the default is that both function definitions and data definitions are *inherited* by your child class unless you specify otherwise. This turns out to be a powerful technique for reusing class definitions. (As you will see in the "Basic PHP Constructs for OOP" section, OOP in PHP supports inheritance.)

Elaboration: Encapsulation

Part of the point of segregating both data and functions into objects is to reduce the complexity of programming by reducing unnecessary interactions. There is no reason why calendars should have to know about the internals of cooking recipes, or vice versa. So some OOP languages actually enforce information barriers between objects — after the programmer has defined which parts of recipes and calendars are purely internal and private to those classes, the language actually forbids code that is external to an object from messing with an object's internal workings. This kind of information-hiding is called *encapsulation*, and although this sounds restrictive, it can be a good source of clarity. In particular, if the programmer who designed a particular class knows that some parts of its workings have been designed to be private in this sense, the programmer also knows that those parts can be redesigned without checking with everyone who might be using that class's code. Support for encapsulation existed for the first time in PHP5, which incorporates Zend Engine 2. You'll see how to use encapsulation later in this chapter.

Elaboration: Constructors and destructors

After you have defined a class, you can make as many instances of it as you like. Each time you create such an instance, your favorite OOP language allocates memory to store the instance in, and gives you some way to refer to that instance later in the program. There are frequently a number of initialization steps you want to take every time you make an object of that class. *Constructor* functions offer a way to build that set of steps into the class definition. The standard way to create a new instance is to call a constructor function (which usually has the same name as the class and which you can customize to do all the necessary initialization).

Destructors are the opposite of constructors and specify all the cleanup actions that should happen when an object is dispensed with.

PHP has offered constructor functions since version 4.2 (which makes sense, because you can't have object orientation without having constructors). The language acquired explicitly definable and callable destructors only in PHP5 (destruction of classes was handled only in an automatic way before then). Again, these functions are covered later in this chapter.

315

Terminology

There are some standard terms in OOP parlance for all the concepts we have talked about thus far, and we will be using them for the rest of the chapter. (Several of these terms have alternate names, which we include in parentheses.)

- **Class:** This is a programmer-defined data type, which includes local functions as well as local data. You can think of a class as a template (or mold, or form) for making many instances of the same kind (or class) of object.

- **Object:** (Also known as *object instance*, or *instance*.) An individual instance of the data structure defined by a class. You define a class once and then make many objects that belong to it.

- **Member variable:** (Also known as *property*, *attribute*, or *instance variable*.) One of the component pieces of data in a class definition.

- **Member function:** (Also known as *method*.) A member that happens to be a function.

- **Inheritance:** The process of defining a class in terms of another class. The new (child) class has all the member data and member function definitions from the old (parent) class by default but may define new members or "override" parent functions and give them new definitions. We say that class A *inherits from* class B if class A is defined in terms of class B in this way.

- **Parent class** (or *superclass* or *base class*): A class that is inherited from by another class.

- **Child class** (or *subclass* or *derived class*): A class that inherits from another class.

How OO is PHP?

How "object-oriented" is PHP? Your answer to that question probably depends on your particular litmus tests for object-orientedness. In this sidebar, we offer a whirlwind tour of features that typically show up in OOP languages and briefly discuss the extent to which PHP supports them. Some of these issues are explored more broadly in the section "Advanced OOP Features," later in this chapter. (Note: This sidebar is really only of interest to developers who are coming to PHP from a different OO language; everyone else may want to skip this game of buzzword bingo.)

Single inheritance

PHP allows a class definition to inherit from another class, using the `extends` clause. Both member variables and member functions are inherited.

Multiple inheritance

PHP offers no support for multiple inheritance as in Java. Each class inherits from, at most, one parent class (though a class may implement many interfaces).

Constructors

Every class can have one constructor function, which in PHP is called `__construct()`. Note that there are two underscore characters at the front of that function name. Because prior to PHP5 (under Zend Engine 1), a class's constructor function had the same name as the class, PHP still allows (but discourages) that strategy for purposes of backward compatibility. Constructors of parent classes are not automatically called but must be invoked explicitly.

Destructors

PHP supports explicit destructor functions as of version 5. The destructor function of a class is always called `__destruct()`.

Encapsulation/access control

PHP supports public, private, and protected variables and member functions as of version 5.

Polymorphism/overloading

PHP supports polymorphism in the sense of allowing instance of subclasses to be used in place of parent instances. The correct member function will be dispatched to at runtime. There is no support for method overloading, where dispatch happens based on the method's signature — each class only has one member function of a given name. However, PHP's weak typing and support for variable numbers of arguments makes workarounds possible. See the section "Simulating polymorphism" later in this chapter (in the section "Advanced OOP Features").

Early versus late binding

Two equally good answers are: (1) The question doesn't arise, because of PHP being loosely typed, and (2) All binding is late. In PHP, values are typed but variables are not, so there is no question about what method to call when the variable is of a different type than the value.

Static (or class) functions

PHP offers static member variables and static methods as of version 5. It is also possible to call member functions via the `Classname::function()` syntax.

Introspection

PHP offers a wide variety of functions here, including the capability to recover class names, member function names, and member variable names from an instance. (See the section "Introspection Functions," later in this chapter.)

Namespaces

PHP6 offers namespaces; these define the area in which an identifier, such as a variable, is unique. For example, a variable named `$foo` inside of a private namespace is different from a global variable `$foo`.

Basic PHP Constructs for OOP

In this section, we cover the basic PHP syntax for OOP from the ground up, with some simple examples.

Defining classes

The general form for defining a new class in PHP is as follows:

```
class myclass extends myparent {
  public $var1;
  public $var2 = "constant string";
  public function myfunc ($arg1, $arg2) {
    [..]
  }
  [..]
}
```

The form of the syntax is as described, in order, in the following list:

- The special form `class`, followed by the name of the class that you want to define.
- An optional extension clause, consisting of the word `extends` and then the name of the class that should be inherited from.
- A set of braces enclosing any number of variable declarations and function definitions. Variable declarations start with the special form `public`, `private`, or `protected`, which is followed by a conventional $ variable name; they may also have an initial assignment to a constant value. Function definitions look much like standalone PHP functions but are local to the class.

As an example, consider the simple class definition in Listing 20-1, which prints out a box of text in HTML.

LISTING 20-1

TextBox.php

```
class TextBoxSimple {
  public $body_text = "my text";
  function display() {
    print("<TABLE BORDER=1><TR><TD>$this->body_text");
    print("</TD></TR></TABLE>");
  }
}
```

This is an extremely simple class definition. It has no parent (and, therefore, no `extends` clause). It has a single member variable (the variable `$body_text`) and a single member function (the function `display()`). The display function simply prints out the text variable, wrapped up in an HTML table definition.

Accessing member variables

In general, the way to refer to a member variable from an object is to follow a variable containing the object with `->` and then the name of the member. So if we had a variable `$box` containing an object instance of the class `TextBox`, we could retrieve its `body_text` variable with an expression like:

```
$text_of_box = $box->body_text;
```

However, when we are writing code within a member function, we haven't yet created the object instance, and so we have no variable like `$box` to refer to. The answer is the magic variable `$this`, which (when used inside a member function of a class) refers to the object instance itself. Note that this is how the `display()` function in Listing 20-1 retrieves the text it displays (`$this->body_text`).

This syntax can be a little counterintuitive. You might think that we could simply refer to `$body_text` in functions within our `TextBox` class because we have declared it in the class definition, but in fact the only way to get to members from within a member function definition is via `$this`. Notice also that the syntax for this access does not put a `$` before the member variable name itself, only the `$this` variable.

Creating instances

After we have a class definition, the default way to make an instance of that class is by using the `new` operator. If we have already defined the class `TextBox` as in Listing 20-1, we can make an instance of it, and then use it, like this:

```
$box = new TextBoxSimple;
$box->display();
```

The result of evaluating this code will be to print an HTML fragment containing a table definition enclosing the text `my text`. (Not especially useful, but it's a start.)

Constructor functions

One way in which our `TextBox` class is not very useful is that its instances do not contain any data when they are created, except for the static initialization of the variable `$body_text`. The point of such a class would be to display arbitrary pieces of text, not the same message every time. It's true that we could make an instance and then install the right data in the instance's internal variables, like this:

```
$box = new TextBoxSimple;
$box->body_text = "custom text";
$box->display();
```

But that would be cumbersome and error-prone as we build more complex objects.

The correct way to arrange for data to be appropriately initialized is by writing a constructor function — a special function called __construct(), which will be called automatically whenever a new instance is created.

Modifying our previous example to include a constructor function gives us Listing 20-2.

LISTING 20-2

TextBox redefined

```
class TextBox {
  public $body_text = "my text";
  // Constructor function
  public function __construct($text_in) {
    $this->body_text = $text_in;
  }
  function display() {
    print("<TABLE BORDER=1><TR><TD>$this->body_text");
    print("</TD></TR></TABLE>");
  }
}
// creating an instance
$box = new TextBox("custom text");
$box->display();
```

As the preceding code is executed, the output is an HTML table enclosing the text custom text.

> **NOTE** There should be only one constructor function per class definition. Defining more than one such function is syntactically legal, but pointless, as only the definition that occurs last will be in effect. If you'd like to have different constructors to handle different numbers and types of input arguments, see the section "Simulating Polymorphism" later in this chapter.

Inheritance

PHP class definitions can optionally *inherit* from a parent class definition by using the extends clause. The syntax is:

```
class Child extends Parent {
 <definition body>
}
```

The effect of inheritance is that the child class (or subclass or derived class) has the following characteristics:

- Automatically has all the member variable declarations of the parent class (or superclass or base class)

- Automatically has all the same member functions as the parent, which (by default) will work the same way as those functions do in the parent

In addition, the child class can add on any desired variables or functions simply by including them in the class definition in the usual way.

In Listing 20-2, we defined a class called `TextBox`; now we'll define a class called `TextBoxHeader` that extends `TextBox` (see Listing 20-3). `TextBoxHeader` has two member variables: one (`$body_text`) that it receives through inheritance from `TextBox`, and another (`$header_text`) that it defines itself. Like `TextBox`, it has a constructor function and a function called `display`. This function definition overrides the display function in `TextBox`.

LISTING 20-3

TextBoxHeader

```php
class TextBoxHeader extends TextBox
{
  public $header_text;

  // CONSTRUCTOR
  public function __construct($header_text_in,
                    $body_text_in) {
    $this->header_text = $header_text_in;
    $this->body_text = $body_text_in;
  }

  // MAIN DISPLAY FUNCTION
  public function display() {
    $header_html =
      $this->make_header($this->header_text);
    $body_html = $this->make_body($this->body_text);
    print("<TABLE BORDER=1><TR><TD>\n");
    print("$header_html\n");
    print("</TD></TR><TR><TD>\n");
    print("$body_html\n");
    print("</TD></TR></TABLE>\n");
  }

  // HELPER FUNCTIONS
  public function make_header ($text) {
    return($text);
  }
  public function make_body ($text) {
    return($text);
  }
}
```

Overriding functions

Function definitions in child classes *override* definitions with the same name in parent classes. This just means that the overriding definition in the more specific class takes precedence and will be the one actually executed. In the example in Listing 20-3, the TextBoxHeader class defines a function called display(), which means that executing the following code:

```
$text_box_header = new TextBoxHeader("The Header", "The Body");
$text_box_header->display();
```

will result in a call to TextBoxHeader's display() function, not the display() function in TextBox. The resulting HTML output prints a box with a header of The Header and a body of The Body. The more specific display() function takes total responsibility here; there is no call, either explicit or implicit, to the display() function defined in the TextBox class. (Although PHP makes no such implicit calls, it is possible to explicitly call functions that have been defined in a parent class — see "Calling parent functions" in the "Advanced OOP Features" section later in the chapter.)

The flip side of overriding functions, however, is that whenever a subclass does not override a parental definition, the parent's definition will be in effect. Note that the "helper" functions in the definition of TextBoxHeader don't really do anything interesting, and you might wonder why we bothered to separate them out. The answer is that this provides an opportunity for an inheriting class to do something interesting with those functions by selectively overriding them — or not, as they see fit.

PHP5 (as a result of Zend Engine 2) introduced the final keyword. If, in the previous example, the definition of display() in class TextBox had looked like this:

```
final function display() {
    print("<TABLE BORDER=1><TR><TD>$this->body_text");
    print("</TD></TR></TABLE>");
}
```

then the method could not have been overridden by a definition in TextBoxHeader.

It is possible to declare whole classes final and individual methods, but not individual properties.

Chained subclassing

PHP does not support multiple inheritance but does support *chained subclassing*. This is a fancy way of saying that, although each class can have only a single parent, classes can still have a long and distinguished ancestry (grandparents, great-grandparents, and so on). Also, there's no restriction on family size; each parent class can have an arbitrary number of children.

As example, see Listing 20-4, where our definition of TextBoxBoldHeader inherits from TextBoxHeader, which in turn inherits from TextBox.

LISTING 20-4

TextBoxBoldHeader

```
class TextBoxBoldHeader extends TextBoxHeader {

  // CONSTRUCTOR
  public function __construct($header_text_in,
                  $body_text_in) {
    $this->header_text = $header_text_in;
    $this->body_text = $body_text_in;
  }

  // HELPER FUNCTIONS
  // make_header overrides parent
  public function make_header ($text) {
    return("<B>$text</B>");
  }
}
```

This definition of `TextBoxBoldHeader` is minimal; it defines no new member variables and defines only one function besides its constructor. That new function (`make_header()`) overrides the definition in its parent. Now what happens when we actually use this definition in the usual way?

```
$text_box_bold_header =
   new TextBoxBoldHeader("The Header", "The Body");
$text_box_bold_header->display();
```

It's worth looking in a bit of detail to see exactly what happens when we make these two function calls.

First, when we call the constructor (`TextBoxBoldHeader()`), the constructor sets variables that were defined in the grandparent (`TextBox`) and the parent (`TextBoxHeader`), respectively, and returns a new instance of `TextBoxBoldHeader`.

Second, when we call `$text_box_bold_header->display()`, the call sequence is:

1. No `display()` function is found in `TextBoxBoldHeader`, so the version from `TextBoxHeader` is called.

2. The first function call in that version of `display()` is to `$this->make_header()`. Remember that `$this` refers to the object instance that we started with, which happens to be an instance of `TextBoxBoldHeader`, so PHP looks first of all for a definition from that class. It finds one and uses it to return the header string wrapped up in the HTML bold text construct (`<B></B>`).

3. The second function call is to `$this->make_body()`. This time, though, there is no overriding definition in `TextBoxBoldHeader`, so the version from `TextBoxHeader` is used.

323

The upshot is that, in defining `TextBoxBoldHeader`, we mostly exploited the behavior of the parent class but were able to change its behavior slightly by overriding a single member function.

Modifying and assigning objects

Prior to PHP5, when you assigned an object to a variable or passed it to a function, that object was actually copied, bit for bit, into the variable or function scope. That caused tremendous hassles, and programmers had to be careful to devise clever workarounds for the problems.

The problem was solved with PHP5, which incorporates Zend Engine 2. Zend Engine 2 copies by reference, rather than explicitly. That is, several variables can point to the exact same object and expect changes made via one reference to be reflected in the others.

Scoping issues

Before we move onto the more advanced features of PHP's version of OOP, it's important to nail down issues of scope — that is, which names are meaningful in what way to different parts of our code. It may seem as though the introduction of classes, instances, and member functions have made questions of scope much more complicated. Actually, though, there are only a few basic rules we need to add to make OOP scope sensible within the rest of PHP:

- Names of member variables and member functions are never meaningful to calling code on their own — they must always be reached via the `->` construct (or, as we'll see in the "Advanced OOP Features" section, the `::` construct). This is true both outside the class definition and inside member functions.

- The names visible within member functions are exactly the same as the names visible within global functions — that is, member functions can refer freely to other global functions but can't refer to normal global variables unless those variables have been declared global inside the member function definition.

These rules, together with the usual rules about variable scope in PHP, are respected in the intentionally confusing example in Listing 20-5. What number would you expect that code to print when executed?

LISTING 20-5

Confusing scope

```
$my_global = 3;

public function my_function ($my_input) {
  global $my_global;
  return($my_global * $my_input);
}
```

```
class MyClass {
  protected $my_member;
  function __construct($my_constructor_input) {
    $this->my_member =
      $my_constructor_input;
  }
  public function myMemberFunction ($my_input) {
    global $my_global;
    return($my_global *
           $my_input *
           my_function($this->my_member));
  }
}

$my_instance = new MyClass(4);
print("The answer is: " .
      $my_instance->myMemberFunction(5));
```

The answer is: 180 (or 3 * 5 * (3 * 4)). If any of these numerical variables had been undefined when multiplied, we would have expected the variable to have a default value of 0, making the answer have a value of 0 as well. This would have happened if we had:

■ Left out the global declaration in `my_function()`

■ Left out the global declaration in `myMemberFunction()`

■ Referred to `$my_member` rather than `$this->my_member`

Advanced OOP Features

In the previous section, we presented a minimal subset of PHP's object-oriented constructs that let you use the most basic OOP techniques. In this section, we look at some of the slightly more unusual constructs, techniques, and gotchas that can get you into more trouble. (We defer any discussion of the functions that give meta-information about classes and objects to the section "Introspection Functions," later in this chapter.)

Public, Private, and Protected Members

Unless you specify otherwise, properties and methods of a class are public. That is to say, they may be accessed in three possible situations:

■ From outside the class in which it is declared

■ From within the class in which it is declared

■ From within another class that implements the class in which it is declared

If you wish to limit the accessibility of the members of a class, you should use `private` or `protected`.

Private members

By designating a member private, you limit its accessibility to the class in which it is declared. The private member cannot be referred to from classes that inherit the class in which it is declared and cannot be accessed from outside the class.

Making a member private is straightforward:

```
class MyClass {

private $colorOfSky = "blue";
$nameOfShip = "Java Star";

public function __construct($incomingValue) {
// Statements here run every time an instance of the class
// is created.
}

public function myPublicFunction ($my_input) {
    return("I'm visible!");
}

private function myPrivateFunction ($my_input) {
    global $my_global;
    return($my_global *
            $my_input *
            my_function($this->my_member));
  }

}
```

When that class is inherited by another class (using `extends`), `myPublicFunction()` will be visible, as will `$nameOfShip`. The extending class will not have any awareness of or access to `myPrivateFunction`, because it is declared `private`.

Protected members

A protected property or method is accessible in the class in which it is declared, as well as in classes that extend that class. Protected members are not available outside of those two kinds of classes, however.

Here is a different version of `MyClass`:

```
class MyClass {

protected $colorOfSky = "blue";
```

```
        $nameOfShip = "Java Star";

        public function __construct($incomingValue) {
        // Statements here run every time an instance
        // of the class is created.
        }

        public function myPublicFunction ($my_input) {
            return("I'm visible!");
        }

        protected function myProtectedFunction ($my_input) {
            global $my_global;
            return($my_global *
                    $my_input *
                    my_function($this->my_member));
        }

    }
```

If we had another class that extended `MyClass`, it would be able to see and use `$colorOfSky` and `myProtectedFunction()`, just as if they were public. It would not, however, be possible to call `MyClass::$colorOfSky`. You'll read more about the `::` syntax later in this chapter.

Interfaces

In large object-oriented projects, there is some advantage to be realized in having standard names for methods that do certain work. For example, if many classes in a software application needed to be able to send e-mail messages, it would be desirable if they all did the job with methods of the same name and had the same number and type of arguments.

```
        interface Mail {
        public function sendMail();
        }
```

Then, if another class implemented that interface, like this:

```
        class Report implements Mail {
        // Definition goes here
        }
```

it would be required to have a method called `sendMail`. It's an aid to standardization.

Constants

A constant is somewhat like a variable, in that it holds a value but is really more like a function because a constant is immutable. Once you declare a constant, it does not change. Declaring one is easy, as is done in this version of `MyClass`:

```
class MyClass {

const requiredMargin = 1.3;

function __construct($incomingValue) {
// Statements here run every time an instance of the class
// is created.
}

}
```

In that class, `requiredMargin` is a constant. It is declared with the keyword `const`, and under no circumstances can it be changed to anything other than `1.3`. Note that the constant's name does not have a leading `$`, as variable names do.

Abstract Classes

An abstract class is one that cannot be instantiated, only inherited. You declare an abstract class with the keyword `abstract`, like this:

```
abstract class MyAbstractClass {

abstract function myAbstractFunction() {
}

}
```

Note that function definitions inside an abstract class must also be preceded by the keyword `abstract`. It is not legal to have `abstract` function definitions inside a nonabstract class.

Simulating class functions

Some other OOP languages make a distinction between *instance* member variables, on the one hand, and *class* or *static* member variables on the other. Instance variables are those that every instance of a class has a copy of (and may possibly modify individually); class variables are shared by all instances of the class. Similarly, instance functions depend on having a particular instance to look at or modify; class (or static) functions are associated with the class but are independent of any instance of that class.

In PHP, there are no declarations in a class definition that indicate whether a function is intended for per-instance or per-class use. But PHP does offer a syntax for getting to functions in a class even when no instance is handy. The `::` syntax operates much like the `->` syntax does, except that it joins class names to member functions rather than instances to members. For example, in the following implementation of an extremely primitive calculator, we have some functions that depend on being called in a particular instance and one function that does not:

```
class Calculator
```

```
{
  public $current = 0;
  public function add($num) {
    $this->current += $num;
  }
 public function subtract($num) {
    $this->current -= $num;
  }
  public function getValue() {
    return($current);
  }
  public function pi() {
    return(M_PI); // the PHP constant
  }
}
```

We are free to treat the `pi()` function as either a class function or an instance function and access it using either syntax:

```
$calc_instance = new Calculator;
$calc_instance->add(2);
$calc_instance->add(5);
print("Current value is " .
        $calc_instance->current ."<BR>");
print("Value of pi is " .
        $calc_instance->pi() . "<BR>");
print("Value of pi is " .
        Calculator::pi() . "<BR>");
```

This means that we can use the `pi()` function even when we don't have an instance of `Calculator` at hand. The `Calculator` class has to be accessible in either case, though, meaning that it has to have been imported with a `require_once` statement, or something similar.

Calling parent functions

Asking an instance to call a function will always result in the most specific version of that function being called, because of the way overriding works. If the function exists in the instance's class, the parent's version of that function will not be executed.

Sometimes it is handy for code in a subclass to explicitly call functions from the parent class, even if those names have been overridden. It's also sometimes useful to define subclass functions in terms of superclass functions, even when the name is available.

Calling parent constructors

In the section "Inheritance" earlier in this chapter, we showed you code (see Listing 20-3) where both subclass and superclass had constructors, and both constructors set a variable that was defined

329

by the superclass. This might be stylistically dodgy, but more importantly, we would like to avoid duplicating work across the two constructors, especially if a lot of code is involved.

Instead of writing an entirely new constructor for the subclass, let's write it by calling the parent's constructor explicitly and then doing whatever is necessary in addition for instantiation of the subclass. Here's a simple example:

```
class Name
{
  public $_firstName;
  public $_lastName;

  public function __construct($first_name, $last_name)
  {
    $this->_firstName = $first_name;
    $this->_lastName = $last_name;
  }
  public function rename() {
    return($this->_lastName .
           ", " .
           $this->_firstName);
  }
}

class NameSub1 extends Name
{
  public $_middleInitial;
  public function NameSub1($first_name, $middle_initial,
                  $last_name) {
    Name::Name($first_name, $last_name);
    $this->_middleInitial = $middle_initial;
  }
  public function rename() {
    return(Name::rename() . " " .
           $this->_middleInitial);
  }
}
```

In this example, we have a parent class (`Name`), which has a two-argument constructor, and a subclass (`NameSub1`), which has a three-argument constructor. The constructor of `NameSub1` functions by calling its parent constructor explicitly using the `::` syntax (passing two of its arguments along) and then setting an additional field. Similarly, `NameSub1` defines its nonconstructor `rename()` function in terms of the parent function that it overrides.

It might seem strange to call `Name::Name()` here, without reference to `$this`. The good news is that both `$this` and any member variables that are local to the parent are available to a parent function when invoked from a child instance.

Automatic calls to parent constructors

In a sense, constructor functions in a subclass override the constructors in superclasses. (We say "in a sense" because we usually only say that one function overrides another if the two functions have the same name; a subclass constructor and a superclass constructor always have different names.)

As you saw in the previous section, if you want both the subclass constructor and the superclass constructor to be called, you must include code in the subclass to call the superclass code explicitly. Beginning with PHP4, if a subclass lacks a constructor function and a superclass has one, the superclass's constructor will be invoked. The most specific constructor that can be found (if any) will be called — anything else is up to the programmer.

Simulating method overloading

One neat trick offered by some OOP languages (and not offered by PHP) is automatic overloading of member functions. This means that you can define several different member functions with the same name but different signatures (number and types of arguments). The language itself takes care of matching up calls to those functions with the right version of the function, based on the arguments that are given.

PHP does not offer such a capability, but the loose typing of PHP lets you take care of one half of the overloading equation — you can define a single function of a given name that behaves differently based on the number and types of arguments it is called with. The result looks like an overloaded function to the caller (but not to the definer).

Here's an example of an apparently overloaded constructor function:

```php
class MyClass
{
public $string_var = "default string";
public $num_var = 42;

  public function __construct($arg1) {
    if (is_string($arg1)) {
      $this->string_var = $arg1;
    }
    elseif (is_int($arg1) ||
            is_double($arg1)) {
      $this->num_var = $arg1;
    }
  }
}

$instance1 = new MyClass("new string");
$instance2 = new MyClass(5);
```

The constructor of this class will look to its caller as though it is overloaded, with different behavior based on the type of its inputs. You can also vary behavior based on the number of arguments by testing the number of arguments supplied by the caller.

Serialization

Serialization of data means converting it into a string of bytes in such a way that you can produce the original data again from the string (via a process known, unsurprisingly, as *unserialization*). After you have the ability to serialize/unserialize, you can store your serialized string pretty much anywhere (a system file, a database, and so on) and recreate a copy of the data again when needed.

PHP offers two functions, `serialize()` and `unserialize()`, which take a value of any type (except type `resource`) and encode the value into string form and decode again, respectively. The PHP3 implementation of object serialization wasn't very useful because member function definitions didn't survive the serialization/unserialization process; beginning with version 4, however, PHP robustly recreates all important aspects of the instance from the string, as long as the class definition is available to the code where `unserialize()` is called.

Here is a quick example, which we'll extend later in this section:

```
class ClassToSerialize {
  public $storedStatement = "data";
  public function __construct($statement) {
    $this->storedStatement = $statement;
  }
  public function display ()
  {
    print($this->storedStatement . "<BR>");
  }
}

$instance1 =
  new ClassToSerialize("You're objectifying me!");
$serialization = serialize($instance1);
$instance2 = unserialize($serialization);
$instance2->display();
```

This class has just one member variable and a couple of member functions, but it's sufficient to demonstrate that both member variables and member functions can survive serialization. We create an object, convert it to a serialized string, convert it back to a new instance, and the printed result is the accurate complaint (`You're objectifying me!`).

Of course, there is no point in serializing and unserializing an object in the same script. Serialization is only worthwhile when we expect the serialized string to outlive the script (and the

variable) that it currently lives in and be reincarnated in another execution. This may be because we store the serialization in a file or a database and read it back in again. It can also happen automatically as a result of PHP's session mechanism — variables that are registered as belonging to a session will be serialized and unserialized from page to page.

CROSS-REF For more on how the session mechanism uses serialization, see Chapter 26.

Sleeping and waking up

PHP provides a hook mechanism so that objects can specify what should happen just before serialization and just after unserialization. The special member function __sleep() (that's *two* underscores before the word sleep), if defined in an object that is being serialized, will be called automatically at serialization time. It is also required to return an array of the names of variables whose values are to be serialized. This offers a way to not bother serializing member variables that are not expected to survive serialization anyway (such as database resources) or that are expensive to store and can be easily recreated. The special function __wakeup() (again, two underscores) is the flip side — it is called at unserialization time (if defined in the class) and is likely to do the inverse of whatever is done by __sleep() (restore database connections that were dropped by __sleep() or recreate variables that __sleep() said not to bother with).

You may wonder why these functions are necessary — couldn't the code that calls serialize() also just do whatever is necessary to shut down the object? Actually, it's very much in keeping with OOP to include such knowledge in the class definition rather than expecting the code using the objects to know about their special needs. Also the calling code may have no knowledge of the object's internals at all (as in the code that serializes all session objects). The author of the class is uniquely qualified to say what should happen when an instance is sent away or revived.

As an example of how to use these functions, here is the previous serialization example, augmented with an extra variable, and the __sleep() and __wakeup() functions:

```php
class ClassToSerialize2 {
  public $storedStatement = "data";
  public $easilyRecreatable = "data again";
  public function __construct($statement) {
    $this->storedStatement = $statement;
    $this->easilyRecreatable =
      $this->storedStatement . " Again!";
  }
  public function __sleep() {
    // Could include DB cleanup code here
    return array('storedStatement');
  }
  public function __wakeup() {
    // Could include DB restoration code here
    $this->easilyRecreatable =
      $this->storedStatement . " Again!";
  }
```

```
  public function display ()
  {
    print($this->easilyRecreatable . "<BR>");
  }
}

$instance1 =
  new ClassToSerialize2("You're objectifying me!");
$serialization = serialize($instance1);
$instance2 = unserialize($serialization);
$instance2->display();
```

The variable called $easilyRecreatable is meant to stand in for a piece of data that is (1) expensive to store and (2) implied by the other data in the class anyway. The definition of __sleep() does no cleanup itself, but it returns an array that contains only one variable name and does not include easilyRecreatable. At serialization time, only the value of the variable storedStatement is included in the string. When the object is recreated, the __wakeup() function assigns a value into $this->easilyRecreatable, which is then displayed: You're objectifying me! Again!

Serialization gotchas

The serialization mechanism is pretty reliable for objects, but there are still a few things that can trip you up:

- The code that calls unserialize() must also have loaded the definition of the relevant class. (This is also true of the code that calls serialize() too, of course, but that will usually be true because the class definition is needed for object creation in the first place.)

- Object instances can be created from the serialized string only if it is really the same string (or a copy thereof). A number of things can happen to the string along the way, if stored in a database (make sure that slashes aren't being added or subtracted in the process), or if passed as url or form arguments. (Make sure that your URL-encoding/decoding is preserving exactly the same string and that the string is not long enough to be truncated by length limits.)

- If you choose to use __sleep(), make sure that it returns an array of the variables to be preserved; otherwise no variable values will be preserved. (If you do not define a __sleep() function for your class, all values will be preserved.)

Introspection Functions

While PHP lacks some features of full OO languages like Java or C++, it is surprisingly good in the esoteric area of *introspection*. (It's the classes and objects that get introspective here, not the programmer.) Introspection allows the programmer to ask objects about their classes, ask classes about

their parents, and find out all the parts of an object without have to crunch the source code to do it. Introspection also can help you to write some surprisingly flexible code, as you will see.

Function overview

Most of this section will be example-driven, but we begin by looking at the introspection functions provided by PHP. Table 20-1 summarizes these functions, what they do, and what version of PHP introduced them. (This table is essentially a reframing of information from the online manual; we offer it here mainly because it highlights features that we found somewhat confusing the first time we studied the manual.)

TABLE 20-1

Class/Object Functions

Function	Description	Operates on Class Names	Operates on Instances	As of PHP Version
get_class()	Returns the name of the class an object belongs to.	No	Yes	4.0.0
get_parent_ class()	Returns the name of the parent class of the given instance or class.	Yes (as of PHP v.4.0.5)	Yes	4.0.0, 4.0.5
class_ exists()	Returns TRUE if the string argument is the name of a class, FALSE otherwise.	Yes	No	4.0.0
get_ declared_ classes()	Returns an array of strings representing names of classes defined in the current script.	N/A	N/A	4.0.0
is_subclass_ of()	Returns TRUE if the class of its first argument (an object instance) is a subclass of the second argument (a class name), FALSE otherwise	No	Yes	4.0.0
get_class_ vars()	Returns an associative array of var/ value pairs representing the name of variables in the class and their default values. Variables without default values will not be included.	Yes	No	4.0.0
get_object_ vars()	Returns an associative array of var/ value pairs representing the name of variables in the instance and their default values. Variables without values will not be included.	No	Yes	4.0.0

TABLE 20-1	*(continued)*			
Function	**Description**	**Operates on Class Names**	**Operates on Instances**	**As of PHP Version**
method_ exists()	Returns TRUE if the first argument (an instance) has a method named by the second argument (a string) and FALSE otherwise.	No	Yes	4.0.0
get_class_ methods()	Returns an array of strings representing the methods in the object or instance	Yes	Yes (as of v4.0.6)	4.0.0, 4.0.6
call_user_ method_ array()	Same as call_user_ method(), except that it expects its third argument to be an array containing the arguments to the method.	No	Yes	4.0.5

These functions break down into the following four broad categories:

- Getting information about the class hierarchy
- Finding out about member variables
- Finding out about member functions
- Actually calling member functions

The first group of functions (get_class() through instanceof()) deal with discovering what classes exist, asking an object about its class, and discovering class inheritance relationships. Some of these functions start with an instance of an object, some start with the class name as a string, and some are happy with either one. (We've included columns in the table to try to clarify this.) Note that after we have the get_class() function, it's easy to satisfy functions that require a class as input; for example, if get_parent_class() insists on a class name, and we want to know the parent class of an object instance, we could just wrap it like this:

```
$parent_class = get_parent_class(get_class($my_instance));
```

Bear in mind that as of PHP4.3, the constant __CLASS__ exists. It contains the class name.

Going in the other direction (trying to satisfy a function that wants an instance when all we have is a class) would be more problematic because you don't want to instantiate a class just to ask questions of it.

The second group of functions (get_class_vars(), get_object_vars()), return an associative array containing member variables and their values. The keys of these arrays are the names of the

variables as strings (without leading $ symbols), and the array values are the values of those variables in the object or class. In both cases (for reasons unknown to your authors), only member variables that actually have a value are returned.

The difference between `get_class_vars()` and `get_object_vars()` is subtle, but it's more than just a question of what type of input they prefer. The `get_class_vars()` function returns information about variables and default values as they exist in the class definition itself, independent of any instance; `get_object_vars()` returns information about the current state of a particular instance. For example, consider this class definition and use:

```php
class Example {
  public $var1 = "initialized";
  public $var2 = "initialized";
  public $var3;
  public $var4;
  public function __construct() {
    $this->var3 = "set";
    $this->var1 = "changed";
  }
}

$example = new Example();
print_r(get_class_vars("Example"));
print_r(get_object_vars($example));
```

For the first call (to `get_class_vars()`), we should expect to find `var1` and `var2` both bound to `"initialized"` as in the class definition itself. The second call (to `get_object_vars()`) should return bindings of `var1`, `var2`, and `var3` to `"changed"`, `"initialized"`, and `"set"`, respectively. In neither case will either function retrieve `var4`.

The third group of functions (`method_exists()`, `get_class_methods()`) manipulate member function names as strings. The first allows you to ask an instance if it contains a given function, and the second recovers all function names from an instance or class. (Notice that we don't need two separate functions as we did with `get_class_vars()` and `get_object_vars()`; PHP doesn't offer you a way to add or delete member functions from instances on the fly.)

Finally, the fourth group lets you apply method names (presumably recovered using functions from the third group) to instances. But these are probably best explained by example, so let's dive in.

Example: Class genealogy

Consider the following, somewhat confusing, class hierarchy.

```php
class Color {}
class Control extends UIelement {}
class Widget extends Control { }
class Button extends Widget {}
class Pulldown extends Widget {}
```

337

```
class Clicker extends Button {}
class Blue extends Color {}
class Displayer extends UIelement {}
class UIElement {}
class LightBlue extends Blue {}
```

Now imagine that we'd like to have a better visualization of this tangle, just for purposes of documentation. For starters, it's pretty easy to use the get_parent_class() function to figure out the classes that a given class descends from:

```
public function print_ancestry($class_name) {
  print("Class ancestry: ");
  print_ancestry_aux($class_name);
  print("<BR>");
}

public function print_ancestry_aux ($class_name) {
  print("$class_name");
  if ($parent = get_parent_class($class_name)) {
    print(" => ");
    print_ancestry_aux($parent);
  }
}
print_ancestry("Clicker");
```

Which gives us the somewhat informative output:

```
Class ancestry: Clicker => button => widget => control => uielement
```

(Notice that our retrieved class names have become lowercase. This happens to user-defined classes, whereas prior to PHP 6, built-in classes should have their capitalization intact.)

Getting a view of the entire class tree is a little bit harder, because PHP doesn't offer a straight-forward way to retrieve child classes given a parent class. Our recourse is the get_declared_classes function, which tells us all the classes that are defined in the current script — we can then somewhat inefficiently do paternity tests on all known classes to discover the children of a given class (see Listing 20-6).

LISTING 20-6

Class genealogy

```
public function same_class_name ($string1, $string2) {
  return ((strtolower($string1)) ==
          (strtolower($string2)));
}

public function get_child_classes ($parent) {
```

```
  $all_classes = get_declared_classes();
  $children = array();
  foreach ($all_classes as $candidate) {
    if (same_class_name($parent,
        get_parent_class($candidate)) &&
      !same_class_name($parent, $candidate)) {
      array_push($children, $candidate);
    }
  }
  return($children);
}

public function print_class_tree () {
  $all_classes = get_declared_classes();
  print("<PRE>");
  print("CLASS HIERARCHY:\n");
    foreach ($all_classes as $candidate) {
      if (!get_parent_class($candidate)) {
        print_class_tree_aux($candidate, 0);
      }
    }
  print("</PRE>");
}

public function print_class_tree_aux ($parent, $level) {
  for ($x = 0; $x < $level; $x++) {
    print("     ");
  }
  print("$parent<BR>");
  $children = get_child_classes($parent);
  foreach ($children as $child) {
    print_class_tree_aux($child, $level + 1);
  }
}
print_class_tree();
```

We start off this listing by defining what it means for two class names to be the same. This may be overkill, but converting every name to lowercase before comparison lets us stop worrying about whether we'll be tripped up by case issues. Then we define a general function to retrieve child classes (inefficiently, but it should make no difference unless your class hierarchy grows to be very, very large). The print_class_tree() function essentially recovers all *orphans* or *roots* (classes without parents) and prints each one individually as a tree. The auxiliary function handles printing a rooted tree — first the parent and then indented children. Finally, we wrap the whole thing in a <PRE></PRE> construct so we can just use spaces for indenting. The result looks like this:

```
CLASS HIERARCHY:
stdClass
```

```
__PHP_Incomplete_Class
OverloadedTestClass
Directory
color
    blue
        lightblue
uielement
    control
        widget
            button
                clicker
            pulldown
    displayer
```

The first few classes printed are unfamiliar and not defined in your code file. These either belong to the PHP implementation itself or to auxiliary packages that you have compiled — the precise classes that you see when you execute this code may vary.

Example: matching variables and DB columns

One frequent use for PHP objects in database-driven systems is as a wrapper around the entire database API. The theory is that the wrapper insulates the code from the specific database system, which will make it trivial to swap in a different RDBMS when the technical needs change. (We've never seen it work out quite this way in practice, but . . . don't get us started.) Another use that is almost as common (and that your authors like better) is to have object instances correspond to database result rows. In particular, the process of reading in a result row looks like instantiating a new object that has member variables corresponding to the result columns we care about, with extra functionality in the member functions. As long as the fields and columns match up (and as long as you can afford object instantiation for every row), this can be a nice abstraction away from the database.

A repetitive task that arises when writing this kind of code is assigning database column values to member variables, in individual assignment statements. This feels like it should be unnecessary, especially when the columns and the corresponding variables have exactly the same names. In this section, we write a hack to automate this process.

For concreteness, let's start with an actual database table. Following are the MySQL statements necessary to create a simple table and insert one row into it:

```
mysql> create table book
   (id int not null primary key auto_increment,
    author varchar(255), title varchar(255),
    publisher varchar(255));
mysql> insert into book (author, title, publisher)
       values ("Robert Zubrin", "The Case For Mars",
               "Touchstone");
```

Because the id column is auto-incremented, it will happen to have the value 1 for this first row.

The code in Listing 20-7 assumes a database called `oop` with the table created as above, and also that we have a file called `dbconnect_vars` that sets `$host`, `$user`, and `$pass` appropriately for our particular MySQL setup. There is also little or no error checking (the code assumes the connection works, that the row was retrieved successfully, and so on). The main point we want to highlight is the hack in the middle of the `Book` constructor.

LISTING 20-7

Matching variables and columns

```php
<?php
include_once("dbconnect_vars.php");

class Book
{
public $id;

// variables corresponding to DB columns
public $author = "DBSET";
public $title = "DBSET";
public $publisher = "DBSET";

  public function __construct($db_connection, $id) {
    $this->id = $id;
    $query = "select * from book " .
             "where id = $id";
    $result = mysql_query($query, $db_connection);
    $db_row_array =
      mysql_fetch_array($result);
    $class_var_entries =
      get_class_vars(get_class($this));
    while ($entry = each($class_var_entries)) {
      $var_name = $entry['key'];
      $var_value = $entry['value'];
      if ($var_value == "DBSET") {
        $this->$var_name =
          $db_row_array[$var_name];
      }
    }
  }

  public function rename () {
    $return_string = "BOOK<BR>";
    $class_var_entries =
      get_class_vars(get_class($this));
    while ($entry = each($class_var_entries)) {
      $var_name = $entry['key'];
      $var_value = $this->$var_name;
```

```
        $return_string .=
            "$var_name: $var_value<BR>";
    }
    return($return_string);
  }
}
$connection =
  mysql_connect($host, $user, $pass)
  or die("Could not connect to DB");
mysql_select_db("oop");
$book = new Book($connection, 1);
$book_string = $book->rename();
?>
<HTML><HEAD></HEAD><BODY>
<?php echo $book_string ?>
</BODY></HTML>
```

The database query returns all columns from the `book` table, and the values are indexed in the result array by the column names. The constructor then uses `get_class_vars()` to discover all the variables that have been set in the object, tests them to see if they have been bound to the string `"DBSET"`, and then sets those variables to the value of the column of the same name.

The result is the output:

```
BOOK
Author: Robert Zubrin
Title:  The Case For Mars
Publisher:  Touchstone
```

If we add fields to the database table definition, the only change we will need to make to Listing 20-7 is to add appropriately named variables to the class definition and initialize them to `"DBSET"`. (We use this initialization to be clear about which variables should be overwritten, but also because we cannot retrieve the variables at all unless they have been initialized.)

Example: Generalized test methods

As a final introspection example, suppose that we are working on a large OOP project, with complex objects that need to maintain a lot of internal state. Testing is extremely important, because bugs will creep in and waste our time if we don't catch them early on.

So let's adopt some testing conventions for this project. As one of them, let's agree that any class in our system can (optionally) define a member function called `selfTest()`. The point of this function is to test the object instance it is called on to make sure the data in the object is valid and consistent across the instance. The `selfTest()` function should always return `FALSE` if everything is okay and a diagnostic string if something is wrong. The coders of the objects agree that they will write these tests in such a way that a test can be applied at any time during execution.

If we agree on such a framework, we can write a general object tester. The tester simply calls selfTest() on any object it is pointed at, if such a method has been defined for that object. To make it easier to apply, we'll also make the object tester accept arrays of objects, and test each component object individually. Such an object tester is in Listing 20-8, along with some sample class definitions that have selfTest() defined.

LISTING 20-8

ObjectTester

```php
class Namestring {
  public $name;
  public $nameLength;
  public $checksum;

 public function __construct($string_in) {
    $this->name = $string_in;
    $this->nameLength = strlen($string_in);
    $this->checksum =
       $this->computeChecksum($string_in);
  }

  public function setName ($new_string) {
    $this->name = $new_string;
    $this->nameLength = strlen($new_string);
    $this->checksum =
      $this->computeChecksum($new_string);
  }

  public function computeChecksum ($string) {
    // not a good checksum in practice
    $sum = 0;
    for ($x = 0;
         $x < strlen($string);
         $x++) {
      $sum += ord($string[$x]);
    }
    return($sum % 100);
  }

  public function selfTest () {
    // returns FALSE if everything is OK
    if ($this->nameLength !=
        strlen($this->name)) {
      return("Name $this->name not of ".
             "length $this->nameLength!");
```

```
      }
      elseif
        ($this->checksum !=
            $this->computeChecksum($this->name)) {
        return("Name $this->name fails checksum!");
      }
      else {
        return(FALSE);
      }
    }
  }

  class NonTestingObject {
  }

  class ObjectTester {
    public function ObjectTester() {
      // empty constructor
    }

    public function test ($thing) {
      if (is_object($thing)) {
        if (method_exists($thing, 'selfTest')) {
          $this->handleTest(
            call_user_func('selfTest', $thing));
        }
      }
      elseif (is_array($thing)) {
        foreach ($thing as $component) {
          $this->test($component);
        }
      }
      // ignore if not an array or object
    }
    public function handleTest ($result) {
      if ($result) {
        print("Warning: $result");
      }
    }
  }
```

The Namestring object in Listing 20-8 has several pieces of data, which must be kept consistent with each other. Using the constructor to build an instance of Namestring keeps them consistent, as does changing the name with setName. Namestring also defines selfTest(), which cross-checks the name, the length of the name, and a primitive checksum.

Now let's see how to use the `ObjectTester` class with some sample `Namestring` data:

```
$object_list = array();
array_push($object_list, new Namestring("Jordan"));
array_push($object_list, new Namestring("Rodman"));
array_push($object_list, new NonTestingObject);
array_push($object_list, new Namestring("Pippen"));

$tester = new ObjectTester($object_list);

print("Running test..<BR>");
$tester->test($object_list);

print("Changing name..<BR>");
$current_object = &$object_list[0]; // note reference!
$current_object->setName("Michael");
print("Running test..<BR>");
$tester->test($object_list);

print("Changing name..<BR>");
$current_object = &$object_list[1]; // note reference!
$current_object->name = "Jordan";

print("Running test..<BR>");
$tester->test($object_list);
```

The results of running this code are:

```
Running test..
Changing name..
Running test..
Changing name..
Running test..
Warning:  Name Jordan fails checksum!
```

This warning resulted because we messed with the object's data directly the second time, rather than using the approved method for changing the name.

We've used toy self-testing classes here, but the basic approach extends easily to more complex classes. Among possible extensions is more interesting handling of the warning messages (and possibly interrupting execution). Another extension would be to use introspection on member variables themselves, as well as array components, to find contained objects and test those. This would mean defining the test runner recursively so that a thing passes a `selfTest()` if (1) its own `selfTest()` method (if it exists) finds no problem, and (2) any components (member variables, array slots) also pass `selfTest()`. (Watch out for circularities though! If the tester is ever called on objects that mutually refer to each other, it would have to be rewritten to track the identities of previously seen objects and would only test each object once.)

Extended Example: HTML Forms

All the OOP code you've seen so far in this chapter has been fairly short, so in this chapter we present an extended piece of code for your enjoyment, shown in Listing 20-9.

The point of this class is to semiautomate the production of HTML forms, which one of your authors has always found to be a bit of a pain to generate. The top-level class represents a form, while other classes represent inputs, text areas, and hidden variables (just the ones that your author uses most frequently). The idea is that you can make a form by adding input fields to an existing object and display the form upon request. The resulting form will be not be especially pretty (every element is displayed sequentially down the left-hand side of the page), but it's good enough for situations where, say, you want to enter some information into your own database yourself.

LISTING 20-9

form_printer.php

```php
<?php

// ---- The form class itself ---

class HtmlForm {

  // suitable for generating quick & dirtyforms

  public $actionTarget; // path to receiving page
  private  $inputForms;  // array of HtmlFormInput
  public $hiddenVariables; // associative name/val

  // CONSTRUCTOR
  public function __construct($action_target) {
    $this->actionTarget = $action_target;
    $this->inputForms = array();
    $this->hiddenVariables = array();
  }

  // PUBLIC METHODS
  public function rename () {
    $return_string = "";
    $return_string .=
      "<FORM METHOD=\"POST\" ".
      "ACTION=\"$this->actionTarget\">\n";
    $return_string .= $this->inputFormsString();
    $return_string .= $this->hiddenVariablesString();
    $return_string .= "<BR>\n";
    $return_string .= $this->submitButtonString();
    $return_string .= "</FORM>";
```

```php
      return($return_string);
    }

    // adding elements to form

    public function addInputForm ($input_form) {
      if (!isSet($input_form) ||
          !is_object($input_form) ||
          !is_subclass_of($input_form,
                    'htmlforminput')){
        die("Argument to HtmlForm::addInputForm ".
            "must be instance of HtmlFormInput.".
            "  Given argument is of class " .
            get_class($input_form));
      }
      else {
        array_push($this->inputForms, $input_form);
      }
    }

    public function addInputButton ($input_button) {
      if (!isSet($input_button) ||
          !isObject($input_button) ||
          !is_a($input_button, 'HtmlInputButton')){
        die("Argument to HtmlForm::addInputButton ".
            "must be instance of HtmlInputButton");
      }
      else {
        array_push($this->inputButtons, $input_button);
      }
    }

    public function addHiddenVariable ($name, $value) {
      if (!isSet($value)) {
        die("HtmlForm::addHiddenVariable requires ".
            "two arguments (name and value)");
      }
      else {
        $this->hiddenVariables[$name] = $value;
      }
    }

    public function inputFormsString () {
      $return_string = "";
      $form_array = $this->inputForms;
      foreach ($form_array as $input_form) {
        $return_string .=
          "<B>$input_form->heading</B>";
        if ($this->headingElementBreak()) {
```

347

```
        $return_string .= "<BR>";
      }
      $return_string .= $input_form->rename();
      $return_string .= "<BR>\n";
    }
    return($return_string);
  }

  public function hiddenVariablesString () {
    $return_string = "";
    while ($hidden_var =
            each($this->hiddenVariables)) {
      $var_name = $hidden_var['key'];
      $var_value = $hidden_var['value'];
      $return_string .=
        "<INPUT TYPE=HIDDEN " .
        "NAME=$var_name ".
        "VALUE=$var_value >";
        $return_string .= "\n";
    }
    return($return_string);
  }

  public function headingElementBreak () {
    // override to disable breaks after headings,
    // or to do more complicate layout
    return(TRUE);
  }

  public function submitButtonString () {
    $return_string = "<INPUT TYPE=Submit " .
                     " VALUE=Submit >\n";
    return($return_string);
  }
}

// ---- Classes for parts of a form ----

abstract class HtmlFormInput {
  public $name; // The variable name for form submission
  public $heading; // The visible label on form
  function __construct() {
    die("Class HtmlFormInput intended only " .
        "to be subclassed");
  }
  function rename () {
    die("Subclass of HtmlFormInput missing " .
        "definition of rename()");
  }
}
```

```php
class HtmlFormSelect extends HtmlFormInput
{
  public $_valueArray = array();
  public $_selectedValue;

  public function __construct ($name, $heading,
                              $value_array,
                              $selected_value=NULL) {
    if (!isSet($value_array)) {
      die("HtmlFormSelect needs a minimum of two " .
          "arguments: a name, and value array");
    }
    elseif (!is_array($value_array)) {
      die("Third argument to HtmlFormSelect()" .
          "should be array where keys are values ".
          "submitted, and values are display values");
    }
    else {
      // actual initialization
      $this->name = $name;
      $this->heading = $heading;
      $this->_valueArray = $value_array;
      $this->_selected_value = $selected_value;
    }
  }

  public function rename () {
    $return_string = "";
    $return_string .=
      "<SELECT NAME=\"$this->name\">";
    while ($var_entry =
            each($this->_valueArray)) {
      $submit_value = $var_entry['key'];
      $display_value = $var_entry['value'];
      if ($submit_value == $this->_selected_value) {
        $return_string .=
          "<OPTION VALUE=${submit_value} SELECTED >";
      }
      else {
        $return_string .= "<OPTION VALUE=${submit_value}>";
      }
      $return_string .= $display_value;
    }
    $return_string .=
      "</SELECT>";
    return($return_string);
  }
}

class HtmlFormText extends HtmlFormInput
```

349

```
{
  public $initial_value;

  public function __construct ($name,
                               $heading,
                               $initial_value="")
  {
    // Initialization of member vars
    if (!isSet($name) ||
        !isSet($heading)) {
      die("HtmlFormText constructor needs " .
          "at least two arguments (name, heading)");
    }
    $this->name = $name; // name defined in parent
    $this->heading = $heading; // defined in parent
    $this->initial_value = $initial_value;
  }

  public function rename () {
    $return_string = "";
    $return_string .= "<INPUT TYPE=TEXT ";
    $return_string .= "NAME=\"$this->name\" ";
    $return_string .=
        "VALUE=\"$this->initial_value\" ";
    $return_string .= " >";
    return($return_string);
  }
}

class HtmlFormTextArea extends HtmlFormInput {
  public $initial_value;
  public $rows;
  public $cols;
  public $wrapType;

  public function __construct ($name,
                               $heading,
                               // optional args:
                               $initial_value="",
                               $rows=1, $cols=60,
                               $wrapType="VIRTUAL")
  {
    // Initialization of member vars
    if (!isSet($name)) {
      die("HtmlFormTextArea constructor needs " .
          "at least two arguments (name, heading)");
    }
    $this->name = $name; // name defined in parent
    $this->heading = $heading; // name defined in parent
    $this->initial_value = $initial_value;
```

```php
    $this->rows = $rows;
    $this->cols = $cols;
    $this->wrapType = $wrapType;
  }

  public function rename ()
  {
    $return_string = "";
    $return_string .= "<TEXTAREA ";
    $return_string .= "NAME=\"$this->name\" ";
    $return_string .= "ROWS=$this->rows ";
    $return_string .= "COLS=$this->cols ";
    $return_string .= "WRAP=$this->wrapType ";
    $return_string .= $this->additionalAttributes();
    $return_string .= ">";
    $return_string .= $this->initial_value;
    $return_string .= "</TEXTAREA>";
    return($return_string);
  }

  public function additionalAttributes () {
    // OVERRIDE THIS to return a string with
    // TextArea attributes other than
    // NAME, ROWS, COLS, and WRAP
    return("");
  }
}
?>
```

The basic design for all these objects includes a constructor function with default arguments and a rename() method that returns HTML for the form or piece thereof. Forms store pieces of input (which might conceivably be reordered or laid out by a more sophisticated version), and recursively call rename() on these pieces. The HTML form elements that are supported are: TEXTAREA, TEXT, and SELECT.

Here is an example of calling this code to generate a simple form page:

```php
<HTML><HEAD></HEAD><BODY>
<?php include("form_printer.php");
$my_form = new HtmlForm($PHP_SELF);
$my_form->addInputForm(
        new HtmlFormText("firstname",
                        "First Name"));
$my_form->addInputForm(
        new HtmlFormText("lastname",
                        "Last Name"));
$my_form->addInputForm(
        new HtmlFormSelect(
            "age",
            "Age",
```

351

```
                    array(0 => "0 - 9",
                          1 => "10 - 19",
                          2 => "20 - 29",
                          3 => "Senior citizen"),
                    2));
        $my_form->addInputForm(
                new HtmlFormTextArea(
                    "feedback",
                    "What's on your mind?",
                "[Please fill in your own personal rant]",
                    5));
        print($my_form->rename());
        ?>
        </BODY>
        </HTML>
```

Much of the form-producing code is straightforward and is concerned with churning out various kinds of HTML syntax. There are two interesting things to notice from the point of view of OOP-in-PHP, however.

The first is that the `HtmlFormInput` class is *designated* `abstract`. That is, it exists not to be instantiated but only to be inherited from. The second point of interest is that the `HtmlForm` class has an array that is intended to hold `HtmlFormInput` objects. Of course, because PHP is loosely typed, we cannot enforce that in any way at compile time, although the manufacturer-approved way to insert new forms (`addInputForm()`) does some type-checking on insertion. If users of this class rely only on this method, we can be assured that everything that ends up in that array will be an instance of `HtmlFormInput` (or subclass thereof) and so should be a well-behaved form element when display time comes around. The `private` designation guarantees that the array cannot be manipulated from outside the class at runtime.

Gotchas and Troubleshooting

In the spirit of Chapter 10, we offer in the following sections the top-two most likely symptoms of problematic OOP code, along with the most likely cause.

Symptom: Member variable has no value in member function

This could have many causes, of course, but the most common is simply a confusion about the right way to refer to member variables. The syntax is:

```
$this->member_name
```

If, instead, your function simply refers to $member_name, that will usually be an unbound variable and, at any rate, will never succeed in referring to the member variable. Similarly, if your function refers to $this->$member_name, you are asking for the field named by the string in the variable $member_name (which is probably unbound).

Symptom: Parse error, expecting T_VARIABLE . . .

There are of course many ways to munge a class definition so that PHP will complain when it tries to parse it. One of the most common errors again has to do with placement of those $ symbols. A class declaration like the following:

```
class MyClass {
  public my_var;  // WRONG
}
```

inevitably gives you a parse error of some sort because the syntax requires a $ before my_var.

OOP Style in PHP

The topic of OOP programming style is a huge one (because it includes OOP design!) and is well beyond the scope of this book. In the spirit of Chapter 32, however, we offer in the following sections some brief notes on writing readable, maintainable PHP OOP code.

Naming conventions

In this section, we simply pass along the parts of the PEAR coding style that pertain to objects.

CROSS-REF For more information on the PEAR project and the PEAR coding style, see Appendix E or the PEAR web site (at `http://pear.php.net`).

PEAR recommends that class names begin with an uppercase letter and (if in a PEAR-approved directory hierarchy of packages) have that inclusion path in the class name, separated by underscores. So your class that counts words, and that belongs to a PEAR package called TextUtils, might be called TextUtils_WordCounter. If building large OOP packages, you may want to emulate this underscore convention with your own package names; otherwise, you can simply give your classes names like WordCounter.

Member variables and member function names should have their first real letter be lowercase and have *word* boundaries be delineated by capitalization. In addition, names that are intended to be private to the class (that is, they are used only within the class, and not by outside code) should start with an underscore. So the variable in your WordCounter class that holds the count of words might be called wordCount (if intended to be messed with from the outside) or _wordCount (if it is intended to be private to the class).

Accessor functions

Another style of documenting your intent about use of internal variables is to have your variables marked as private, in general, and provide "getter" and "setter" functions to outside callers. For example, we might define a class like this:

```
class Customer
{
    private var _name;
    private var _creditCardNumber;
    private var _rating;

    function getName ()
    {
        return($this->_name);
    }

    function getRating ()
    {
        return($this->_rating);
    }

    function setRating($rating)
    {
        $this->_rating = $rating;
    }
    [... more functions ]
}
```

This class definition has three private variables: one (_creditCardNumber) that should neither be set nor retrieved from outside code, another (_name) that outside code should be able to retrieve but not set, and a third (_rating) that outside code should feel free to both get and set.

Although PHP class syntax lets you interleave variables with function definitions, it's a good idea, in general, to organize your code so that similar items with similar usage intent are located together in the class definition. For example, you might develop the habit of laying out class functions like this:

```
class myClass
{
// Public variables:
..
// Private variables
..
// Constructor

// Public functions
..
// Private functions
..
}
```

Designing for inheritance

The question of exactly how to design a class hierarchy is, as we've said, a vast area of study unto itself, and we're not about to try to contribute to it here. Just as a stylistic matter, though, it's worth thinking about whether you intend your class to be inherited from, and then try to indicate your decision, either with comments or in the structure of the definition.

For example, you may intend that your class should never breed, in which case you might just indicate that in comments, and then stop worrying about inheritance issues. (There is currently no way in PHP to enforce that a class cannot be inherited from.) At the other end of the spectrum, you might have all or part of your class intended only for inheritance. You can indicate this in comments, or you can use the trick we used in the definition of `HtmlFormInput` in Listing 20-9: Provide methods that die informatively when called directly in the base class. Finally, of course, you may have some methods that can be called directly in the base class but are especially intended for overriding. You may want to group these "hook" methods together in a clearly marked section of your class definition, so that the later writer of a derived class can quickly figure out what options are available for specializing the class's behavior. (Remember that the clueless coder of the future that you are helping may well be yourself.)

Summary

PHP provides the basics to support object-oriented programming. Among other things, the OOP syntax in PHP allows for programmer-defined classes with member variables and member data and offers single inheritance, constructor functions, object serialization, and functions for introspection. Nothing in PHP requires that you write in an object-oriented style, but if you prefer that style you can write almost all your code that way. PHP was not originally intended to be an object-oriented language, and developers with OOP experience will miss some aspects of more mature OOP languages. On the other hand, the OOP extension is usable, fairly mature, pretty stable, and widely used. It provides an extra layer of organization that can be helpful when maintaining complex code and offers a nice way to package code for distribution and reuse.

Chapter 21

Advanced Array Functions

I n Chapter 8 we introduced you to arrays, their uses, and some handy functions for working with them. In some subsequent chapters, we saw how PHP returns many of its results as arrays, particular when working with database function sets. This chapter will look at some of the more advanced functions for working with PHP arrays.

Transformations of Arrays

PHP offers a host of functions for manipulating your data once you have it nicely stored in an array. What the functions in this section have in common is that they take your array, do something with it, and return the results in another array. (We will defer the array-sorting functions until a later section.)

CROSS-REF Not covered in this chapter are `explode()` and `implode()`, which convert strings into arrays and vice versa. We cover these very handy functions in Chapter 22.

In Chapter 8, we incrementally developed a function to print out the entire contents of an array, and in this section we will use the last of these

(print_keys_and_values_each()) to show the arrays that are being returned in examples. We'll list this function again here, in a more generic form:

```
function print_keys_and_values_each($array_to_test)
{  // reliably prints everything in array
  reset($array_to_test);
  while ($array_cell = each($array_to_test))
  {
    $current_value = $array_cell['value'];
    $current_key = $array_cell['key'];
    print("Key: $current_key; Value: $current_value<BR>");
  }
}
```

Retrieving keys and values

The array_keys() function returns the keys of its input array in the form of a new array where the keys are the stored values. The keys of the new array are the usual automatically incremented integers, starting from 0. The array_values() function does exactly the same thing, except the stored values are the values from the original array. If we start with an array like the following:

```
$pizza_requests = array('Alice' => 'pepperoni',
                        'Bob' => 'mushrooms',
                        'Carl' => 'sausage',
                        'Dennis' => 'mushrooms');
```

and then we print the arrays resulting from calls to the these two functions:

```
print("Array keys:<BR>");
print_keys_and_values_each(array_keys($pizza_requests));
print("Array values:<BR>");
print_keys_and_values_each(array_values($pizza_requests));
```

we get output like this:

```
Array keys:
Key: 0; Value: Alice
Key: 1; Value: Bob
Key: 2; Value: Carl
Key: 3; Value: Dennis
Array values:
Key: 0; Value: pepperoni
Key: 1; Value: mushrooms
Key: 2; Value: sausage
Key: 3; Value: mushrooms
```

The second of these (array_values()) may seem uninteresting because we have essentially taken our old array and produced a new one with the keys renamed to successive numbers.

We can do something slightly more useful (and more helpful for ordering) with the function `array_count_values()`. This takes an array and returns a new array, where the old values are now the new keys and the new values are the number of times each old value occurs in the original array.

```
print_keys_and_values_each(
    array_count_values($pizza_requests));
```

gives us:

```
Key: pepperoni; Value: 1
Key: mushrooms; Value: 2
Key: sausage; Value: 1
```

Flipping, reversing, and shuffling

A function that is even more odd is `array_flip()`, which changes the keys of an array into the values, and vice versa. For example:

```
print_keys_and_values_each(array_flip($pizza_requests));
```

gives us:

```
Key: pepperoni; Value: Alice
Key: mushrooms; Value: Dennis // what happened to Bob?
Key: sausage; Value Carl
```

Notice that, although array keys are guaranteed to be unique, array values are not — because of this, any duplicate values in the original array become the same key in the new array. Only one of the original keys will survive to become the corresponding new value.

Reversing an array is more simple: `array_reverse()` returns a new array with the key/value pairs in reverse order. So, with the usual printing test:

```
print_keys_and_values_each(array_reverse($pizza_requests));
```

we get the result:

```
Key: Dennis; Value: mushrooms
Key: Carl; Value: sausage
Key: Bob; Value: mushrooms
Key: Alice; Value: pepperoni
```

In this case, although the internal order has been reversed, all the key/value pairs end up being the same. However, this function (like several other PHP array functions) treats integer keys somewhat special. It assumes that the ordering of integer keys on those key/value pairs should also be reversed for the later use of code that pays attention to the ordering of keys, rather than using

the internal linked-list ordering. So, `array_reverse()` swaps integer keys to make the new key ordering match the internal list. Dennis, in other words, is now actually at position 0.

If you need some extra randomness in your life, the `shuffle()` function can give it to you — `shuffle()` takes an array argument and pseudo-randomizes the order of the elements in the array. It uses `rand()`, a function that generates successive pseudo-random numbers. Before you use `shuffle()`, you need to have seeded the random-number generator with a call to `srand()`. (See the discussion of random-number generation in Chapter 9.) A reasonable calling sequence looks like this:

```
srand((double)microtime() * 1000000);  // for random # gen
shuffle($pizza_requests);
print_keys_and_values_each(array_flip($pizza_requests));
```

which *might* give us output like:

```
Key: Carl; Value: sausage
Key: Bob; Value: mushrooms
Key: Dennis; Value: mushrooms
Key: Alice; Value: pepperoni
```

CAUTION Unlike many of the array functions in this chapter, `shuffle()` is *destructive,* meaning that it operates directly on its array argument and changes it, rather than returning a newly created array. (Functions that return a new thing without disturbing their arguments might be called *constructive*, or just *nondestructive*.) Among other things, this means that the correct way to call the shuffle function is not:

```
$my_new_array = shuffle($my_old_array); //WRONG!
```

especially because the `shuffle()` function does not return a value. Instead, the right call is:

```
shuffle($my_array);   // change the array itself
```

Merging, padding, slicing, and splicing

If we want to combine two arrays for a more complete list, the function to use is `array_merge()`. This function takes two or more arrays as arguments and returns a renumbered new array that is the second array tacked onto the end of the first. If we create a new array containing some additional pizza requests like this:

```
$more_pizza_requests = array('Ted' => 'anchovies',
                             'MrWilson' => 'pineapple',
                             'Dagwood' => 'ham');
```

then we can use `array merge();` as:

```
$all_requests = array_merge($pizza_requests, $more_requests);
```

and then use our handy array inspecting function:

```
print_keys_and_values_each($all_requests);
```

We should see:

```
Key: Alice; Value: pepperoni
Key: Bob; Value: mushrooms
Key: Carl; Value: sausage
Key: Dennis; Value: mushrooms
Key: Ted; Value: anchovies
Key: MrWilson; Value: pineapple
Key: Dagwood; Value: ham
```

The `array_pad()` function is used to create some leading or following key/value pairs increasing the size of an array. It takes an input array as its first argument, then a number of elements to increase the array to, and then a value to assign to the added elements. A positive integer in the second argument will pad the end of the array; a negative integer will pad the beginning. If the second argument is smaller than the size of the array, no padding is performed.

```
$requests = array_pad($pizza_requests, 10, 'mushrooms')
//do we have any mushroom fans in the audience tonight?
```

With our function, we'd get:

```
Key: Alice; Value: pepperoni
Key: Bob; Value: mushrooms
Key: Carl; Value: sausage
Key: Dennis; Value: mushrooms
Key: 0; Value: mushrooms
Key: 1; Value: mushrooms
Key: 2; Value: mushrooms
Key: 3; Value: mushrooms
Key: 4; Value: mushrooms
Key: 5; Value: mushrooms
```

If we make the second argument negative, the new elements appear at the beginning of the array. Note that the automatically assigned keys start at 0, even though they are in the fifth position.

Somewhat more complicated are the `array_slice()` and `array_splice()` functions. The first of these returns a subset of an input array by accepting an offset and a length as its second and third arguments, respectively:

```
$subset = array_slice($pizza_requests, 1, 2);
// returns mushrooms and sausage
```

The `array_splice()` function is similar, but it accepts a fourth argument, which can be an array of any length, to splice into the input array, again returning an all new array:

```
$super_set = array_splice($pizza_requests, 2, 0,
$more_requests);
```

which will return an array like:

```
Key: Alice; Value: pepperoni
Key: Bob; Value: mushrooms
Key: Ted; Value: anchovies
Key: MrWilson; Value: pineapple
Key: Dagwood; Value: ham
Key: Carl; Value: sausage
Key: Dennis; Value: mushrooms
```

These array-manipulating functions are summarized in Table 21-1.

TABLE 21-1

Array Transformation Functions

Function	Behavior
`array_keys()`	Takes a single array argument and returns a new array where the new values are the keys of the input array, and the new keys are the integers incremented from zero.
`array_values()`	Takes a single array argument and returns a new array where the new values are the original values of the input array, and the new keys are the integers incremented from zero.
`array_count_values()`	Takes a single array argument and returns a new array where the new keys are the old array's values, and the new values are a count of how many times that original value occurred in the input array.
`array_flip()`	Takes a single array argument and changes that array so that the keys are now the values and vice versa.
`array_reverse()`	Takes a single array argument and changes the internal ordering of the key/value pairs to reverse order. Numerical keys will also be renumbered.
`shuffle()`	Takes a single array argument and randomizes the internal ordering of key/value pairs. Also renumbers integer keys to match the new ordering. This function itself uses the random-number generator `rand()`, so `srand()` must be called to seed the generator before the call to `shuffle()`.
`array_merge()`	Takes two array arguments, merges them, and returns the new array, which has (in order) the first array's elements and then the second array's elements. (*Note:* This is most useful for arrays that are being used for simple linked lists rather than for their associative keys, because keys that appear in both arrays will have one of the values overwritten. Also, numerical keys will be renumbered from 0 to reflect the new ordering.)

Function	Behavior
`array_pad()`	Takes three arguments: an input array, a pad size, and a value to pad with. Returns a new array that is "padded" by the following rules: If the pad size is greater than the length of the input array, the array is lengthened with the pad value to the pad size, as though by successive assignments like `$my_array[] = $pad_value`. A negative pad size will act the same way with the absolute value of that pad size, except that the padding will occur at the beginning of the array rather than the end. If the array is already longer than the (absolute value of) the pad size, the function has no effect.
`array_slice()`	Takes three arguments: an input array, an integer offset, and an (optional) integer length. Returns a new array that is a "slice" of the old one — a subsequence of its list of key/value pairs. The starting and stopping points of the slice are determined by the offset and length. A positive offset means that the starting point is that number of elements after the beginning; a negative offset means that it is that many elements before the end. The optional length argument specifies how long the resulting slice is (if positive) or how many elements before the end it should stop (if negative). If the length argument is not present, the slice continues to the end of the array.
`array_splice()`	Removes a chunk (or a slice) of an array and replaces it with the contents of another array. Takes four arguments: an input array, an offset, an optional integer length, and an optional replacement array. Returns a new array containing the slice that was removed from the input array. The rules for using the offset and length arguments to determine the slice that is removed are the same as in the previous `array_slice()` function. If no replacement array is supplied, this function simply (destructively) removes a slice of the input array and returns it. If there is a replacement array, the elements of that array are inserted in place of the removed slice.

Stacks and Queues

Stacks and *queues* are abstract data structures, frequently used in computer science, that enforce a certain kind of access discipline on the objects they contain, without necessarily committing to what those objects are. PHP arrays are well suited to imitating other kinds of data structures, and the loose typing of PHP array elements makes it easy for them to imitate stacks and queues. PHP provides some array functions specifically for this purpose — if you use them exclusively, you can forget that arrays are involved at all.

A *stack* is a container that stores values and supports last-in–first-out (LIFO) behavior. This means that the stack maintains an order on the values you store, and the only way you can get a value back is by retrieving (and removing) the most recently stored value. The usual analogy is a stack of cafeteria trays in one of those dispensers that keeps the top tray at a constant level. You can push new trays down on top of the old ones, and you can take trays off the top, but you can't grab an older tray without taking the newer ones first. The act of adding into the stack is called *pushing* a value onto

363

the stack, whereas the act of taking off the top is called *popping* the stack. Another analogy is the way some web browsers store the pages you have visited for use by the Back button; visiting a new page pushes a new URL onto that stack, and using the Back button pops the stack.

A *queue* is similar to a stack, but its behavior is first in, first out (FIFO). The usual analogy here is what the British call a *queue* and what Americans call a *line*, where people line up in order to wait for something. The rule is that whoever has been in the queue the longest is the next to be served.

The stack functions are `array_push()` and `array_pop()`. The `array_push()` function takes an initial array argument and then any number of elements to push onto the stack. The elements will be inserted at the end of the array, in order from left to right. The `array_pop()` function takes such an array and removes the element at the end, returning it. Take the following fragment:

```
$my_stack = array();  // needed--array_push() will not create
array_push($my_stack, "the first", "the middle");
array_push($my_stack, "the last");
while ($popped = array_pop($my_stack))
  print("Popped the stack and got: $popped<BR>");
```

This will produce the browser output:

```
Popped the stack and got: the last
Popped the stack and got: the middle
Popped the stack and got: the first
```

PHP also offers functions that behave exactly the same way as `array_push()` and `array_pop()`, except that they work at the other end, adding to and removing from the beginning of the array. The `array_unshift()` function is analogous to `array_push()`, and `array_shift()` is like `array_pop()`. If you choose one function from column A and one from column B, you can get the behavior of a queue. For example, we can rewrite our previous example to push into the beginning of the array (using `array_unshift()`) and pop from the end (using `array_pop()`, as before):

```
$my_queue = array();// needed--array_unshift() will not create
array_unshift($my_queue, "the first");
array_unshift($my_queue,"the middle");
array_unshift($my_queue, "the last");
while ($popped = array_pop($my_queue))
  print("Popped the queue and got: $popped<BR>");
```

It produces the output:

```
Popped the queue and got: the first
Popped the queue and got: the middle
Popped the queue and got: the last
```

CAUTION The `array_unshift()` and `array_shift()` functions are somewhat different from `array_push()` and `array_pop()` in that the former do some renumbering of the array indices if the indices are integers. The idea is that some people may be relying on the numerical indices to order the array contents, so using `array_unshift()` to insert a new element at the beginning should assign an index of 0 to the new element, and renumber those above. Similarly, popping an element from the beginning with `array_shift()` causes integral indices of other elements to be reduced. (This is not an issue with `array_push` and `array_pop`, because changes are at the end, and no renumbering is needed.) If you are using string indices exclusively, this renumbering has no effect. This is a general pattern with PHP array functions: Some of them treat integer indices like any other associative indexes, whereas others assume that integers imply order, and redo them if the order has changed.

The stack and queue functions are summarized in Table 21-2.

TABLE 21-2

Stack and Queue Functions

Function	Arguments	Side Effect	Returns
`array_ push()`	An initial array argument, then any number of values to be pushed onto the stack.	Modifies the array by adding the elements in order to the end of the array.	Returns the number of elements in the array after the push.
`array_pop()`	A single array argument.	Removes the element at the end of the array.	Returns the last (removed) value, or a false value if the array is empty.
`array_ unshift()`	An initial array argument, then any number of values to be pushed onto the front of the array.	Modifies the array by adding the successive elements to the beginning. (The last argument will be at the beginning of the array.)	Returns the number of elements in the array after the new elements are added.
`array_ shift()`	A single array argument.	Removes the element at the beginning of the array.	Returns the first (removed) value or a false value if the array is empty.

Translating between Variables and Arrays

PHP offers a couple of unusual functions for mapping between the name/value pairs of regular variable bindings and the key/value pairs of an array. The `compact()` function translates from variable bindings to an array, and the `extract()` function goes in the opposite direction. These are summarized briefly in Table 21-3.

TABLE 21-3

Array/Variable-Binding Functions

Function	Behavior
compact()	Takes a specified set of strings, looks up bound variables (if any) in the current environment that are named by those strings, and returns an array where the keys are the variable names, and the values are the corresponding values of those variables.
	This function takes any number of arguments, each of which is either a string or an array that contains strings at some level of index depth. The entire set of strings that are included in the argument(s) is used as the candidate set of variable names. Strings that do not correspond to bound variables are ignored.
extract()	Takes an array (plus two optional arguments explained in the next paragraph) and imports the key/value pairs into the current variable-binding context. The array keys become the variable names, and the corresponding array values become the values of the variables. Any keys that do not correspond to a legal variable name will not produce an assignment.
	The optional arguments are an integer (intended to receive one of a small set of constants) and a prefix string. The point of these arguments is to specify what should happen in the case of a collision between the name of an existing variable and one that would be created from an array key.
	The intended possible constants for the optional integer arguments include (1) EXTR_OVERWRITE, (2) EXTR_SKIP, (3) EXTR_PREFIX_SAME, and (4) EXTR_PREFIX_ALL. The corresponding behaviors are (1) go ahead and overwrite existing variables, (2) skip any new assignments that would require overwriting, (3) use the optional prefix string to distinguish the new variable from the old one, or (4) prefix all the new variables with the string. For example, extract(array('my_var' => 4), EXTR_PREFIX_SAME, 'diff_'); would cause $my_var to be 4 if $my_var were not already bound; otherwise, it would assign the value 4 to $diff_my_var. Other constants exist, though are less commonly used. See http://php.net/extract for more information.

Sorting

Finally, PHP offers a host of functions for sorting arrays. As you saw earlier, a tension sometimes arises between respecting the key/value associations in an array and treating numerical keys as ordering info that should be changed when the order changes. Luckily, PHP offers variants of the sorting functions for each of these behaviors and also allows sorting in ascending or descending order and by user-supplied ordering functions. The function names are terse, but each letter (other than the sort part) has its meaning. The decoder ring is something like:

- An initial a means that the function sorts by value but maintains the association between key/value pairs the way it was.

- An initial k means that it sorts by key but maintains the key/value associations.
- A lack of that initial a or k means that it sorts by value but doesn't maintain the key/value association. In particular, numerical keys will be renumbered to reflect the new ordering.
- An r before the sort means that the sorting order will be reversed.
- An initial u means that a second argument is expected: the name of a user-defined function that specifies the ordering of any two elements that are being sorted. (See the description in Table 21-4.)

TABLE 21-4

Array Sorting Functions

Function	Behavior
asort()	Takes a single array argument. Sorts the key/value pairs by value but keeps the key/value mapping the same. Good for associative arrays.
arsort()	Same as asort(), but sorts in descending order.
ksort()	Takes a single array argument. Sorts the key/value pairs by key but maintain the key/value associations the same.
krsort()	Same as ksort(), but sorts in descending order.
sort()	Takes a single array argument. Sorts the key/value pairs of an array by their values. Keys may be renumbered to reflect the new ordering of the values.
rsort()	Same as sort(), but sorts in descending order.
uasort()	Sorts key/value pairs by value using a comparison function. Similar to asort(), except the actual ordering of the values is determined by the second argument, which is the name of a user-defined ordering function. That function should return a negative number if its first argument is before the second (according to the comparison function), a positive number if the first argument comes after the second, and zero if the elements are the same.
uksort()	Sorts key/value pairs by key, using a comparison function. Similar to uasort(), except that the ordering is by key, rather than by value.
usort()	Sorts an array by value using a supplied comparison function. Similar to uasort(), except that (as in sort()), the key/value associations are not maintained.

Printing Functions for Visualizing Arrays

Before we leave this subject entirely, we should mention a couple of printing functions that are very useful for visualizing and debugging arrays, especially multidimensional arrays.

The first function is print_r(), which is short for *print recursive*. This takes an argument of any type and prints it out, which includes printing all its parts recursively. For a simple value (a number or string), this means simply that the value is printed; for compound types like arrays and objects it means that all elements (and all parts of those elements) are printed. The layout that makes the compound structure clear involves spaces, so it's best to wrap its output in an HTML <pre></pre> construct so that the spaces are printed literally.

CROSS-REF For more detail on the var_dump **function and other ways to visualize data structures, see Chapter 31 on debugging.**

The var_dump() function is similar, except that it prints additional information about the size and type of the values it discovers. An example is worth a thousand words here, so we will create a simple multidimensional array and print it using both functions:

```php
<?php

$my_array = array("key1" => "value1",
                  "key2" => array("subkey1" => "value2"));

print("The result of print_r:<BR><pre>");
print_r($my_array);
print("</pre><BR>");
print("The result of var_dump:<BR><pre>");
var_dump($my_array);
print("</pre><BR>");
?>
```

The resulting output from this sample looks like this:

```
The result of print_r:
Array
(
    [key1] => value1
    [key2] => Array
        (
            [subkey1] => value2
        )

)

The result of var_dump:
array(2) {
  ["key1"]=>
  string(6) "value1"
  ["key2"]=>
  array(1) {
    ["subkey1"]=>
    string(6) "value2"
  }
}
?>
```

Summary

The transformation functions are designed to do interesting things to your arrays. With the exception of `shuffle()`, these functions return their results as a newly created array. To treat an array as a stack is to give it a last-in–first-out property. You can treat an array as a stack by using the `array_push()` and `array_pop()` functions in tandem. Alternatively, `array_unshift()` and `array shift()` used in tandem will have a similar effect, though they work on the opposite end of the array. By choosing one function from each pair, you can effectively cause an array to act like a queue.

The `compact()` function maps variable names and values onto array keys and values, while `extract()` reverses the process, even if the array was not created with compact. Finally, a variety of functions in two major classes will sort and reorder arrays. The first major class will do it without reordering integral keys; the second will reorder your integral keys according to the new sorted order.

Chapter 22

Examining Regular Expressions

In Chapter 7 we covered PHP strings — how to create them, print them, and (to some extent) how to examine and modify them. In this chapter, we delve into more advanced string-manipulation techniques, starting off with functions to split up (or tokenize) strings into parts. We'll soon run into limitations of the basic tokenization functions, which show the need for regular expressions.

Finally, we'll cover some of the more advanced string functions that enhance the effectiveness of regular expressions and the use of strings in general.

Tokenizing and Parsing Functions

Sometimes you need to take strings apart at the seams, and you have your own notions of what should count as a seam. The process of breaking up a long string into *words* is called *tokenizing,* and among other things it is part of the internals of interpreting or compiling any computer program, including PHP. PHP offers a special function for this purpose, called `strtok()`.

The `strtok()` function takes two arguments: the string to be broken up into tokens and a string containing all the *delimiters* (characters that count as boundaries between tokens). On the first call, both arguments are used, and the string value returned is the first token. To retrieve subsequent tokens, make the same call, but omit the source string argument. It will be remembered as the current string, and the function will remember where it left off. For example:

```
$token = strtok(
    "open-source HTML-embedded server-side Web
scripting",
    " ");
```

371

```
while($token){
  print($token . "<BR>");
  $token = strtok(" ");
}
```

produces the browser output:

```
open-source
HTML-embedded
server-side
Web
scripting
```

The original string would be *broken* at each space. At our discretion, we could change the delimiter set, like this:

```
$token = strtok(
    "open-source HTML-embedded server-side Web scripting",
    "-");
while($token){
  print($token . "<BR>");
  $token = strtok("-");
}
```

This gives us (less sensibly):

```
Open
source HTML
embedded server
side Web scripting
```

Finally, we can break the string at all these places at once by giving it a delimiter string like " -", containing both a space and a dash. The code:

```
$token = strtok(
    "open-source HTML-embedded server-side Web scripting",
    " -");
while($token){
  print($token . "<BR>");
  $token = strtok(" -");
}
```

prints this output:

```
open
source
HTML
embedded
server
side
Web
scripting
```

Notice that in every case the delimiter characters do not show up anywhere in the retrieved tokens.

The `strtok()` function doles out its tokens one by one. You can also use the `explode()` function to do something similar, except that it stores the tokens all at once in an array. After the tokens are in the array, you can do anything you like with them, including sort them.

The `explode()` function takes two arguments: a separator string and the string to be separated. It returns an array where each element is a substring between instances of the separator in the string to be separated. For example:

```
$explode_result = explode("AND", "one AND a two AND a three");
```

results in the array `$explode_result` having three elements, each of which is a string: `"one "`, `" a two "`, and `" a three"`. In this particular example, there would be no capital letters anywhere in the strings contained in the array, because the `AND` separator does not show up in the result.

The separator string in `explode()` is significantly different from the delimiter string used in `strtok()`. The separator is a full-fledged string, and all its characters must be found in the right order for an instance of the separator to be detected. The delimiter string of `strtok()` specifies a set of single characters, any one of which will count as a delimiter. This makes `explode()` both more precise and more brittle — if you leave out a space or a newline character from a long string, the entire function will be broken.

Because the entire separator string disappears into the ether when `explode()` is used, this function can be the basis for many useful effects. The examples given in most PHP documentation use short strings for convenience, but remember that a string can be almost any length — and `explode()` is especially useful with longer strings that might be tedious to parse some other way. For instance, you can use it to count how many times a particular string appears within a text file by turning the file into a string and using `explode()` on it, as in this example (which uses some functions we haven't explained yet, but we hope make sense in context).

```php
<?php
//First, turn a text file into a string called $filestring.
$filename = "complex_layout.html";
$fd = fopen($filename, "r");
$filestring = fread($fd, filesize($filename));
fclose ($fd);

//Explode on the beginning of the <TABLE> HTML tag
$tables = explode("<TABLE", $filestring); // assumes uppercase
//Count the number of pieces
$num_tables = count($tables);

//Subtract one to get the number of <TABLE> tags, and echo
echo ($num_tables - 1);
?>
```

The `explode()` function has an inverse function, `implode()`, which takes two arguments: a "glue" string (analogous to the separator string in `explode()`) and an array of strings like that returned by `explode()`. It returns a string created by inserting the glue string between each string element in the array.

You can use the two functions together to replace every instance of a particular string within a text file. Remember that the separator string will vanish into the ether when you perform an `explode()` — if you want it to appear in the final file, you have to replace it by hand. In this example, we're changing the font tags on a web page.

```php
<?php
//Turn text file into string
$filename = "someoldpage.html";
$fd = fopen($filename, "r");
$filestring = fread($fd, filesize($filename));
fclose ($fd);
$parts = explode("arial, sans-serif", $filestring);
$whole = implode("arial, verdana, sans-serif", $parts);

//Overwrite the original file
$fd = fopen($filename, "w");
fwrite($fd, $whole);
fclose ($fd);
?>
```

Why Regular Expressions?

The string-comparison and substring-finding functions we saw here and in Chapter 7 are fine as far as they go, but they are on the literal-minded side. As an example of their weakness, let's say that you want to test strings to see if they are a particular kind of web hostname: addresses that start with `www.` and end with `.com`, and have one lowercase alphabetic word in the middle. For example, these are strings we want:

```
'www.ibm.com'
'www.zend.com'
```

And the following are not:

```
'java.sun.com'
'www.java.sun.com'
'www.php.net'
'www.IBM.com'
'www.Web addresses can't have spaces.com'
```

With a little thought, it's obvious that there is no convenient way to simply use string and substring comparison to build the test that we want. We can test for the presence of `www.` and `.com`, but it is difficult to enforce what should be happening between them. This is what regular expressions are good for.

Regex in PHP

Regular expressions (or *regex*, pronounced with a soft *g* by your authors, but with no consensus pronunciation) are patterns for string matching, with special wildcards that can match entire portions of the target string. There are two broad classes of regular expression that PHP works with: *POSIX* (extended) regex and *Perl-compatible* regex. The differences mostly have to do with syntax, although there are some functional differences, too.

POSIX-style regular expressions are ultimately descended from the regex pattern-matching machinery used in Unix command-line shells; Perl-compatible regex is a more direct imitation of regular expressions in Perl. We've already waxed poetic about the utility of arrays. We're about to do it again with regex. If you're planning on doing any substantial coding in a web environment, sooner or later you will bump up against regex.

> **NOTE** Note that for PHP6, the ereg functions are no longer included.

An example of POSIX-style regex

Here are a few of the rules for POSIX-style regular expressions, simplified:

■ Characters that are not *special* are matched literally. The letter a in a pattern, for example, matches the same letter in a target string.

■ The special character ^ matches the beginning of a string only, and the special character $ matches the end of a string only.

■ The special character . matches any character.

■ The special character * matches zero or more instances of the previous regular expression, and + matches one or more instances of the previous expression.

■ A set of characters enclosed in square brackets matches any of those characters — the pattern [ab] matches either a or b. You can also specify a range of characters in brackets by using a hyphen — the pattern [a-c] matches a, b, or c.

■ Special characters that are escaped with a backslash (\) lose their special meaning and are matched literally.

We can use the preceding rules to construct an expression that matches the kind of web address we want in the section "Why Regular Expressions?" earlier in this chapter. Our chosen expression is:

```
^www\.[a-z]+\.com$
```

In this expression we have the '^' symbol, which says that the www portion must start at the beginning of the string. Then comes a dot (.), preceded by a backslash that says we really want a dot, not the special . wildcard character. Then we have a bracket-enclosed range of all the lowercase alphabetic letters. The following + indicates that we are willing to match any number of these lowercase letters in a row, as long as we have at least one of them. Then another literal ., the com, and the special $ that says that com is the end of it.

Now let's use that expression as an argument to the function `ereg()`, which takes as arguments a pattern string and a string to match against. We can use an `ereg()` call to build a test function for our kind of web address.

```
function simple_dot_com ($url)
{
  return(ereg('^www\\.[a-z]+\\.com$', $url));
}
```

Confusingly, we have to put two backslashes in the pattern string, because PHP treats the first slash as an escape character for the second backslash. (You can get away with just one backslash, but that behavior is not guaranteed to continue in future versions of PHP.) The second backslash (escaped by the first), in turn, is a regex escape character for the following character.

This function will return TRUE or FALSE, depending on whether it successfully matches our pattern. Now we can use our function to test some of the addresses listed earlier.

```
$urls_to_test =
   array('www.ibm.com', 'www.java.sun.com',
         'www.zend.com', 'java.sun.com',
         'www.java.sun.com', 'www.php.net',
         'www.IBM.com',
         'www.Web addresses can\'t have spaces.com');
while($test = array_pop($urls_to_test)){
  if (simple_dot_com($test))
    print("\"$test\" is a simple dot-com<BR>");
  else
    print("\"$test\" is NOT a simple dot-com<BR>");

}
```

The results of our tests are:

```
"www.Web addresses can't have spaces.com" is NOT a simple dot-com
"www.IBM.com" is NOT a simple dot-com
"www.php.net" is NOT a simple dot-com
"www.java.sun.com" is NOT a simple dot-com
"java.sun.com" is NOT a simple dot-com
"www.zend.com" is a simple dot-com
"www.java.sun.com" is NOT a simple dot-com
"www.ibm.com" is a simple dot-com
```

This is the kind of discriminating behavior we are looking for.

TIP   **On many Unix systems, typing** `man 7 regex` **will lead you to a guide to POSIX regular expressions. If that does not work, try** `man regex` **and follow any pointers to related pages.**

Regular expression functions

The POSIX-style regular expression functions in PHP are summarized in Table 22-1. These are included for legacy applications where you might find them still being used. These functions are no longer in PHP6 and have been replaced with preg functions, discussed later in this chapter.

> **TIP** If you find yourself using a regular expression function with a pattern that has no special characters, you are probably using an expensive tool where a cheap one would do. If you are trying to match a simple string to a simple string, you need only one of the more basic (and faster) functions that we cover earlier in this chapter and in Chapter 7.

TABLE 22-1

POSIX Regular Expression Functions

Function	Behavior
ereg()	Takes two string arguments and an optional third-array argument. The first string is the POSIX-style regular expression pattern, and the second string is the target string that is being matched. The function returns TRUE if the match was successful and FALSE otherwise. In addition, if an array argument is supplied and portions of the pattern are enclosed in parentheses, the parts of the target string that match successive parenthesized portions will be copied into successive elements of the array.
ereg_replace()	Takes three arguments: a POSIX regular expression pattern, a string to do replacement with, and a string to replace into. The function scans the third argument for portions that match the pattern and replaces them with the second argument. The modified string is returned.
	If there are parenthesized portions of the pattern (as with ereg()), the replacement string may contain special substrings of the form \\digit (that is, two backslashes followed by a single-digit number), which will themselves be replaced with the corresponding piece of the target string.
eregi()	Identical to ereg(), except that letters in regular expressions are matched in a case-independent way.
eregi_replace()	Identical to ereg_replace(), except that letters in regular expressions are matched in a case-independent way.
split()	Takes a pattern, a target string, and an optional limit on the number of portions to split the string into. Returns an array of strings created by splitting the target string into chunks delimited by substrings that match the regular expression. (Note that this is analogous to the explode() function, except that it splits on regular expressions rather than literal strings.)
spliti()	Case-independent version of split().

Perl-Compatible Regular Expressions

Perl-compatible regex in PHP has a completely distinct set of functions and a slightly different set of rules for patterns.

Perl-compatible regex patterns are always bookended by one particular character, which must be the same at beginning and end, indicating the beginning and end of the pattern. By convention, this is most often the / character, although you can use a different character if you so desire. The Perl-compatible pattern:

```
/pattern/
```

matches any string that has the string (or substring) pattern in it. To make things slightly more complicated, these patterns are typically strings, and PHP needs its own quotes to recognize such strings. So if you are putting a pattern into a variable for later use, you might well do this:

```
$my_pattern = '/pattern/';
```

This variable would now be suitable for passing off to a Perl-compatible regex function that expects a pattern as argument.

Although we don't have time or space to cover Perl-compatible regex patterns in detail, Table 22-2 shows a list of the most commonly used constructs.

TABLE 22-2

Common Perl-Compatible Pattern Constructs

Construct	Interpretation
Simple literal character matches	If the character involved is not special, Perl will match characters in sequence. The example pattern /abc/ matches any string that has the substring 'abc' in it.
Character class matches: [<*list of characters*>]	Will match a single instance of any of the characters between the brackets. For example, /[xyz]/ matches a single character, as long as that character is either x, y, or z. A sequence of characters (in ASCII order) is indicated by a hyphen, so that a class matching all digits is [0-9].
Predefined character class abbreviations	The patterns \d will match a single digit (from the character class [0-9]), and the pattern \s matches any whitespace character.
Multiplier patterns	Any pattern followed by * means: "Match this pattern 0 or more times."
	Any pattern followed by ? means: "Match this pattern exactly once."
	Any pattern followed by + means: "Match this pattern 1 or more times."

Construct	Interpretation
Anchoring characters	The caret character ^ at the beginning of a pattern means that the pattern must start at the beginning of the string; the $ character at the end of a pattern means that the pattern must end at the end of the string. The caret character at the beginning of a character class [^abc] means that the set is the complement of the characters listed (that is, any character that is not in the list).
Escape character '\'	Any character that has a special meaning to regex can be treated as a simple matching character by preceding it with a backslash. The special characters that might need this treatment are: `. \ + * ? [ ] ^ $ ( ) { } = ! < >	:`
Parentheses	A parenthesis grouping around a portion of any pattern means: "Add the substring that matches this pattern to the list of substring matches."

Take, as an example, the following pattern:

```
/phone number\s+(\d\d\d\d\d\d\d)/
```

It matches any string that contains the literal phrase phone number, followed by some number of spaces (but at least one), followed by exactly seven digits (no spaces, no dash). In addition, because of the parentheses, the seven-digit number is saved and returned in an array containing substring matches if it is called from a function that returns such things.

The Perl-compatible functions are summarized in Table 22-3.

The most widely used of these functions are probably preg_match() and preg_match_all(). The first is best for simply answering whether a pattern matches a string, and the latter is best for either counting matches or collecting portions that match.

The optional fourth argument to preg_match_all() requires a little more explanation. The array that contains the returned matches is going to be two levels deep, with one level being the iteration of the match (the first match, the second, and so on) and the other level being the position of the match in the pattern. (The entire match is always first, followed by any parenthesized subpatterns in order.) The question is: Which level is on top? Will the array be a list of positions, each of which contains a list of iterations, or the other way around? If the argument is PREG_PATTERN_ORDER, the first element will contain all matches of the entire pattern, the second element will contain all matches of the first parenthesized pattern, and so forth. If the argument is PREG_SET_ORDER, the first argument will be all the substrings from the first match (first the total match, then parenthesized bits in order), the second element will contain all the substrings from the second match, and so on. (See the following example to clarify.)

TABLE 22-3

Perl-Compatible Regular Expression Functions

Function	Behavior
preg_match()	Takes a regex pattern as first argument, a string to match against as second argument, and an optional array variable for returned matches. Returns 0 if no matches are found, and 1 if a match is found. If a match is successful, the array variable contains the entire matching substring as its first element, and subsequent elements contain portions matching parenthesized portions of the pattern. As of PHP 4.3.0, an optional flag of PREG_OFFSET_CAPTURE is also available. This flag causes preg match to return into the specified array a two-element array for each match, consisting of the match itself and the offset where the match occurs.
preg_match_all()	Like preg_match(), except that it makes all possible successive matches of the pattern in the string, rather than just the first. The return value is the number of matches successfully made. The array of matches is not optional (If you want a true/false answer, use preg_match()).
	The structure of the array returned depends on the optional fourth argument (either the constant PREG_PATTERN_ORDER, or PREG_SET_ORDER, defaulting to the former). (See further discussion following the table.) PREG_OFFSET_ CAPTURE is also available with this function.
preg_split()	Takes a pattern as first argument and a string to match as second argument. Returns an array containing the string divided into substrings, split along boundary strings matching the pattern. (Analogous to the POSIX-style function split().) An optional third argument (limit) controls how many elements to split before returning the list; -1 means no limit. An optional flag in the fourth position can be PREG_SPLIT_NO_EMPTY causing the function to return only nonempty pieces, PREG_SPLIT_DELIM_CAPTURE causing any parenthesized expression in the delimiter pattern to be returned, or PREG_SPLIT_OFFSET_ CAPTURE, which does the same as PREG_OFFSET_CAPTURE.
preg_replace()	Takes a pattern, a replacement string, and a string to modify. Returns the result of replacing every matching portion of the modifiable string with the replacement string. An optional limit argument determines how many replacements will occur (as in preg_split()).
preg_replace_ callback()	Like preg_replace(), except that the second argument is the name of a callback function, rather than a replacement string. This function should return the string that is to be used as a replacement.
preg_grep()	Takes a pattern and an array and returns an array of the elements of the input array that matched the pattern. Surviving values of the new array have the same keys as in the input array.
preg_quote()	A special-purpose function for inserting escape characters into strings that are intended for use as regex patterns. The only required argument is a string to escape; the return value is that string with every special regex character preceded by a backslash.

Example: A simple link-scraper

As an example of what regex can do for us, let's write a simple function to grab and print links from an arbitrary web page. The input will be a URL for the page we're interested in analyzing; the output will be a printed list of the links on the page, split into the target URL for the link and the descriptive text that appears in the link (the *anchortext*). We will do this using Perl-compatible regex functions.

Such a function might be the very first step in writing a web crawler for a search engine. Search engines download the contents of web pages to analyze and index them, but they also need to discover links to other pages, if only to discover new content.

The regular expression

The heart of our little function will be the regular expression itself. What we need to do is design an expression that will match HTML links (and nothing else) and that is suitable for using to extract pieces of such links.

HTML links generally look something like this:

```
<A HREF="http://mysite.com/mypage.php">My cool page on my cool
site</A>
```

That is, an anchor tag that has an HREF attribute, and which encloses the anchortext between the start tag (<A>) and the end tag (</A>). We'll construct a pattern to match this simplified view of an anchortext element. (This won't capture everything that the HTML spec permits as legal anchor links — in particular, you are allowed attributes in anchors other than HREFs, but we will ignore that for our purposes.)

Now, regular expressions are famously unreadable when considered all at once. So we will grow this one in several drafts as we explain what's going on.

First, let's start with a minimal expression to catch a beginning anchor tag. Our first draft looks like this:

```
/<A\sHREF="[^"]+">/
// first draft of a pattern to match anchor links
```

(Note that this is not yet intended to be working PHP code; we're drafting an expression that we'll plug into PHP code later.)

In English, our first-draft definition of an anchor tag is left angle bracket, followed by A, followed by a space, followed by the string HREF=, followed by a double-quotation mark, followed by any number of characters that are not quotation marks, followed by a closing quotation mark, followed by a right angle bracket. Then the whole expression is enclosed in a pair of slashes, indicating to the regex engine the start and end of the expression.

The [^"]+ construction in the middle of this expression breaks down like this: The brackets indicate a character set, and the caret (^) immediately after the left bracket indicates that we are negating the set — that the set contains every character that is *not* in the subsequent list. Finally, the + after that bracketed class means that we expect at least one nonquote character.

As we've said, we're not trying to capture the precise syntax prescribed by the HTML specification. But there are a couple of ways that we can make this expression less strict. For one thing, as far as we know, there may be spaces between the initial < character and the A tag. Similarly, there may be an arbitrary number of spaces between the A and the HREF or the closing double-quote and the right angle bracket. Adding these, the expression becomes:

```
/<A\s+HREF="[^"]+"\s*>/
// second draft, allowing more spaces
```

Here, \s+ means one or more spaces.

Now we add the anchortext itself and the closing </A> tag:

```
/<A\s+HREF="[^\"]+"\s*>[^>]*<\/A>/
 // third draft, with text and close tag
```

We are allowing the anchortext to be anything up until a closing anchor tag, so we make an anything-but-right-angle-bracket character class ([^>]) and indicate that it can repeat zero or more times. Finally, we add the subpattern to match the closing anchor tag (<\/A>).

This is fine as far as it goes, but it will only match anchors where the tag name (A) and attribute (HREF) are in uppercase. Lowercase tags should be legal as well, so we add an i modifier after the entire expression, to specify case-independent matching.

```
/<A\s+HREF="[^\"]+"\s*>[^>]*<\/A>/i
// fourth draft, case-independent
```

This draft is nearly final and could be used to give true/false answers to the question of whether a page contains the kind of links we like. But we want to go further and extract certain portions of any string that does match. We signify this by adding parentheses to enclose the portions we're interested in:

```
/<A\s+HREF="([^\"])+"\s*>([^>]*)<\/A>/i
// final draft, extracts portions
```

They may be hard to see by this point, but we've added a pair of parentheses to enclose the target of the HREF (between the quotes) and another pair around the anchortext area (between the tags). These parentheses tell the calling function to save the string portion that matches the enclosed area, so that it can be added to the return array.

Using the expression in a function

With an anchor-tag-matching expression in hand, our goal now is to write a function to scrape links from an HTML page. We'll need to:

- Take a URL as argument
- Open up an HTTP connection to the URL and grab its contents as a string
- Iterate through the string, applying our regex pattern wherever we can, saving what matches
- Print the extracted portions (target URL and anchortext)

Such a function is shown in Listing 22-1.

LISTING 22-1

A print_links function

```php
<?php

function print_links ($url)
{
  $fp = fopen($url, "r")
    or die("Could not contact $url");
  $page_contents = "";
  while ($new_text = fread($fp, 100)) {
    $page_contents .= $new_text;
  }
  $match_result =
    preg_match_all('/<A\s+HREF="([^\"]+)"\s*>([^>]*)<\/A>/i',
               $page_contents,
               $match_array,
               PREG_SET_ORDER);

  foreach ($match_array as $entry) {
    $href = $entry[1];
    $anchortext = $entry[2];
    print("<B>HREF</B>: $href;
          <B>ANCHORTEXT</B>:  $anchortext<BR>");
  }
}

?>
```

This function is easier to write than you might expect because PHP takes care of several parts of it for us. We do not need to write anything special to make an HTTP connection to download a web page because fopen() will accept a URL as argument and do the right thing. All we need to do after

calling `fopen()` on the URL is to read characters until we are out of them, appending what we get onto a constructed string.

The iteration through the HTML page's contents is taken care of by `preg_match_all()`, which applies the regex pattern as many times as possible, starting from the previous match each time, and saving the matches in `$match_array`. We chose to have the array arranged by `PREG_SET_ORDER`, meaning that each entry in the top-level array is the portion from a particular match in the iteration, rather than across matches.

Applying the function

The only argument the function requires is a URL. In testing the function before including it in the book, we pointed it at link-rich, top-level pages like `http://slashdot.org`, `www.cnn.com`, and `www.php.net`. Those results would be fun to display, but all of those sites have copyright notices, and publishers are understandably wary of allowing authors to put other people's copyrighted material into their copyrighted book without permission. So, instead, we pointed it at the top-level place-holder page for our own vanity site (`www.troutworks.com`), like this:

```
print_links("http://www.troutworks.com/");
```

You get the following result (approximately):

```
HREF: http://www.mysteryguide.com; ANCHORTEXT: MysteryGuide
HREF: http://www.sciencebookguide.com; ANCHORTEXT:
ScienceBookGuide
HREF: /Joycelog/joycelog.php; ANCHORTEXT: Troutgirl weblog
HREF: /Timlog/timlog.php; ANCHORTEXT: Timboy weblog
HREF: http://www.troutworks.com/phpbook; ANCHORTEXT: code
download site
HREF: http://www.amazon.com/exec/obidos/tg/detail/-/0764549553/;
HREF: http://www.mysteryguide.com; ANCHORTEXT: MysteryGuide
HREF: http://www.sciencebookguide.com; ANCHORTEXT:
ScienceBookGuide
ANCHORTEXT: PHP Bible
HREF: http://www.troutworks.com/phpbook; ANCHORTEXT: code
download site
```

Just because we didn't feel that we could print the results of the links from those more interesting sites doesn't mean that you can't apply this code to them (however, see the warnings in the sidebar "Writing Well-Behaved Spiders").

Extending the code

As we've said, code like Listing 22-1 is the very beginning of writing a web search spider. If you want to make it more real, you could:

- Convert the relative links to absolute (`http://`) links by remembering the URL that you are scraping and splicing that base URL appropriately with the relative path

- Add a more graceful way to bounce back from an unreachable site rather than immediately dying

- Expand the regex pattern to match HREFs that have quotation marks around the URL as well as HREFs that do not

- Add capability for recursive calls so that, rather than simply printing a child link, you apply the same function again to it and explore its own links

Writing Well-Behaved Spiders

A note of caution, however (informed by the experience of one of your authors in the search engine business). There are two rules that you should observe, though, before writing any kind of spider that does more automated crawling. When you crawl any site, you should:

- Check to see if there is a `robots.txt` file (at `http://sitename/robots.txt`). If there is no such file, the site owners are implicitly saying the site is okay to crawl. If there is such a file, you should either not crawl the site or, if you do, you should make sure that you are not crawling pages that match the patterns laid out in that file. (For more on this, do a web search for "robot exclusion standard".)

- Make sure that you don't request files from any particular site too frequently. A decent interval to wait between requests is 10 seconds or so. (You can implement this delay on a per-site basis, or simply by sleeping for 10 seconds between every request.) It is not OK to simply create a recursive version of the preceding code and then unleash it on a large site, grabbing new links and pages as fast as your code can loop. Remember: One man's search engine is another's denial-of-service attack.

Advanced String Functions

We have now covered the most basic things to do with strings, as well some more sophisticated means of working with them via regular expressions. Now, we'll delve into some more exotic string functions, which we've categorized by type and/or purpose. These are the sort of functions that might only be relevant to you if you're working on a particular kind of project. Some of these sections might make you want to say, "Why would anyone want to do that?" If so, please ignore them until you the day that you suddenly realize that you need to do that thing exactly.

HTML functions

PHP offers a number of web-specific functions for string manipulation, which are summarized in Table 22-4.

TABLE 22-4

HTML-Specific String Functions

Function	Behavior
htmlspecialchars()	Takes a string as argument and returns the string with replacements for four characters that have special meaning in HTML. Each of these characters is replaced with the corresponding HTML entity, so that it will look like the original when rendered by a browser. The & character is replaced by & "" (the double-quote character) is replaced by "; < is replaced by <; > is replaced by >.
htmlentities()	Goes further than htmlspecialchars(), in that it replaces all characters that have a corresponding HTML entity with that HTML entity.
get_html_ translation_ table()	Takes one of two special constants (HTML_SPECIAL_CHARS and HTML_ ENTITIES), and returns the translation table used by htmlspecialchars() and htmlentities(), respectively. The translation table is an array where keys are the character strings and the corresponding values are their replacements.
nl2br()	Takes a string as argument and returns that string with inserted before all new lines (\n, \r or \r\n). This is helpful, for example, in maintaining the apparent line length of text paragraphs when they are displayed in a browser.
strip_tags()	Takes a string as argument and does its best to return that string stripped of all HTML tags and all PHP tags.

Hashing using MD5

MD5 is a string-processing algorithm that is used to produce a digest or signature of whatever string it is given. The algorithm boils its input string down into a fixed-length string of 32 hexadecimal values (0,1,2, . . . 9,a,b, . . . f). MD5 has some very useful properties:

■ MD5 always produces the same output string for any given input string, so it is not appropriate to use MD5 to store passwords.

■ The fixed-length results of applying MD5 are very evenly spread over the range of possible values.

■ It may be possible produce an input string corresponding to a given MD5 output string or to produce two inputs that yield the same output.

PHP's implementation of MD5 is available in the function md5(), which takes a string as input and produces the 32-character digest as output. For example, evaluating this:

```
print("md5 of 'Tim' is " . md5('Tim') . "<BR>");
print("md5 of 'tim' is " . md5('tim') . "<BR>");
print("md5 of 'time' is " . md5('time') . "<BR>");
```

gives us the browser output:

```
md5 of Tim is dc2054afd537ddc98afd9347136494ac
md5 of tim is b15d47e99831ee63e3f47cf3d4478e9a
md5 of time is 07cc694b9b3fc636710fa08b6922c42b
```

Although the input strings seem close to each other in some sense, there is no apparent similarity in the output strings. And since the range of possible output values is so huge (16 to the 32nd power), the chances that any two distinct strings will *collide* by producing the same MD5 value is vanishingly small.

The characteristics of MD5 make it useful for a wide variety of tasks, including:

- **Checksumming a message or file**: If you are worried about errors that might happen in transfer, you can transmit an MD5 digest, along with the message, and run the message through MD5 again after transfer. If the two versions of the digest do not match, then something is amiss.

- **Detecting if a file's contents have changed**: Similar to checksumming, MD5 is often used in this way by search engines as a check on whether a web page has changed, making re-indexing necessary. It is cheaper to store the MD5 digest than the entire original file.

- **Splitting strings or files into buckets**: If you want to divide a set of strings into *N* randomly dispersed sets, you can MD5 the strings, take the first few hex characters, translate them into a number, and take that number modulo the number of bins you want.

In addition to the md5() function, PHP offers md5_file(), which takes a filename as argument and returns an MD5 hash of the file's contents.

Strings as character collections

PHP offers some pretty specialized functions that treat strings more as collections of characters than as sequences.

The first is strspn(), which you can use to see what portion of a string is composed only of a given set of characters. For example:

```
$twister = "Peter Piper picked a peck of pickled peppers";
$charset = "Peter picked a";
print("The segment matching '$charset' is " .
    strspn($twister, $charset) . " characters long");
```

gives us:

```
The segment matching 'Peter picked a' is 26 characters long
```

because the first character not found in $charset is the o in of, and there are 26 characters that precede it.

The `strcspn()` function (where that internal `c` stands for *complement*) does the same thing, except that it accepts characters that are *not* in the character set argument. For example, the statement:

```
echo(strcspn($twister, "abcd"));
```

prints the number 14, because it accepts a 14-character sequence with the last character being the `c` in `picked`.

Finally, hark back to Chapter 8 on arrays and check out the following for an acute analysis of alliteration:

```
$twister = "Peter Piper picked a peck of pickled peppers";
print("$twister<BR>");
$letter_array = count_chars($twister, 1);
while ($cell = each($letter_array)){
  $letter = chr($cell['key']);
  $frequency = $cell['value'];
  print("Character: '$letter'; frequency:  $frequency<BR>");
}
```

This gives the browser output:

```
Peter Piper picked a peck of pickled peppers
Character: ' '; frequency: 7
Character: 'P'; frequency: 2
Character: 'a'; frequency: 1
Character: 'c'; frequency: 3
Character: 'd'; frequency: 2
Character: 'e'; frequency: 8
Character: 'f'; frequency: 1
Character: 'i'; frequency: 3
Character: 'k'; frequency: 3
Character: 'l'; frequency: 1
Character: 'o'; frequency: 1
Character: 'p'; frequency: 7
Character: 'r'; frequency: 3
Character: 's'; frequency: 1
Character: 't'; frequency: 1
```

The `count_chars()` function returns a report on the occurrences of characters in its string argument, packaged as an array where the keys are the ASCII values of characters, and the values are the frequencies of those characters in the string. The second argument to `count_chars()` is an integer that determines which of several modes the results should be returned in. In mode 0, an array of key/value pairs is returned, where the keys are every ASCII value from 0 to 255, and the corresponding values are the frequencies of each character in the string. Modes 1 and 2 are variants that include only ASCII values that occurred in the string (mode 1) or that did not occur (mode 2).

Finally, modes 3 and 4 return a string instead of an array, where the string contains all characters that occur (mode 3) or do not occur (mode 4).

These functions are summarized in Table 22-5.

CROSS-REF For an explanation of how to take apart array formats like that returned by `count_chars()`, see Chapter 8. The `chr()` function used in the preceding example, which maps from ASCII numbers to the corresponding characters, is covered in Chapter 5.

TABLE 22-5

Functions for Examining Character Contents

Function	Behavior
`count_chars()`	Takes a single string argument and an integer mode argument from 0 to 4. Returns a report about frequencies of characters in the string argument, as either an array or a string. (See the preceding text for more detail.)
`strspn()`	Takes two string arguments and returns the length of the initial substring of the first argument that is composed entirely of characters found in its second argument.
`strcspn()`	Takes two string arguments and returns the length of the initial substring of the first argument that is composed entirely of characters that are *not* found in its second argument.

String similarity functions

How similar is this string to that string? Well, it depends what you mean by *similar*, right?

If the kind of similarity you want is similarity of spelling, consider the Levenshtein metric. The `levenshtein()` function takes two strings and returns the minimum number of additions, deletions, and replacements of letters needed to transform one into the other. For example:

■ `levenshtein('Tim', 'Time')` returns 1.

■ `levenshtein('boy', 'chefboyardee')` returns 9.

■ `levenshtein('never', 'clever')` returns 2.

If the similarity you are interested in is phonetic, consider the functions `soundex()` and `metaphone()`. Both of them take an input string and return a key string representing the pronunciation category of the word (in English). If two input word strings map to exactly the same output key, they most likely have a similar pronunciation.

Summary

PHP has a wealth of built-in functions for handling strings — functions to create them, stick them together, chop them up, and do various kinds of analysis. The simplest of these were covered in Chapter 7, and in this chapter we saw functions for tokenizing, hashing, character counting, and determining similarity, as well as HTML-specific functions.

Simple string matching is all very well, but when you need industrial-strength pattern matching, nothing less than regex will do. PHP6 removes the ereg functions, preferring instead the preg functions for pattern matching.

Chapter 23

Working with the Filesystem

This chapter contains information on the multiplicity of system functions built into PHP. Many of these functions duplicate system functions via HTTP. Among the most useful are file-reading and -writing functions and those that return dates or times.

CAUTION Many of the functions in this chapter have serious security implications. You are inviting bad news if you use them without thinking pretty hard about the consequences! We'll try to point out the scariest ones as we go, but nothing that allows the system to be altered via HTTP should be undertaken lightly.

Some of these functions are Unix-only. The Windows system is deliberately made less available to users, especially to non-administrator users, and lacks many utilities that Unix-heads take for granted. If you're having problems and you run on Windows, make sure the function is enabled on your platform.

Understanding PHP File Permissions

Many PHP users, who have a developer orientation rather than any sysadmin experience, unfortunately do not take the time to understand Unix filesystem permissions. You really need to have a firm grasp of the basics to make good decisions about using many of the functions in this section. If you already do, feel free to skip the rest of this section.

Unfortunately, most explanations of the subject are quite general and user's eyes can easily glaze over in a hail of `rwx`es and three-digit numbers. So we're going to break it down for you into two simple default rules specifically for PHP users.

- Unless you have a good reason to do otherwise, the PHP files that you wish to make public should all be set to 644 (`rw-r--r--`).

- Unless you have a good reason to do otherwise, the PHP-enabled directories that you wish to make public should all be set to 751 (`rwxr-x--x`).

For some reason, many users seem to believe that PHP files need to be executable. This is only true for files that you write with the intention of their being called on the command line (for example, `./myscript.php`). Files that will be run through a web server only need to be readable by the web server user (usually Nobody, or some other user with very limited permissions). It's rather inconvenient to make the files not writable by you (and doesn't really matter if you own the parent directory), which is why our default recommendation is 644 (`rw-r--r--`) rather than 444 (`r--r--r--`), but this is a matter of convenience only.

Directory permissions are also very often misunderstood. Many users seem to believe that directories need to be readable for files to run. Actually the *read* directory permission means a user can list the contents of that directory (via the `ls` command, for instance). The *execute* directory permission is closer to what we think of as readable. For your PHP scripts to run, the directory needs only to be world-executable (751 or `rwxr-x--x`). Do not make the directory writable by others unless you know what you're doing.

File Reading and Writing Functions

This is a supremely useful set of functions, particularly for data sets too small or scattered to merit the use of a database. File reading is pretty safe unless you keep unencrypted passwords lying around, but file writing can be quite unsafe.

> **TIP** Remember that although the web server (and client-side languages such as JavaScript) can only act on files located under the document root, unless the `open_basedir` value or another chroot mechanism is set, PHP can access files at any location in the file system — including those above or entirely outside the web server document root — as long as the file permissions and `include_path` are set correctly. For instance, if your web server document root is located at `/usr/local/apache/htdocs`, Apache will be able to serve only files from this directory and its subdirectories, but PHP can open, read, and write to files in `/usr/local`, `/home/php`, `/export/home/httpd` or any other directory that you make readable and includable by the PHP and/or web server user.

A file manipulation session might involve the following steps:

1. Open the file for read/write.
2. Read in the file.
3. Close the file (may happen later).

4. Perform operations on the file contents.

5. Write the results out.

Each step has a corresponding PHP filesystem function.

This archetypal example illustrates some subtleties of the syntax for manipulating file contents:

```
$fd = fopen($filename, "r+")
  or die("Can't open file $filename");
$fstring = fread($fd, filesize($filename));
$fout = fwrite($fd, $fstring);
fclose($fd);
```

The effect of this particular example will be to *double* the file — in other words, the end result will be a file with the original contents of the file written out twice. This function will not overwrite the file, as you might expect. In the following sections, we walk you through this archetypal file manipulation session, step by step.

File open

It's essentially mandatory to assign the result of `fopen()` to a variable (traditionally `$fd` for *file descriptor*, or `$fp` for *file pointer*).

CAUTION Note that `fopen()` **does not return an integer on success. In fact, it returns a resource that says** `Resource id #n`**, where** `n` **is the number of the currently opened stream. Do not attempt to test the success of your file open by using** `is_int()` **or** `is_numeric()`**. Use** `die()` **instead.**

If it's successful in opening the file, PHP will return a resource ID, which it requires for further operations such as `fread` or `fwrite`. Otherwise, the value will be `false`.

CAUTION The system makes only a certain number of file descriptors available, which is a good argument for closing files as soon as you can. If you anticipate a large demand and have access to system settings, you may increase the number. However, if you fail to close a file descriptor, PHP will do it for you when the script ends.

Files may be opened in any of six modes (similar to permissions levels). If you try to do mode-inappropriate things, you will be denied. The modes are:

- Read-only (`"r"`).
- Read and write if the file exists already (`"r+"`): will write to the beginning of the file, doubling the original contents of the file if you read the file in as a string, edit it, and then write the string out to the file.
- Write-only (`"w"`) will create a file of this name, if one doesn't already exist, and will erase the contents of any file of this name before writing! You cannot use this mode to read a file, only to write one.

- Write and read even if the file doesn't exist already (`"w+"`) will create a file of this name, if one doesn't already exist, and will erase the contents of any file of this name before writing!

- Write-only to the end of a file whether it exists or not (`"a"`).

- Read and write to the end of a file whether it exists or not (`"a+"`), "doubling" original contents of the file if you read the file in as a string, edit it, and then write the string out to the file.

You need to be very sure you have read in the contents of any preexisting file before using w or w+ on it. Your chances of losing data with the other modes is much less.

There are several types of file connections that can be opened, including HTTP(S), FTP(S), standard I/O, filesystem, and others as shown at `http://php.net/manual/en/wrappers.php`.

> **TIP**   Some users have reported problems with the "+" modes. Many of these problems actually appear to be caused by slightly faulty understanding of the six modes. When in doubt, try opening in separate read and write modes. See the section on file writing later in this chapter.

HTTP fopen

An HTTP `fopen()` tries to open an HTTP 1.0 connection to a file of a type 0 that would normally be served by a web server (such as HTML, PHP, ASP, and so on). PHP actually fakes out the web server so that it thinks the request is coming from a normal web browser surfing the 'Net rather than a file-open operation.

You should be able to use forward slashes like this on either Unix or Windows, since the addresses are URLs rather than filepaths:

```
$fd = fopen("http://www.example.com/openfile.html/", "r");
```

Remember that technically a URL without a trailing slash is malformed, but through incorrect usage most web servers will automatically rewrite the URL with the slash and try redirecting it. Versions of PHP before 4.0.5 did not support redirects, so all HTTP `fopen()` requests would fail without the trailing slash. After 4.0.5, the trailing slash became optional.

Remember that you need not necessarily use an HTTP connection just because you're looking at an HTML file. If you have filesystem access, you can open from the filesystem instead and treat the file as a text file. The HTTP `fopen()` alternative is mostly useful for getting HTML pages from remote web servers — as when you try to "scrape" data from an HTML page. The effect will be much like viewing an HTML page and saving the source code.

PHP versions older than 4.3.0 were unable to make HTTPS `fopen`s. Now, you can accomplish this simply by using "`https://`" rather than "`http://`".

HTTP `fopen()`s are read-only. You will not be able to write to a remote HTML file using this type of file manipulation.

FTP fopen

An FTP `fopen()` attempts to establish an FTP connection to a remote server by pretending to be an FTP client. This is the trickiest of the four options because you need to use an FTP username and password in addition to the hostname and path.

```
$fd = fopen("ftp://username:password@example.com/openfile.txt/",
"r");
```

The FTP server must support passive mode for this method to work correctly. Also, FTP file opens can only be read *or* write, not both at once, and writes can only be to new files, not to existing ones.

PHP has many specific FTP functions, sufficient to implement a complete FTP client in PHP. If you want to do anything except a simple FTP file download, you should probably use them instead. See the PHP manual at `www.php.net/manual/en/ref.ftp.php`.

Standard I/O fopen

Standard I/O read/writes are indicated by `php://stdin`, `php://stdout`, or `php://stderr` (depending on the desired stream). The standard I/O `fopen()` comes into play mostly when PHP is used on the command line or as a system scripting language, à la Perl, because standard I/O is usually associated with terminal windows.

A command-line script using a standard I/O `fopen` looks like this:

```
#! /usr/local/bin/php
<?php
$fp = fopen("php://stdin", "r");
while (!feof($fp)) {
  echo fgets($fp, 4096);
}
echo "\n";
?>
```

You would run it like this from a Linux/Unix command line (Windows would require the PHP interpreter to be called from the command prompt itself).

```
echo "goo goo ga ga" | ./stdin_test.php
```

Filesystem fopen

The most common and useful way to use `fopen()` is from the filesystem. Unless specifically directed otherwise, PHP will attempt to open from the filesystem.

On Windows systems, you can choose to use the Windows format with backslashes if desired — but remember to escape them:

```
$fp = fopen("c:\\php\\phpdocs\\navbar.inc", "r");
```

You can use forward slashes from both Windows and Unix. You should not use a trailing slash for filesystem `fopen()` calls.

> **TIP** Remember that your files, and potentially your directories, need to be readable or writable by the PHP (or web server, if module) process UID rather than by you as a system user. If you share a server, this means any of the other legitimate PHP users may be able to read and/or write to your files.

File read

The `fread()` function takes a file-pointer identifier and a file size in bytes as its arguments. If the file size given is not sufficient to read in the whole file, you will have mysterious problems (unless you're passing in a smaller file size on purpose, which is useful when reading huge files in chunks). Unless you have a reason to do otherwise (such as a huge, unwieldy file), it's best just to let PHP fill in the file size itself, by using the `filesize()` function with the name of the file (or a variable) as the argument:

```
$fstring = fread($fd, filesize($filename));
```

A common error is to type `filesize($fd)` rather than `filesize($filename)`. You may not remember this from the initial example, because in the intervening paragraphs, we've called the used `fopen()` with an actual filename rather than a variable to which that name has been assigned, as in the first example.

This is an extremely useful function because it allows you to turn any file into a string, which can then be manipulated with PHP's large variety of useful string functions. Any string can also be broken up into an array through use of a function like `file()` or `explode()`, which gives you access to the large arsenal of PHP array-manipulation functions. PHP gives you more slicing and dicing functions than a whole set of Ginsu knives!

If you wish to send the entire contents of a file to standard output (meaning, for most PHP installations, echoing it to the web browser window), use `readfile()` instead. This function has file opening built in, so you need not use a separate function to open the file first. The `readfile()` function is equivalent to the combination of `fopen()` and `fpassthru()`.

Beginning with PHP 4.3.0, a new function called `file_get_contents()` was made available. This function returns the entire contents of a file as a string, including the `fopen()`. It is equivalent to `fopen()` and `fread()`, or to `readfile()` except returning the contents as a string rather than straight to standard output.

If you wish to read in and perform operations on a file line by line, you can use `fgets()` instead of `fread()`. Beginning in PHP 4.2.0, if you do not specify a line length as the second argument to `fgets()`, the function will default to 1024 bytes per line.

```
$fd = fopen("samplefile.inc", "r") or die('Cannot find file');
while (!feof($fd)) {
  $line = fgets($fd, 4096);
```

```
   if ($line === $targetline)
     {echo "A match was found!";}
 }
 fclose($fd);
```

If you would rather read the file in as an array, you can use the function `file()` instead. You might want to do this if you're reading one of the many types of data files that use newlines to indicate rows — such as a spreadsheet saved to tab-delimited text format. This creates an array, each element of which is a line from the original file, including an ending newline character. The function `file()` does not require a separate file open or file close step. A single operation using `file()`, such as:

```
$line_array = file($filename);
```

is the equivalent of this:

```
$fd = fopen($filename, "r") or die("Can't open file $filename");
$fstring = fread($fd, filesize($filename));
$line_array = explode("\n", $fstring);
```

CAUTION   **The `file()` function will work correctly only when PHP recognizes newlines. Hopefully, PHP will handle newlines from other operating systems correctly — current Windows and Unix versions of PHP seem to identify newline characters from the other operating system — but we cannot guarantee that this will be true of every case.**

Finally, if you'd like to read in a file character by character, you can use the `fgetc()` function. This will return a character from the file pointer, until the end-of-file. In practice, this function is not used very much, because it's so inefficient to read in a file one character at a time. You'd probably use `fgetc()` only in situations where you wanted to test the first or second character in the file.

Constructing file downloads by using fpassthru()

Besides reading in a file for manipulation by PHP, you can use `fpassthru()` in combination with the PHP `header()` construct to assemble and send file downloads. For instance, let's say you keep lots of tab-delimited data lying around in files, and occasionally you need to let someone download some data from them. Your users are typical businesspeople, not techies, so you know they use IE and would prefer the data as an Excel spreadsheet. So you give the user an HTML form that he or she can use to ask for the data from a particular day, and when it submits you assemble a download and send it like so:

```
<?php

// This example assumes there is one data file per day,
// and your form lets the user specify the date they want to
// see.

$file = $_POST['date'].'.txt';
$fp = $fopen($file, "r");
header("Content-Type:application/xls");
```

```
header("Content-Disposition:attachment;
filename=$_POST['date'].xls");
// Notice we changed the file name and type
header("Content-Transfer-Encoding:binary");
fpassthru($fp);
?>
```

CAUTION **File downloads in PHP are surprisingly tricky because every browser implements the file download behavior differently — even different versions of the same browser can have different behaviors. The preceding method works fine in IE 6.0, but in Mozilla 1.0 the file will claim to be of type** application.xls **but will download as** 20020526.xls.php**. Most of the methods necessary to get a perfect file download are hacks and involve tricking the browser into thinking it's downloading the file directly — for instance by tacking** /$_POST['data'].xls **onto the end of the URL (for example,** http://example.com/sample.php/20020526.xls**). Also, if you saved the script above as** data.xls**, and jiggered your web server into parsing** .xls **files as PHP, you could get a great download in just about every browser. No single perfect method exists for every browser, but this is one situation where you can't just go by what you read in the PHP online manual.**

File write

NOTE **For file writing via PHP, directory permissions must be set to at least 703.**

File writing is pretty straightforward if you've successfully opened in the correct mode for your intended purpose. The function `fwrite()` takes arguments of a file pointer and a string, with an optional length in bytes, which should not be used unless you have a specific reason to do so. It returns the number of characters written.

```
$fout = fwrite($fp, $fstring);
if ($fout != strlen($fstring)){
    echo "file write failed!";
}
```

The function `fputs()` is identical to `fwrite()` in every way. They are simply aliases for one another, but `fputs()` is the C-style function name.

Keep in mind that opening a file in `w` or `w+` modes will result in the complete and utter obliteration of any file contents. These modes are meant for clean overwrites only. If you want to write to the beginning or end of a file, use `r+` or `a+`, respectively.

Probably the most common error with PHP file writing modes involves using a web interface (in other words, an HTML form) to edit a text file. If you want to open a file, read in and view the contents, then write an edited version to the same filename, you cannot depend on `w+` mode. The `w` modes erase the contents of the file *immediately upon opening it* — thus, although you can `fread()` from a `w+` file, there will be no text to read until after you write to it. To get around this issue, you need to open once in read mode and once in write mode, as in the following example:

```
<?php
if (IsSet($_POST['submitted'])) {
```

```
        $fd = fopen($filename, "w+")
          or die("Can't open file $filename");
        $fout = fwrite($fd, $_POST['newstring']);
        fclose($fd);
    }
    $fd = fopen($filename, "r") or die("Can't open file $filename");
    $initstring = fread($fd, filesize($filename));
    fclose($fd);
    echo "<HTML>";
    echo "<FORM METHOD='POST' ACTION=\"$_SERVER['PHP_SELF']\">";
    echo "<INPUT TYPE='text' SIZE=50 NAME='newstring'
    VALUE=\"$initstring\">";
    echo "<INPUT TYPE='HIDDEN' NAME='submitted' VALUE=1>";
    echo "<INPUT TYPE='SUBMIT'>";
    echo "</FORM>";
    echo "</HTML>";
    ?>
```

Let us reiterate that file writing is not at all a good idea unless you can control your environment very tightly! In other words, a well-hardened intranet server might be appropriate, but file writing on a production web site can be a security risk. For more information, see Chapter 28.

As we explain in Chapter 29, in PHP there is now a very easy mechanism to disable the capability to file write. This is a great idea especially if your site is entirely database-driven, in which case you don't have any legitimate need to write to the filesystem with PHP anyway. To disable file writing, simply add fwrite to the list of disabled functions in php.ini:

```
disabled_functions = "fwrite"
```

If you don't use php.ini and need to set this value in Apache httpd.conf, remember that it requires a php_admin_value flag (rather than php_value):

```
php_admin_value disabled_functions="fwrite";
```

File close

File closing is straightforward:

```
fclose($fd);
```

Unlike fopen(), the result of fclose() does not need to be assigned to a variable. File closing may seem like a waste of time, but your system has only so many file descriptors available, and you may run out if you do not close your files. On the other hand, PHP will close all open files when your script ends, and at least one version of PHP3 had a buggy fclose() function, which would crash the server. You know your own setup best, and you can make the call.

Filesystem and Directory Functions

Most of these functions will be quite familiar to Unix users, as they closely replicate common system commands.

> **CAUTION**   **Many of the functions in this section are dangerous. Because they duplicate functions that can and should be performed from the local system, they can be a cracker's bonanza without providing much value to legitimate users. Strongly consider disabling them using PHP's `disable_functions` directive (as discussed in the preceding section on file writing)!**

The one piece of good news is that some of these functions will only work if the PHP process is running as the superuser. Because this is not the default case in the web browser, presumably these functions are intended to be used by the scripting version of PHP, and only trusted users who know what they're doing are even in a position to shoot themselves in the foot this way. Of course, if you are foolish enough to run your web server as root, you are doubly screwed.

The most common and safest functions are listed first in the following sections; the less common and less safe are listed in Table 23-1.

feof

The `feof` function tests for end-of-file on a file pointer and takes a filename as argument. It's used mostly in a `while` loop to perform the same function on each line in a file:

```
while (!feof($fd)) {
  $line = fgets($fd, 4096);
  echo $line;
}
```

file_exists

The `file_exists` function is a simple function you will use again and again if you use filesystem functions at all. It simply checks the local filesystem for a file of the specified name.

```
if (!file_exists("testfile.php")) {
  $fd = fopen("testfile.php", "w+");
}
```

The function returns `true` if the file exists, `false` if not found. The results of this test are stored in a cache, which may be cleared by use of the function `clearstatcache()`.

filesize

Another simple but useful function is `filesize`, which returns and caches the size of a file in bytes. We use it in all the `fread()` examples earlier in this chapter. Never pass in a filesize as an integer if you can do it by using `filesize()` instead.

TABLE 23-1

Filesystem Functions

Function	Description
basename (filepath, [suffix])	Returns the filename portion of a stated path.
chgrp(file, group)	Changes file to any group to which the PHP process belongs. Inoperative on Windows systems.
chmod(file, mode)	Changes to the stated octal mode. Inoperative on Windows systems.
chown(file, user)	If executed by the superuser, changes file owner to stated owner. Inoperative but returns true on Windows systems.
clearstatcache	Clears cache of file status info.
copy(file, destination)	Copies file to stated destination.
delete(file)	See "unlink."
dirname(path)	Returns the directory portion of a stated path.
disk_free_space("/dir")	Returns the number of free bytes in a given directory.
fgetcsv(fp, length, delimiter [, enclosure])	Reads in a line and parses it for CSV format.
fgetss(fp, length [, allowable_tags])	Gets a file line (delimited by a newline character) and strips all HTML and PHP tags except those specifically allowed.
fileatime(file)	Returns (and caches) last time of access.
filectime(file)	Returns (and caches) last time of inode change.
filegroup(file)	Returns (and caches) file group ID number. Names can be determined by using posix_getgrgid().
fileinode(file)	Returns (and caches) file inode.
filemtime(file)	Returns (and caches) last time of modification.
fileowner(file)	Returns (and caches) owner ID number. Names can be determined by using posix_getpwuid().
fileperms(file)	Returns (and caches) file permissions level.

continued

401

TABLE 23-1 *(continued)*

Filesystem Functions

Function	Description
filetype(file)	Returns (and caches) one of: fifo, char, dir, block, link, file, unknown.
flock(file, operation [,&wouldblock])	Advisory file locking. Operation value must be LOCK_SH (shared), LOCK_EX (exclusive), LOCK_UN (release), or LOCK_NB (don't block while locking). The optional third parameter is set to true if enforcing the lock would block existing access.
fpassthru(fp)	Standard output of all data from file pointer to EOF.
fseek(fp, offset, whence)	Moves file pointer offset number of bytes into file stream from the position indicated by whence.
ftell(fp)	Returns offset position into file stream.
stream_set_write_ buffer(fp [, buffersize])	Sets a buffer for file writing; the default is 8K.
Is_dir(directory)	Returns (and caches) true if named directory exists.
Is_executable(file)	Returns (and caches) true if named file is executable.
Is_file(file)	Returns (and caches) true if named file is a regular file.
Is_link(file)	Returns (and caches) true if named file is a symlink.
Is_readable(file)	Returns (and caches) true if named file is readable by PHP.
is_writable (file/ directory)	Returns (and caches) true if named file or directory is writable by PHP.
link(target, link)	Creates hard link. Inoperative on Windows systems.
linkinfo(path)	Confirms existence of link. Inoperative on Windows systems.
mkdir(path, mode)	Makes directory at location *path* with the given permissions in octal mode.
pclose(fp)	Closes process file pointer opened by popen().
popen(command, mode)	Opens process file pointer.
readlink(link)	Returns target of a symlink. Inoperative on Windows systems.
rename(oldname, newname)	Renames file.

Function	Description
`rewind(fp)`	Resets file pointer to beginning of file stream.
`rmdir(directory)`	Removes an empty directory.
`stat(file)`	Returns a selection of info about file.
`lstat(file)`	Returns a selection of info about file or symlink.
`symlink(target, link)`	Creates a symlink from target to link. Inoperative on Windows systems.
`touch(file, [time])`	Sets modification time; creates file if it does not exist.
`umask(mask)`	Returns `umask`, and sets to `mask & 0777`. With no argument passed, it simply returns the `umask`.
`unlink(file)`	Deletes file.

Network Functions

The network functions are a bunch of relatively little used functions that provide network information or connections. Many of these may be more useful from the command line than the web page, unless you're writing some kind of monitoring tool.

Syslog functions

The `syslog` functions allow you to open the system log for a program, generate a message, and close it again.

- `openlog([ident], option, facility)` is entirely optional when used with `syslog()`. The `ident` value is generated automatically.
- `syslog(priority, message)` generates a system log entry.
- `closelog()` is entirely optional when used with `syslog()`. It takes no arguments.

DNS functions

PHP offers some very slick DNS-querying functions, outlined in the Table 23-2. These functions allow PHP scripts to do some jiggering between IP address (which is available via the Apache `REMOTE_ADDR` variable, for instance) and hostname, or vice versa.

403

TABLE 23-2

DNS Functions

Function	Description
checkdnsrr($host, [$type])	Checks for existence of DNS records. Default is MX; other types are A, ANY, CNAME, NS, SOA, PTR and AAAA. Doesn't exist on Windows.
gethostbyaddr($Ipaddress)	Gets hostname corresponding to address.
gethostbyname($hostname)	Gets address corresponding to hostname.
gethostbynamel($hostname)	Gets list of addresses corresponding to hostname.
getmxrr($hostname, [mxhosts array], [weight])	Checks for existence of MX records corresponding to hostname, places in mxhosts array, fills in weight info. This function doesn't exist on Windows.

Socket functions

A *socket* is a kind of dedicated connection that allows different programs (which may be on different machines) to communicate by sending text back and forth. PHP socket functions allow scripts to establish such connections to socket-based servers. For instance, web and database servers communicate via fsockopen() — so you could theoretically write your own web server in PHP using this function, if you had lost all contact with reality. The connection can then be read from or written to with the standard file-writing functions (fputs(), fgets(), and so on.)

The standard socket-opening function is fsockopen(). The pfsockopen() function is identical except that sockets are not destroyed when your script exits; instead, the connection is pooled for later use. The blocking behavior of socket connections can be toggled with set_socket_block-ing(). When blocking is enabled, functions that read from sockets will hang until there is some input to return; when it is disabled, such functions will return immediately if there is no input. These functions are summarized in Table 23-3.

TABLE 23-3

Socket Functions

Function	Description
fsockopen($hostname, $port, [error number], [error string], [timeout in seconds])	Opens the socket connection to specified port on the host, and returns a file pointer suitable for use by functions like fgets().

Function	Description
`getservbyname($service, $protocol)`	Returns the port number of the specified service.
`getservbyport($port, $protocol)`	Returns service name on port.
`pfsockopen($hostname, $port, [error number], [error string], [timeout in seconds])`	Opens the specified persistent socket connection.
`stream_set_blocking($socket descriptor, $mode)`	`TRUE` for blocking mode, `FALSE` for nonblocking. Default is nonblocking.

Date and Time Functions

These functions are basic tools used in many self-defined functions. You may use them simply to output the date or time, to keep track of microtime for a PHP performance-tracking utility, or to initiate a function over a particular date range (such as putting a *Happy Holidays* message on your site during holiday seasons).

These are pretty straightforward to use if you understand the Unix timestamp. They fall into three main categories: functions that *return* date or time, functions that *format* date or time, and functions that *validate* date.

> **TIP** The Unix timestamp measures time as a number of seconds since the beginning of the Unix epoch (midnight Greenwich Mean Time on January 1, 1970). Despite the name, these functions mostly work on Windows also.

If you don't know either date or time

The fastest way to get a time is to use the function `time()`. This will return the Unix timestamp for your locale, which will look something like `101906652`. If you plan to pass this timestamp to another function or program, this is the best format. Alternatively, you can then use one of the functions in the next section to format the timestamp into something a bit more human-readable.

You could also use `microtime()` to return the current time in seconds and microseconds since the Unix epoch. This can be supremely helpful for utilities that are designed to measure performance. The format is `0.74321900 961906846`, where the first part is microseconds and the second is the Unix timestamp. If you're trying to (for instance) measure the performance of different parts of your web page, you really just want the microseconds part, which can be cut out like this:

```php
<?php
$stampmebaby = microtime();
$chunks = explode(" ", $stampmebaby);
```

```
$microseconds = $chunks[0];
echo $microseconds;
?>
```

A function used to return date information is `getdate($timestamp)`. When used with the argument `time()`, as in `getdate(time())`, it returns an associative array with the following numeric elements derived from the Unix timestamp:

- Seconds
- Minutes
- Hours
- Mday (day of the month, for example 1–31)
- Wday (day of the week, for example 1–7)
- Mon (month, for example 1–12)
- Year (numeric, for example 1984)
- Yday (day of the year, for example 1–365)
- Weekday (day of the week, for example `Sunday–Saturday`)
- Month (for example `January–December`)

You can also use the `getdate()` function with a Unix timestamp other than that representing the current time.

If you want to get the time and format it in one step, you can use `date()` instead of `getdate()`. In the absence of a Unix timestamp argument, `date()` will default to the current local date. This has the advantage of allowing nicer formatting, as we will explain in the next subsection. The function `strftime()` will also format the current Unix timestamp for you (as we explain in the next subsection) unless another is specified.

If you've already determined the date/time/timestamp

The functions in this section come into play if you already have a timestamp and merely wish to format the information more finely. For instance, you may like to express your dates European style (2000.20.04) rather than American (4/20/2000).

The main method to format a timestamp is using `date($format...$formatn[, $timestamp])`. You pass a series of codes indicating your formatting preferences, plus an optional timestamp. For instance:

```
date('Y-m-d');
```

returns a string like `2002-05-27`. You can choose a date with two-zero day identifiers or strictly numeric date identifiers, 12- or 24-hour format, or abbreviated month name. (See the PHP manual for all the options.) An analogous function is `gmdate($format...$formatn [, $timestamp])`, which will return a Greenwich Mean Date.

The function `strftime($format...$formatn[, $timestamp])` is similar but specializes in formatting the time rather than the date; `gmstrftime($format...$formatn [, $timestamp])` returns the time in formatted Greenwich Mean Time.

The function `mktime()` allows you to convert any date into a timestamp. It's subtly different in the order of arguments from the Unix command of the same name, so pay attention. The function `gmmktime()` gives the Greenwich alternative to your own time zone.

Finally, `checkdate($month, $day, $year)` allows you to quickly ensure that a particular date is a valid one. This is great for leap-year questions.

Calendar Conversion Functions

Finally, we have some optional calendar conversion routines, which are now available as an extension.

> **TIP** **Many new users have made the mistake of thinking *calendar functions* mean *date functions*. Not so. These functions strictly convert between different (largely historical) calendar systems. See "Date and Time Functions" earlier in this chapter if you feel you have entered this section in error.**

If you happen to be a French historian, you'll be happy to know that PHP can automatically convert between the French Revolutionary calendar and the Gregorian calendar with but a couple of commands. What can we say to that but: *Bon Thermidor, citoyens et citoyennes!*

Seriously, these functions have real uses — particularly on the global Internet. (And not to be ungrateful or anti-Judeo-Christian-centric . . . but Joyce is patiently and lazily waiting for someone to add the Chinese lunar calendar to PHP, so she can always know when Chinese New Year celebrations will occur.)

Conversion between systems is made possible because all the calendar functions share a universal referent, the so-called "Julian Day Number" (aka "Julian Day Count"). This is an integer that represents the days since noon on the first of January, 4713 BC by the Julian calendar (which wasn't in use at the time, but why nitpick?). This date would be the 14th of January in the Gregorian calendar, which is commonly used in secular societies today. The so-called "Julian Date" is a double that represents the days and hours since Julian Day Zero — but PHP does not allow this level of specificity; we're just mentioning it here in case anyone is looking for this information.

> **TIP** **Remember that the Julian day changes at noon rather than midnight, which is the convention today.**

PHP's calendar conversion functions translate a date in some calendar into or out of Julian Day Count. To convert between two calendars, you will need to use two separate functions: one to give the date from one calendar as a Julian Day Number, and the other to convert JD into another calendar's date. In this example, we are converting a Gregorian date into its equivalent in the Jewish calendar.

```
$jd_no = gregoriantojd(8, 11, 1945);
$hebrew = jdtojewish($jd_no);
echo $hebrew;
```

This will return a date of 2, 6 [Elul], 5705. Conversion to the Jewish calendar is somewhat complicated by the fact that it uses lunar months and its days begin at sunset rather than midnight.

The calendars offered at the moment are:

- French Republican
- Gregorian
- Jewish
- Julian
- Unixian

Each of these calendars has associated "JDTo*X*" and an "*X*ToJD" functions.

Finally, there are two other pairs of miscellaneous calendar functions. JDMonthName() and JDDayofWeek() return the month and day of week of any Julian Day Number in any of the supported calendars, whereas easter_date() and easter_days() will tell you when (Western or Catholic, as opposed to Eastern or Orthodox) Easter falls/fell/will fall in any given year. easter_date() is the more straightforward method but can only be used within a Unix date range (1970–2037). It returns the Unix timestamp of Easter midnight in the specified year.

Summary

PHP has numerous filesystem and system functions built in, which can be extremely handy, although sometimes potentially insecure. A large number of PHP functions duplicate Unix systems utilities, such as chmod() and copy(). PHP can also boast some extra-clever functions such as those for DNS querying. Although we recommend turning off some of these functions, others can be useful in trusted hands and a well-planned environment.

PHP's file opening, reading, and writing functions are extremely powerful tools. Most problems with these functions result from a slightly incorrect understanding of the file-opening modes. In addition to filesystem fopen(), PHP supports very slick HTTP, HTTPS, FTP, and standard I/O file opening.

Finally, PHP offers a plethora of time, date, and calendar functions so you always know what time it is.

Chapter 24

Working with Cookies and Sessions

This chapter might as well have been called "Keeping Track," because its theme throughout is the problem of tracking interactions with users over longer periods of time than it takes to generate a single web page. We explain the extent of PHP support for extended user sessions and for setting and checking cookies, and then cover a couple of related techniques involving directly sending HTTP headers.

Sessions and cookies are closely allied concepts in PHP and in web programming more generally, largely because the best way to actually implement sessions is by using cookies. Sessions are a higher-level concept than cookies, and for this chapter we plan to start at the top and work our way down.

What's a Session?

What do we mean by a session? Informally, a session of web browsing is a period of time during which a particular person, while sitting at a particular machine, views a number of different web pages in his or her browser program and then calls it quits, either for the night or because the person in question actually has a life. If you run a web site that this person visits during that time, for your purposes the session runs from that person's first download of a page from your site through the last. For example, a Caribbean hotel's web site might enjoy a session of five pages duration in the middle of a real user's session that began with a travel portal and ended with that user booking his or her vacation with a competitor.

So what's the problem?

Why is the idea of a session tricky enough that we're just talking about it now, even though PHP is at version 6 already? It's because the HTTP protocol by which browsers talk to web servers is *stateless*, with the result that your web server has less long-term memory than your housecat. That is, your web server reacts independently to each individual request it receives and has no way to link requests together even if it is logging requests. If I sit at my computer in Chicago, and you sit at yours in Monterey, and we both ask for page one and then page two of the Caribbean hotelier's site, the HTTP protocol offers no help toward figuring out that two people looked at two pages each — what it sees is four individual requests for pages, with various information attached to each request. Not only does this information not identify you personally (by name, e-mail address, phone number, or any other traceable identification); it offers nothing reliable to identify your two page requests as being from the same person.

Why should you care?

If our web site's only mission in life is to offer various pages to various users, we may, in fact, not care at all where sessions begin and end. On the other hand, there are a number of reasons why we might in fact care. For example:

- We want to customize our users' experiences as they move through the site, in a way that depends on which (or how many) pages they have already seen.

- We want to display advertisements to the user, but we do not want to display a given ad more than once per session.

- We want the session to accumulate information about users' actions as they progress — as in an adventure game's tracking of points and weapons accumulated or an e-commerce site's shopping cart.

- We are interested in tracking how people navigate through our site in general — when they visit that interior page, is it because they bookmarked it, or did they get there all the way from the front page?

For all of these purposes, we need to be able to match up page requests with the sessions they are part of, and for some purposes it would be nice to store some information in association with the session as it progresses. PHP sessions solve both of these problems for us.

Home-grown Alternatives

Before we look at PHP's treatment of sessions, let's look at a few alternative ways the problem can be handled. As you'll see, the PHP treatment combines a couple of these techniques.

IP address

Web servers usually know either the Internet hostname or the IP address of the client that is requesting a page. In many configurations of PHP, these show up for free as variables — `$_SERVER['REMOTE_ HOST']` and `$_SERVER['REMOTE_ADDR']`, respectively. Now you might think that the identity of the machine at the other end is a reasonable stand-in for the person at the other end, at least over the short term. If you get two requests in quick succession from the same IP address, your code can safely conclude that the same person followed a link or form from one of your site's pages to another.

Unfortunately, the IP address your browser knows about may not belong to the machine your user is browsing from. In particular, AOL and other large operations employ proxy servers, which act as intermediaries. Your user's browser actually requests a URL from the proxy server, which in turn requests the page from your server and then forwards back the page to the user. The result is that many different AOL users might be browsing your site simultaneously, all apparently from the same address. IP addresses are not unique enough to form a basis for session tracking.

Hidden variables

Every HTTP request is dealt with independently, but each time your user moves from page to page within your site, it is usually via either a link or a form submission. If the very first page a user visits can somehow generate a unique label for that visit, every subsequent "handoff" of one page to another can pass that unique identifier along.

For example, here is a hypothetical code fragment that you might include near the top of every page on your site:

```
if (!IsSet($_GET['my_s_id']))
   $my_session_id = generate_session_id();
   // warning! hypothetical function
```

This fragment checks to see if the `$_GET['my_s_id']` variable is bound — if it is, we assume that it has been already set and we are in the middle of a session. If it is not, we assume that we are the first page of a new session, and we call a hypothetical function called `generate_session_id()` to create a unique identifier.

After we have included the preceding code, we assume that we have a unique identifier for the session, and our only remaining responsibility is to pass it along to any page we link or submit to. Every link from our page should include the `$my_s_id` as a `GET` argument, as in:

```
<A HREF="next.php?my_s_id=<?php echo $_GET['my_s_id'];?>">Next</A>
```

And every form submission should have a hidden `POST` argument embedded in it, like this:

```
<FORM ACTION=next.php METHOD=POST>
body of form
<INPUT TYPE=HIDDEN NAME=my_s_id
       VALUE="<?php echo $_GET['my_s_id'];?>" >
</FORM>
```

What's wrong with this technique? Nothing. It works just fine as a way to keep different sessions straight (as long as you can generate unique identifiers). And once we have unique labels for the sessions, we can use a variety of techniques to associate other kinds of information with each session, such as using the session ID as a key for database storage. However, this approach to sessions is a pain to maintain — you must make sure that *every* link and form submission propagates the information as described, or the session identifier will be dropped. Also, if you send the information as GET arguments, your session-tracking machinery will be visible in the web-address box of your user's browser, and such arguments are easily edited by the user. Passing around unique identifiers in GET requests is probably the least secure method of maintaining state in web development, as well as possibly causing problems when your users try to cut and paste links — for instance, if they want to send a link to their friends via e-mail.

Cookie-based home-grown sessions

Another approach to session tracking is to use a unique session identifier as in the previous section but perform the handoff by setting or checking a cookie.

A cookie is a special kind of file, located in the filesystem of your user's browsing computer, that web servers can read from and write to. Rather than checking for a passed GET/POST variable (and assigning a new identifier if none is found), your script checks the user's machine for a previously written cookie file and stores a new identifier in a new cookie file if none is found or if the old cookie has expired. This method has some benefits over using hidden variables: The mechanism works behind the scenes (typically, not showing any trace of its activity in the browser window), and the code that checks or sets the cookie can be centralized (rather than affecting every form and link).

What's the drawback? Some very old browsers do not support cookies at all, and more recent browsers allow users to deny cookie-setting privileges to web servers. So, although cookies make for a smooth solution, we can't assume that they are always available.

NOTE There is a subtle difference in the "coverage" of cookie-based sessions and that of sessions based on GET/POST variables. A variable-based session will only maintain its identity as long as your user stays within your site, following intrasite links or form postings. However, there are any number of ways that a user might go away and come back again within a short period of time — by visiting a site that your site links to, which in turn links back or by wandering away and then finding your site again with a search engine. Cookie-based approaches will treat returns from these little detours as a continuation of the same session, whereas variable-propagation approaches have to treat them as new visits.

We cover cookies in more detail in the "Cookies" section later in the chapter.

How Sessions Work in PHP

Good session support takes care of the following two things:

- Session tracking (that is, detecting whether two separate script invocations are, in fact, part of the same user session).
- Storing information in association with a session.

Obviously, you need the first capability before you can hope to have the second.

PHP session tracking works by a combination of the hidden-variables method and the cookie method described in the preceding section. Because of the advantages of cookies, PHP will use them when the user's browser supports them and, otherwise, will have recourse to stashing the session ID in GET and POST arguments. Fortunately, though, the session functions themselves operate at a more abstract level and take care of checking for cookie support all by themselves. If your version of PHP has been appropriately configured for sessions, you should be able to use the session functions without worrying which method is being used.

> **NOTE** If you want PHP to transparently handle passing session variables for you when cookies are not available, you need to have configured PHP with both the `--enable-trans-sid` and `--enable-track-vars` options. If PHP is not handling this for you, you must arrange to pass a `GET` or `POST` argument, of the form `session_name=session_id`, in all your links and forms. When a session is active, PHP provides a special constant, `SID`, which contains the right argument/value pair. Following is an example of including this constant in a link:
>
> ```
> <A HREF="my_next_page.php?<?php echo(SID);?>">Next page</A>
> ```

Making PHP aware of your session

The first step in a script that uses the session feature is to let PHP know that a session may already be in progress so that it can hook up to it and recover any associated information. This is done by calling the function `session_start()`, which takes no arguments. (If you want every script invocation to look for a session without having to call this function, set the variable `session.auto_start` to 1 in your `php.ini` file, rather than the usual default of 0.) Also, any call to `session_register()` causes an implicit initial call to `session_start()`.

The effect of `session_start()` depends on whether PHP can locate a previous session identifier, as supplied either by HTTP arguments or in a cookie. If one is found, the values of any previously registered session variables are recovered. If one is not found, then PHP assumes that we are in the first page of a new session, and generates a new session ID.

Propagating session variables

Changes in PHP's treatment of global and external variables starting with version 4.1 have made certain things more inconvenient. In our view, though, these changes will also remove a lot of potential confusion about sessions. Accordingly, we'll list two approaches to propagating variables in sessions: one, which is simple and works in PHP version 4.1 or later, and another which is more complicated and works only in PHP version 4.1 or *earlier* (unless you reenable the `register_globals` setting in `php.ini`). (You can guess which one we recommend.)

The simple approach (using $_SESSION)

The simple approach is this: Assuming that you've made a call to `session_start()` (as early in your script as possible), use the `$_SESSION` superglobal array as your suitcase for storing anything that you want to retrieve again from a later page in the same session. Assume that any other

variables will be left behind when you leave this page and that everything in that suitcase will be there when you arrive at the next page.

So, session code to propagate a single numerical variable can be as simple as this:

```php
<?php
session_start();

$temporary_number = 45;
$save_this_one = 19;
$another_temporary = 33;

$_SESSION['save_this'] = $save_this_one;
?>
```

The receiving code can be as simple as the following example:

```php
<?php
session_start();
$saved_from_prev_page = $_SESSION['save_this'];
[..]
$temporary_number = 45;
$another_temporary = 33;
[..]
?>
```

That's all there is to it. Assignment into the $_SESSION superglobal array implicitly does any registration necessary for the new value to be carried forward to the next page.

Note that we could have given the same name to both the variable ($save_this_one) and the corresponding $_SESSION index (save_this), because the two have nothing to do with one another.

For this simple approach, we assume that register_globals has been turned off (as it is by default in versions 4.2 and later), so that no session variables are being automatically promoted into global variables. Or, more precisely, we don't care whether it is turned on or not; the code will work in either situation.

NOTE   **The** $_SESSION **array is one of the superglobal variables introduced in PHP4.1. The** *superglobal* **adjective means that it can be referenced anywhere in PHP code, even within functions, without first being declared global.**

Where is the data really stored?

There are two things that the session mechanism must hang onto: the session ID itself and any associated variable bindings.

As you have seen, either the session ID is stored as a cookie on the browser's machine, or it is incorporated into the GET/POST arguments submitted with page requests. In the latter case, there is really no storage happening — the ID is submitted as part of a request and is returned folded into HTML code for links and forms, which may generate the next request. The browser and server pass this vital information back and forth like a hot potato, and the session is effectively over if either side drops it.

By default, the contents of session variables are stored in special files on the server, one file per session ID. (It's already slightly rude to store the session ID as a client-side cookie — it would be even more rude to store session variable data on the client disk when it's not necessary.) Doing this kind of storage requires the session code to serialize the data, which means turning it into a linear sequence of bytes that can be written to a file and read back to recreate the data.

Obviously, storing session data on the server like this will cause problems in most clusters since each web server will be writing to files on its own (presumably unshared) disk. Unless your cluster-management scheme enforces all page views per session to be served from a particular host — which is uncommon, since in most cases that conflicts with the goals of load management and seamless failover — a new session will be started every time a page request is routed to a different server.

There are three main methods to solve this issue, none of them easy or foolproof to implement. First, a company can write its own custom session-data-sharing layer. In this case, PHP will think it's making normal session-registration calls, but instead of writing to disk, a custom server will intercept the requests and centralize the data. However, developing and maintaining such a server and the customized version of PHP required is not cheap. Second, it's possible to direct PHP to write session data not to the normal local disk location (that is, `/tmp`) but to some other share which could be mounted by all web servers (such as `/shared/session`). This is the fastest solution if you have good sysadmins, since it requires only a change to the `session.save_path` setting in `php. ini`. Finally, it's possible to configure PHP to store the contents of session variables in a server-side database, rather than in files. This is probably the most common solution to the problem, although it should be kept in mind that this strategy will increase the impact of database failures. For more information, see the section "Configuration Issues" later in this chapter.

CROSS-REF **In the first edition of this book, we warned you that serialization support for objects was still problematic, and so we didn't recommend trying to store object variables in sessions. Fortunately, in PHP version 4.1 and later, session serialization seems to be stable. See Chapter 20 for more about object.**

Sample Session Code

In Listing 24-1, we show a short code file, which really has a dual purpose. The first purpose is to provide an example of a full (short) script that successfully uses session functions; the second is to provide a test script that you can use to make sure that you have session support and that it is doing what you expect.

In this listing, we perform the following tasks:

- Initiate a session (or pick up an existing one) by using `session_start()`.
- Check for the existence of a preexisting entry in `$_SESSION`. If one is not present, we assume that the session is new.
- Increment a counter that tracks how many times that the user has visited this page.
- Store the incremented counter back in `$_SESSION`.
- Provide a link back to the page itself, embedding the session ID as an argument if it is found.

LISTING 24-1

Test script using $_SESSION

```php
<?php
session_start();
?>

<HTML><HEAD><TITLE>Greetings</TITLE></HEAD>
<BODY>
<H2>Welcome to the Center for Content-free Hospitality</H2>
<?php

if (!IsSet($_SESSION['visit_count'])) {
  echo "Hello, you must have just arrived.
Welcome!<BR>";
  $_SESSION['visit_count'] = 1;
}
else {
  $visit_count = $_SESSION['visit_count'] + 1;
  echo "Back again are ya?  That makes $visit_count times now ".
       "(not that anyone's counting)<BR>";
  $_SESSION['visit_count'] = $visit_count;
}

$self_url = $_SERVER['PHP_SELF'];
$session_id = SID;
if (IsSet($session_id) &&
    $session_id) {
  $href = "$self_url?$session_id";
}
else {
  $href = $self_url;
}
echo "<BR><A HREF=\"$href\">Visit us again</A> sometime";
?>
</BODY></HTML>
```

This code should be available at www.troutworks.com/phpbook and is suitable for your use in testing your session support if you are using PHP4.1 or later. (See Listing 24-2, a little later in this section, if you are using a pre-4.1 version or if you prefer the register_globals style of using sessions.) After obtaining the code, you should first simply test that it loads without errors. The page you see should look something like that shown in Figure 24-1. After that, to see if cookie-based session support is working, try simply reloading or refreshing the page in your browser. You should see a page that looks something like Figure 24-2.

FIGURE 24-1

Session test page

FIGURE 24-2

Session test page, second visit

If the result of your second visit is Figure 24-2, cookie-based session support is working. If instead it still looks like Figure 24-1, then PHP did not detect your session. Make sure that the browser you are testing with is configured to accept cookies and take a look at the section "Gotchas and Troubleshooting" at the end of this chapter.

The second half of Listing 24-1 is about constructing a self-link that will propagate session information even without cookie support. You can test it by turning off cookies in your test browser. (This is usually an Advanced or Security option in your browser's preferences or options.) After cookies have been turned off, you should be treated as a first-time visitor when you reload the page. However, with cookies off, the SID constant should now contain the session ID name and value, which our code embeds in the link's URL as a GET argument. Clicking on this link should increment the visit

count appropriately, and thereafter either clicking the link or reloading should increment it again (because the session ID should now be in the URL that is being reloaded).

This embedding of the session ID in the URL is exactly what should be unnecessary if PHP has been compiled with --enable-trans-sid. In this case, you should be able to add another self-link to this page, without embedding anything extra in the URL, and PHP should take care of it for you.

Listing 24-2 shows the same test script as in Listing 24-1, except that it does not use superglobal variables and assumes that the register_globals directive has been turned on. It's appropriate if you happen to be using PHP version 4.0.*x*, or if you prefer the register_globals style. Remember that code you write using register_globals will not be portable to many other PHP servers.

LISTING 24-2

Test script assuming register_globals

```php
<?php
session_start();
session_register('visit_count');
?>

<HTML><HEAD><TITLE>Greetings</TITLE></HEAD>
<BODY>
<H2>Welcome to the Center for Content-free Hospitality</H2>
<?php
if (!IsSet($visit_count)) {
  echo "Hello, you must have just arrived.
Welcome!<BR>";
  $visit_count = 1;
}
else {
  $visit_count++;
  echo "Back again are ya?  That makes $visit_count times now ".
       "(not that anyone's counting)<BR>";
}

$self_url = $_SERVER['PHP_SELF'];
$session_id = SID;
if (IsSet($session_id) &&
    $session_id) {
  $href = "$self_url?$session_id";
}
else {
  $href = $self_url;
}
echo "<BR><A HREF=\"$href\">Visit us again</A> sometime";
?>
</BODY></HTML>
```

Session Functions

Table 24-1 lists the most important session-related functions, with descriptions of what they do. Note that in some cases the behavior of these functions depends on configuration options that we detail in the "Configuration Issues" section.

TABLE 24-1

Session Function Summary

Function	Behavior
`session_start()`	Takes no arguments and causes PHP either to notice a session ID that has been passed to it (via a cookie or `GET/POST`) or to create a new session ID if none is found.
	If an old session ID is found, PHP retrieves the assignments of all variables that have been registered and makes those assigned variables available as regular global variables.
`session_register()`	Takes a string as argument and registers the variable named by the string — for example, `session_register('username')`. (*Note:* The variable-name string should not include the leading $.) It can also be passed an array of string arguments to register multiple variables at once. Unnecessary if using `$_SESSION` or `$HTTP_SESSION_VARS`.
	The effect of registering a variable is that subsequent assignments to that variable will be preserved for future sessions. (After a script completes, the registered variables and their values are serialized and propagated in such a way that later calls to `session_start()` can recreate the bindings.)
	If `session_start()` has not yet been called, `session_register` will implicitly call it before executing.
`session_unregister()`	Takes a string variable name as argument and unregisters the corresponding variable from the session. As a result, the variable binding will no longer be serialized and propagated to later pages. (The variable-name string should not include the leading $.) Unnecessary if using `$_SESSION` or `$HTTP_SESSION_VARS`.
`session_is_registered()`	Takes a variable-name string and tests whether a variable with a given variable name is registered in the current session, returning `TRUE` if so and `FALSE` if not. Unnecessary if using `$_SESSION` or `$HTTP_SESSION_VARS`, use `isset()` instead.
`session_destroy()`	Calling this function gets rid of all session variable information that has been stored. (*Note:* A browser's session ID may still be the same after this function call.) It does not unset any variables in the current script or the session cookie.
`session_unset()`	Takes no arguments, and frees all variables in the session. Dangerous if using `$_SESSION` or `$HTTP_SESSION_VARS`; use `unset()` instead.

TABLE 24-1	(continued)

Session Function Summary

Function	Behavior
session_ write_ close()	Manually close session and release write lock on data file. Useful with frames, some clustering situations, and if you do something that might cause PHP to not realize the session has terminated (such as redirection).
session_ name()	When called with no arguments, returns the current session-name string. This is usually 'PHPSESSID' by default. When called with one string argument, session_name() sets the current session name to that string. This name is used as a key to find the session ID in cookies and GET/POST arguments — for successful retrieval, the session name must be the same as it was when the values were serialized and stored. Note that there is no reason to change the session name unless you have some need to distinguish session types that are being served by the same web server (such as in the case of multiple sites that each track sessions). The session name is reset to the default whenever a script executes, so any name change must happen in every script that uses the name, and before any other session functions are called.
session_ module_ name()	If given no arguments, returns the name of the module that is responsible for handling session data. This name currently defaults to 'files', meaning that session bindings are serialized and then written to files in the directory named by the function session_save_path(). If given a string argument, changes the module name to that string. (This could presumably be, for example, 'user' for a user-defined session database, but it should not be changed unless you know what you are doing.)
session_ save_path()	Returns (or sets, if given an argument) the pathname of the directory to which session variable-binding files will be written (which typically defaults to /tmp on Unix systems). This directory needs to exist and have appropriate permissions for PHP to write files to it. On Windows systems, you must change this value to a valid path before using sessions!
session_id()	Takes no arguments and returns a string, which is the unique key corresponding to a particular session. If given a string argument, will set the session ID to that string.
session_ regenerate_ id()	Takes no arguments and sets a new session ID, setting a new cookie if necessary and returning TRUE on success or FALSE on failure. Unlike session_id(), it does not return a string with the actual new ID.
session_ encode()	Returns a string encoding of the state of the current session, suitable for use by string_decode(). This can be used for saving a session for revival at some later time, such as by writing the encoded string to a file or database.
session_ decode()	Takes a string encoding as produced by session_encode() and reestablishes the session state, turning session bindings into page bindings as session_start() does.

Function	Behavior
session_ get_cookie_ params()	Returns an array with current session cookie data: `lifetime` (in seconds till expiration, or `0` for no expiration), `path` (for which the cookie is valid), `domain` (for which the cookie is valid), `secure` (whether or not the cookie will only be sent over SSL connections). These parameters are normally set in the php.ini file, but can be changed for a single script through the `session_set_cookie_params()` function.
session_ set_cookie_ params()	Takes four arguments: int `lifetime` (in seconds till expiration, or `0` for no expiration), string `path` (for which the cookie is valid), string `domain` (for which the cookie is valid), boolean `secure` (whether or not the cookie will only be sent over SSL connections). Be sure to include a trailing slash on the `path` argument.

Configuration Issues

The variables in Table 24-2 can be set in the `php.ini` file and viewable by calling `phpinfo()`. We offer descriptions and the typical default values. (Some defaults are platform-dependent.)

TABLE 24-2

Session Configuration Variables

Php.ini Variable	Typical Value	Description
session. save_path	/tmp under Unix systems	Pathname for the server-side directory where session datafiles will be written. Must be changed for Windows systems!
session. auto_start	0	When `1`, sessions will initialize automatically every time a script loads. When `0`, no session data will be available unless there is an explicit call to either `session_start()` or `session_register()`.
session. save_handler	'files','user'	String that determines underlying method for saving session variable information. Changing this is not recommended for the casual user.
session. cookie_ lifetime	0	Specifies how long session cookies take to expire and, consequently, the lifetime of a session. The default of `0` means that sessions last until the browser is closed — any other value indicates the number of seconds the session is allowed to live.
session.use_ cookies	1	If `1`, the session mechanism will attempt to propagate the session ID by setting/checking a cookie. (If the browser refuses the cookie, then `GET/POST` vars may be used.) If this variable is `0`, no attempt to use cookies is made.

Cookies

Many uses of cookies amount to session tracking — keeping track of some piece of information as a single user navigates through your site. If you are tempted to use cookies for a purpose like this, and you are using PHP4, you might want to consider simply using the built-in session functions that are covered in the section "Cookie-based home-grown sessions" earlier in this chapter. Not only do they offer a nicer level of abstraction, but they also have a built-in fallback mechanism that deals with refusal of cookies by propagating the information via GET/POST arguments instead.

A *cookie* is a small piece of information that is retained on the client machine, either in the browser's application memory or as a small file written to the user's hard disk. It contains a name/value pair — *setting a cookie* means associating a value with a name and storing that pairing on the client side. *Getting* or *reading* a cookie means using the name to retrieve the value. (See the sidebar "Cookies and Privacy," a little later in this chapter, for a summary of the controversy surrounding the use of cookies.)

NOTE As a general rule, you want to store information only in a client-side cookie when storing it on the server is not an option. This is partly simple politeness — try accepting cookies manually for a week, and you'll see some extreme abuses of the technique — but it is also because there are constraints that prevent server abuses of the client's hard disk. In particular, each browser will typically accept only 20 cookies from each domain before it starts popping old cookie values off the stack. If you need to store a lot of info, consider developing a scheme where the cookie file contains an ID that enables you to look up the rest of that information on the server — in other words, some form of sessions.

In PHP, cookies are set using the setcookie() function, and cookies are read nearly automatically. In PHP4.1 and later, names and values of cookie variables show up in the superglobal array $_COOKIES, with the cookie name as an index, and the value as the value it indexes.

The setcookie() function

There is just one cookie-related function, called setcookie(). Table 24-3 shows its arguments, in order, all but the first of which are optional.

TABLE 24-3

Arguments to setcookie()

Argument Name	Expected Type	Meaning
name	string	The name of your cookie (analogous to the name of a variable).
value	string	The value you want to store in the cookie (analogous to the value you would assign to a variable). If this argument is not supplied, the cookie named by the first argument is deleted.
expire	int	Specifies when this cookie should expire. A value of 0 (the default) means that it should last until the browser is closed. Any other integer is interpreted as an absolute time (as returned by the function mktime()) when the cookie should expire.

Argument Name	Expected Type	Meaning
`path`	string	In the default case, any page within the web root folder would see (and be able to set) this named cookie. Setting the path to a subdirectory (for example, "`/forum/`") allows distinguishing cookies that have the same name but are set by different *sites* or subareas of the web server (in this example, the cookie will only be valid in the forum area). Be sure to include a trailing slash in the `path`.
`httponly`	boolean	Cookies set with this flag are only sent through HTTP requests. Default is FALSE.
`domain`	string	In the default case, no check is made against the domain requested by the client. If this argument is nonempty, then the domain must match. For example, If the same server serves `www.mysteryguide.com` and `forum.mysteryguide.com`, one site's code can ensure that the other site does not read (or set) its cookies by including this argument as "`forum.mysteryguide.com`."
`secure`	boolean (TRUE (1) or FALSE (0))	Defaults to `0 (FALSE)`. If this argument is `1 or TRUE`, the cookie will only be sent over a secure socket (aka SSL or HTTPS) connection. Note that a secure connection must already be running for such a cookie to be set in the first place.

CROSS-REF For details about the representation of time used by the `expire` argument, see Chapter 23 — specifically, the discussions of the functions `time()` and `mktime()`

.

CAUTION Calling `setcookie()` results in sending HTTP header information, which cannot be done after you have already sent some regular PHP output (even if that output consists of a single space or blank line!).

Examples

This section provides some example calls to `setcookie()`, along with comments, such as the following:

```
setcookie('membername', 'timboy');
```

This sets a cookie called `membername`, with a value of `timboy`. Because there are no arguments except for the first two, the cookie will persist only until the current browser program is closed, and it will be read on subsequent page requests from this browser to this server, regardless of the domain name in the request or from where in the web root file hierarchy the page is served. The cookie will also be read regardless of whether the web connection is secure. For example, consider the following call:

```
setcookie('membername', 'troutgirl', time() + (60 * 60 * 24),
          "/", "www.troutworks.com", 1);
```

423

This sets the cookie to have the value `'troutgirl'` and would overwrite the previous example's value if it had been set by a previous page. The expiration time is set to 86,400 seconds (or 1 day) after the current time. The path argument is given the most inclusive path possible ("/"), so this cookie will still be read regardless of where it is in the web directory hierarchy. The host argument is set to `'www.troutworks.com'`, which means that subsequent page views will not cause the cookie to be read unless the user actually is making a request of that host. Finally, the last argument specifies that this cookie will only be read or written over a secure socket connection. (If the very connection used by this page is not secure, presumably the cookie will not be set at all.)

NOTE If you want to specify later arguments to `setcookie()` while leaving the earlier ones with their default values, it is best to give the empty string (`""`) for the domain argument, a string containing a slash character (`"/"`) for the path argument, and `0` for the expiration.

CAUTION Multiple calls to `setcookie()` will typically be interpreted in the opposite order that they appear in your PHP script, although not every browser version does this. The best rule is to never send two different values for the same cookie from a single page execution. (Sending more than one is pointless anyway because one of them will always overwrite the other.)

Cookies and Privacy

Cookies have always been controversial from a privacy point of view, and that controversy heats up again periodically. As we wrote the first edition, DoubleClick (an Internet advertising agency) was being flamed for its announcement that it planned to cross-correlate cookie information with a very large database of consumer names, addresses, and buying habits (in an apparent reversal of earlier promises about such behavior).

The worry was that, after a consumer reveals his or her identity on a site by filling out a form and accepting a cookie, any other site that compares notes with the original site could conceivably know the true identity of the user (and lots of other information as well). If this practice became widespread, every e-commerce site you visit might be able to figure out not only your name, address, and buying habits, but also a list of other pages you have visited on the web.

So, cookies worry some people, but at the same time they are also a reasonable and benign workaround to the statelessness of the HTTP protocol. There are plenty of good reasons to want a web client/server interaction to coherently span a few page requests in a row, rather than covering just a single request. As a web developer, you might well decide to use cookies for such a purpose, comfortable in the knowledge that there is no substantive invasion of privacy occurring.

Your comfort is not the same as the user's comfort, however, and many users have set up their browsers to refuse all cookies, as is their right. (Remember that what is at issue here is not only the user's privacy but also the use of his or her own personal hard disk!) Any server-side code you write should gracefully handle a cookie refusal from the client side, and any web sites you design should have easily found privacy policies, so that your users know what they are getting into. This does not mean, though, that you are obligated to provide the same level of service to users that refuse cookies; there are some kinds of functionality that are just too painful to write without them, and deciding that cookie cooperation is a prerequisite to using a privately provided site seems perfectly legitimate.

Deleting cookies

Deleting a cookie is easy. Simply call setcookie(), with the exact same arguments as when you set it, except the value, which should be set to an empty string. This does not set the cookie's value to an empty string — it actually removes the cookie. Remember: If you used the path or domain arguments to set the cookie, you need to use them to unset the cookie too. Another method to clear cookies is to set the expiration time in the past.

Reading cookies

Cookies that have been successfully set in a browser or user's machine will automatically be read on the next request from that browser. This has the following effects:

■ In PHP4.1 and later, the cookie's name/value pair will be added to the superglobal array $_COOKIE, as though we had evaluated $_COOKIE['name'] = value.

■ If the register_globals directive is turned on (for versions earlier than PHP6), a regular page-level global variable will be set to the cookie's value, named the same as the cookie's name. Because register_globals is turned off by default starting with PHP4.2, this feature is not available in 4.2 or later, unless either you or your ISP's administrator has changed the configuration.

So, for example, you can set a cookie as follows:

```
setcookie('membername', 'timboy');
```

This means that, on a *later* page access, you might be able to print the value again as easily as this:

```
$membername = $_COOKIE['membername'];
print("The member name is $membername<BR>");
```

And, if register_globals has been turned on, the later page's use of the cookie becomes even simpler:

```
print("The member name is $membername<BR>");
```

> **NOTE** **If you set a cookie in a given script, it won't be set on the client until that page (and its HTTP headers) are sent off to the client, which is too late for you to be able to take advantage of it in that very script. This means that the corresponding global variable won't be available to you until the next page request.**

The following code typically does not work as you might expect:

```
setcookie('membername', 'timboy');
print("I set a cookie!  Now I will grab the value<BR>");
// (WRONG - the following membername will most likely be blank)
$membername = $_COOKIE['membername'];
print("The member name is $membername<BR>");
```

This is because, as the preceding note points out, the cookie will not be set until the current page's worth of HTTP headers arrives at the client. Because that has not yet happened in this example, and the variable $membername has not been otherwise set, that variable will probably produce an empty string in the preceding print statement.

The following code gets it right:

```
$cookievalue = 'timboy';
setcookie('membername', $cookievalue);
print("I set a cookie for the benefit of future pages<BR>");
// (RIGHT - only print variables that this page actually set)
print("Its name is membername, its value is $cookievalue<BR>");
```

Any subsequent scripts that are loaded into the same browser can now refer to $membername.

CROSS-REF We have already noted some privacy risks to users of accepting cookies from servers. It's worth noting that there are risks that go the other way as well. If you write scripts that depend on the integrity of data that you include in cookies, you should remember that a clever end user can edit those cookies and install arbitrary values in them. See Chapter 28 for techniques for encrypting sensitive data, even inside cookies.

Cookie pitfalls

It is hard to do much wrong with cookies purely at the PHP level. After all, setting a cookie involves only one function (set_cookie()), and reading cookies involves no functions at all. What could be easier than that? The problems that typically arise are those imposed by the HTTP protocol itself.

Sending something else first

The single most common error in using cookies is trying to set a cookie after some regular HTML content has already been generated. (We may be repeating ourselves here, but we will also repeat it in the "Sending HTTP Headers" section later in the chapter, because this fact applies to other direct HTTP protocol manipulations in addition to cookies and is the cause of a great deal of debugging confusion.)

The reason this doesn't work is that the HTTP protocol requires headers to be sent before the content of the HTML page itself — they can't be intermixed. As soon as any regular content is generated, PHP figures that it must already know about all headers of interest, and so it sends them off and then begins the transmission of HTML content. If it encounters a cookie (or other header information) later on, it is too late, and an error is generated.

It's surprisingly easy to write code that violates this prohibition. Consider the following:

```
 <?php /* A subtle, insidious cookie error */
setcookie('mycookie', 'myvalue');
?>
<HTML><HEAD>
<TITLE>A seemingly benign cookie-setting page</TITLE>
```

```
</HEAD><BODY>
   <H3>This page is so simple it absolutely must be right</H3>
</BODY></HTML>
```

When we load this script, we get browser output indicating `cannot add header information`. The culprit is the very first character in the file: the space before `<?php`. Because PHP files start off in HTML mode by default, this file causes one space's worth of generated content to be sent to the client before PHP mode kicks in, and the attempt is made to set the cookie.

A similar way to accidentally send header information too early is to `include()` or `require()` a file that includes blank lines at the end after the closing PHP tag. Finally, of course, you can violate the prohibition entirely in PHP mode, but only if you include something like a `print` or `echo` statement.

If you ever run into this kind of error, it is relatively easy to debug if you are methodical about it. Try moving HTTP-related code toward the beginning of the script file first — if you still get the error after that, then trace backward from the offending line toward the beginning of the file. Somewhere between the beginning and the failing statement you either have some characters that are being interpreted in HTML mode, or else you have a PHP printing construct. If you have any included PHP files before the offending statement, make sure that there are no characters at all before the start tags or after the end tags.

Reverse-order interpretation

As with most HTTP commands, calls to `setcookie()` may actually be executed in the opposite order from the way that they appear in your PHP script, but this depends on the particular browser your user is running and the version of PHP you're using. This means that a pair of successive statements like the following probably have the counterintutive result of leaving the `"mycookie"` cookie with no value, because the unsetting statement is executed second.

```
setcookie("mycookie");// get rid of the old value (WRONG)
setcookie("mycookie", "newvalue");// set the new value (WRONG)
```

> **TIP** There is no need to remove a cookie before setting it to a different value — simply set it to the desired new value. Among other things, this means that the confusing reverse order of interpretation of `setcookie()` calls should not usually matter — if the effect depends on the order, it may mean that you are doing something wrong (or at least something unnecessary).

Cookie refusal

Finally, be aware that `setcookie()` makes no guarantees that any cookie data will, in fact, be accepted by the client browser — `setcookie()` just agrees to try, by sending off the appropriate HTTP headers. What happens after that is up to the client, and the client may be an older browser that does not accept cookies or a browser whose user has intentionally disabled cookies.

The `setcookie()` function does not even return a value that indicates acceptance or refusal of the cookie. If you think about it, this is imposed by the timing of the script execution and the HTTP protocol. First, the script executes (including the `setcookie()` call), with the result that a page complete with HTTP headers is sent to the client machine. At this point, the client browser

decides how to react to the cookie-setting attempt. Not until the client generates another request can the server receive the cookie's value and detect whether the cookie-setting attempt was successful. The implication of this for scripting is that you must always ensure that something reasonable happens, even in cases where `setcookie()` is called without success. One common technique is to set a test cookie with the name `CookiesOn` and then check on a subsequent page load if the `$_COOKIE['CookiesOn']` variable has been set.

Sending HTTP Headers

The `setcookie()` call provides a wrapper around a particular usage of HTTP headers. In addition, PHP offers the `header()` function, which you can use to send raw, arbitrary HTTP headers. You can use this function to roll your own cookie function if you like, but you can also use it to take advantage of any other kind of header-controlled functionality.

The syntax of `header()` is as simple as it can be: It takes a single string argument, which is the header to be sent.

> **CAUTION** All the cautions from earlier in this chapter (about sending HTTP before any real page content) apply to the `header()` function as well.

Example: Redirection

One useful kind of HTTP header is `"Location:"`, which can act as a redirector. Simply put a fully qualified URL after the `"Location:"` string, and the browser will start over again with the new address instead. Here's an example:

```php
<?php
  if (IsSet($_GET['gender']) && ($_GET['gender'] == "female"))
    {
       header(
"Location: http://www.example.com/secret.php");
       exit;
    }
?>
<HTML><HEAD><TITLE>The inclusive page</TITLE></HEAD></HTML>
<BODY>
<H3>Welcome!</H3>
We welcome anyone to this page, even men!  Talk amongst yourselves.
</BODY></HTML>
```

If we simply enter the URL for this page (`www.example.com/inclusive.php`), we will see the rendering of the HTML at the bottom of the script. On the other hand, if we include the right `GET` argument (`www.example.com/inclusive.php?gender=female`), we find ourselves redirected to a different page entirely. Note that this is significantly different from selectively importing contents with the `include()` statement — we actually end up browsing a different URL than the one we typed in, and that new web address is what shows up in the Location or Address bar of your browser.

This kind of redirection can be useful when you want the structure of your web site to conditionally branch without having to make the user explicitly choose different links.

Example: HTTP authentication

Another useful thing you can do with HTTP is ask the browser to ask the user for a username and password, via a pop-up window. This is done with the `WWW-Authenticate` header, as in the following example:

```php
<?php
  $the_right_user = 'user';  // example only! not recommended
  $the_right_password = 'password';  // example only!

  if(!isset($_SERVER['PHP_AUTH_USER'])) {
    Header("WWW-Authenticate: Basic realm=\"PHP book\"");
    Header("HTTP/1.0 401 Unauthorized");
    echo "Canceled by user\n";
    exit;
  } else {
    if (($_SERVER['PHP_AUTH_USER'] == 'user') &&
        ($_SERVER['PHP_AUTH_PW'] == 'password')) //see caution below
      print("The realm is yours<BR>");
    else
      print("We don't need your kind<BR>");
  }
?>
```

If we visit this script for the first time (and are using the appropriate browser and server versions), we will get a pop-up window. After the user enters the information into the pop-up box, the script is automatically called again with new variables `$_SERVER['PHP_AUTH_USER']` (set to the user string entered), `$_SERVER['PHP_AUTH_PASSWD']` (set to the password string entered), and `$_SERVER['PHP_AUTH_TYPE']` (which will be *Basic* until such time as another type of authentication is supported). The nice thing about this is that these variables will continue to be set by the browser on each request, and you do not need to do anything in your scripts to propagate them — one verification of identity per session should suffice.

CAUTION   **The preceding code is the bare minimum necessary to demonstrate the HTTP authentication mechanism and is not a model for how user/password combinations should really be verified! Our code fragment simply compares the values of the variables delivered to hardcoded strings, which is a bad idea for several reasons. To make this part of a real verification system, you probably want to compare the result of encrypting the password to a similarly encrypted version in a database or password file. See Chapter 28 for more on encryption and real security measures.**

In addition to redirection and authentication, the capability to send real HTTP headers offers finer control of many aspects of the HTTP client/server relationship, which usually are set by default. For example, you can explicitly set the expiration and caching behavior of your page, or send return status codes that tell the client whether whatever is returned should be considered a success or not.

Because PHP is just acting as a channel to the underlying HTTP protocol, most of these techniques are beyond the scope of PHP documentation and this book.

> **NOTE** The `WWW-Authenticate` **mechanism works only under the Apache Web Server, with PHP as a module. It does not currently work in the CGI version or under IIS/PWS.**

Header gotchas

As we have said innumerable times by now, the `header()` function is subject to the same restriction as the `setcookie()` function: No headers may be sent after regular page content is generated, unless you are using a release of PHP that has output buffering enabled.

More generally, be aware that using the header capability requires not only some knowledge of the HTTP protocols but also some knowledge of the extent to which different browser versions conform to them. Unless you are writing for a known population of users that all use the same browser, you will probably need to do more cross-browser testing than with vanilla HTML-generating scripts.

> **TIP** **Most browsers can be set to warn you whenever they are about to accept a cookie. Although this can be annoying when viewing benign yet cookie-intensive sites, it can also be a great debugging tool when writing your own cookie-setting code. Mozilla browsers also feature a tool called Cookie Manager that lists cookies from each site and allows you to manually delete them, which is also handy for debugging.**

Gotchas and Troubleshooting

If you are having trouble with sessions, first make sure that your session support exists and is doing what you think it is. Try downloading the sample session code from `www.troutworks.com/phpbook` and debugging it from the earliest error, if any.

If sessions are not working or are giving errors, check the pathname returned by `session_save_path()`, and make sure that it exists and is PHP writable. If not, you should either make it so or change the value of `'session.save_path'` in `php.ini`.

Remember that session functions that have variable names as arguments do not expect a leading $ in the name.

If you ever run into a complaint that refers to already having sent HTTP headers, it may be that your script is sending some text (even blank lines) before the `session_start()` or `session_register()` functions. Scrutinize any included files for blank lines or move the session functions to the very beginning of your file.

When testing session-related code, remember to try it out both with a browser that accepts cookies and with a browser that is set up to refuse them. If you see no session name in the URL of a link (such as, `'PHPSESSID'`) with a cookie-refusing browser, then either sessions are not working or your version of PHP is not configured to transparently pass session IDs in the `GET/POST` arguments. It's also informative to try session and cookie code with a browser that is configured to alert the user whenever it is setting a cookie.

Summary

Sessions are useful for tracking a user's behavior over interactions that last longer than one script execution or page download. If what you present to the user depends on which previous pages he or she has seen or interacted with, your code must store or propagate that information in a way that distinguishes one user from another. Because the HTTP protocol is stateless, this inevitably entails some kind of workaround technique — usually either *hidden variables* (which impose maintenance headaches) or *cookies* (which are not universally supported by client browsers).

The PHP implementation of sessions encapsulates these messy issues and presents a clean layer of abstraction to the scripter. Unique session identifiers are automatically created and propagated, and variables can be passed from page to page by storing them in the superglobal `$_SESSION` array. Aside from one's having to connect to a session initially and store (or register) the variables that should persist beyond the current page, session use is virtually transparent to the programmer.

PHP offers several ways to use the capabilities of the HTTP protocol, in addition to the obvious one of constructing HTML pages that are transmitted via HTTP. The `setcookie()` function allows you to set and delete cookies in your user's browser, the values of which show up in subsequent page views as ordinary global variables.

The `header()` function allows you to send arbitrary HTTP headers. Among other things, `header()` can be useful for authentication and page-level redirection.

The HTTP functions in PHP are very simple, and the main complexities that arise are a consequence of the HTTP protocol itself. One such complication is the fact that that HTTP requires all headers to be sent before any page content is sent. Remember that any use of header-manipulating functions must happen before even a blank space is sent to the browser.

Chapter 25

Learning PHP Types

In Chapter 4, we covered PHP types in basic terms, outlining the different types and how they might best be used in your programs. Our first purpose in this chapter is to review those types and elaborate a little more on *resources*. (Another type, *objects*, was covered fully in Chapter 20.) We'll also look at some type-testing techniques, and finally, type conversion.

Type Round-up

You should remember from earlier chapters that unlike many other languages, PHP does not require explicit type declarations. PHP is fairly intuitive about the purpose of your various variables and can often infer your intent from the context in which those variables are used. PHP, for example, understands that the statement:

```
$my_value = 4.50;
```

refers to a float, that is, a floating-point number. But if you subsequently create a string, such as:

```
$my_string = "I paid \$$my_value for a box of
twinkies.";
```

PHP understands that it needs to convert the variable $my_value into a string for purposes of concatenating the larger string assigned to

433

the variable `$my_string`. However, this conversion should in no way prevent you from later doing something like:

```
$my_tax = .065;
$my_total = round($my_value + ($my_value * $my_tax));
$my_string .= "However; if I lived in a state where
Twinkies are not considered a food item, the same box
would have cost
\$$my_total";
```

The eight basic PHP types are listed here. If you need more of a reminder, refer to Chapter 4. The more complex types are treated in their own chapters: strings in Chapter 7, arrays in Chapter 8, and objects in Chapter 20.

- *Integers* are whole numbers, without a decimal point, like 495.
- *Floats* (aka *doubles*) are floating-point numbers, like 3.14159, or 49.0.
- *Booleans* have only two possible values: TRUE and FALSE.
- *NULL* is a special type that has only one value: NULL.
- *Strings* are sequences of characters, like `'PHP4.0 supports string operations'`.
- *Arrays* are named and indexed collections of other values.
- *Objects* are instances of programmer-defined classes, which can package up both other kinds of values and functions that are specific to the class.
- *Resources* are special variables that hold references to resources external to PHP (such as database connections).

Resources

As previous chapters provided in-depth coverage of the first seven types, let's have a look at the eighth. *Resources* are special values that refer to memory or state information that is external to the PHP language itself. You don't have to know too much about resources for casual PHP programming — we'll briefly explain what resources are all about, but feel free to skip to the section "How to handle resources."

What are resources?

The resource type is needed when PHP communicates with some external program (which may be a database or a graphics program) that allocates memory in response to requests from PHP. In general, PHP programmers do not have to worry about freeing memory within PHP — if you create a string in a PHP script (which will take up some space in memory), you can forget all about it and let your

script run until the end. PHP (or the web server it is attached to) will reclaim all memory associated with your script when your script is done, if not earlier.

External programs (databases, and so on.) might not be smart enough to do this deallocation. You might have space reserved in your database's memory for your script long after your script has gone to script heaven. The way this problem gets handled in PHP is that all special functions that request memory from such external programs return *resources*, which PHP tracks to see if your script can still get to them. If nobody can reach the resource, PHP makes sure that the external program does the right kind of cleanup. PHP does this by counting references to the resource — if the reference count goes to zero, then the resource can be freed.

How to handle resources

In general, PHP programmers do not create resources by themselves — they call special functions that return values of the resource type, and then pass them on to other functions that require resources. For example, you might call the function `mysql_connect()` (which returns a resource value that refers to a connection to a MySQL database), save the result in a variable, and then pass it on to `mysql_query()` (which uses the connection resource to query the database).

Essentially, all you have to do with this connection resource is store it in a variable and pass that variable to functions that require it. You can depend on PHP to clean up the resource after your script is done. If, for whatever reason, you feel that the resource is tying up enough memory during script execution that you want the memory freed before the script is done, you can usually do something like this:

```
$my_resource = mysql_connect();  // stores  variable
// .. code that uses the connection resource ..
$my_resource = NULL; // variable no longer refers to resource
```

The reassignment of `$my_resource` should cause PHP to check that no other piece of code is using the `MySQL` resource and then free it. Alternatively, most resource-opening functions have resource closing counterparts such as `mysql_close()`, covered in Chapter 15, or `fclose()`, which we used in Chapter 23.

Type Testing

Especially because variables can change types because of reassignment, it is sometimes necessary to find out the type of a value at program execution time. PHP offers both a general type-testing function (`gettype()`) and individual Boolean functions for each of the five types. These functions, some of which have alternate names, are summarized in Table 25-1.

TABLE 25-1

Functions for Type Testing

Function	Behavior
gettype(arg)	Returns a string representing the type of arg: either integer, float, string, array, object, or unknown type
is_int(arg) is_integer(arg) is_long(arg)	Returns TRUE if arg is an integer, and FALSE if not
is_double(arg) is_float(arg) is_real(arg)	Returns TRUE if arg is a float, and FALSE if not
is_bool(arg)	Returns TRUE if arg is a Boolean value (TRUE or FALSE), and FALSE if not
is_null(arg)	Returns TRUE if arg is of the NULL type, and FALSE if not
is_string(arg)	Returns TRUE if arg is a string, and FALSE if not
is_array(arg)	Returns TRUE if arg is an array, and FALSE if not
is_object(arg)	Returns TRUE if arg is an object, and FALSE if not
is_resource(arg)	Returns TRUE if arg is a resource, and FALSE if not
is_binary(arg)	Returns TRUE is arg is a binary string, and FALSE if not
is_buffer(arg)	Returns TRUE is arg is a Unicode or binary string, and FALSE if not
is_scalar(arg)	Returns TRUE is arg is a scalar variable, and FALSE if not
is_unicode(arg)	Returns TRUE is arg is a Unicode string, and FALSE if not
is_numeric(arg)	Returns TRUE is arg is numeric, and FALSE if not

Assignment and Coercion

As we have said, PHP often automatically converts from one type to another when the context demands it, and as it turns out, the PHP programmer can also force some of these conversions to happen. In either situation, the programmer should know what to expect.

Type conversion behavior

Here are some general rules for PHP's conversion from one type to another:

- **Integer to float:** The exact corresponding float is created (for example, the int 4 becomes the float 4.0).

- **Float to integer:** The fractional part is dropped, truncating the number toward zero.
- **Number to Boolean:** FALSE if exactly equal to 0, TRUE otherwise.
- **Number to string:** A string is created that looks exactly the way the number would print. Integers are printed as a sequence of digits, and floats are printed with the minimum precision needed. Extreme float values will be converted to scientific notation.
- **Boolean to number:** 1 if TRUE, 0 if FALSE.
- **Boolean to string:** '1' if TRUE, the empty string if FALSE.
- **Null to number:** 0.
- **Null to boolean:** FALSE.
- **String to number:** Equivalent to *reading* a number from the string, then making a conversion to the given type. If a number cannot be read, the value is zero. Not all of the string needs to be read for the reading to be considered a success.
- **String to Boolean:** FALSE if it is an empty string or the string is '0', TRUE otherwise.
- **Simple type (number or string) to array:** Equivalent to creating a new array with the simple value assigned to index zero.
- **Array to number:** Undefined (see following note).
- **Array to Boolean:** FALSE if the array has no elements, TRUE otherwise.
- **Array to string:** 'Array'.
- **Object to number:** Undefined (see note below).
- **Object to boolean:** TRUE if the object contains any member variables that have a value, and FALSE otherwise.
- **Object to string:** 'Object'.
- **Resource to Boolean:** FALSE.
- **Resource to number:** Undefined (see note below).
- **Resource to string:** Something like 'Resource id #1' (but this should not be relied upon).

In the preceding list, we noted that some types have an undefined result when converted to numerical values. In this context, *undefined* simply means that the PHP developers are not making a commitment as to what kind of behavior you'll get in future versions of PHP, so it would be a bad idea to depend on a particular behavior in your code. You may find that these types can be converted to numbers in expressions in your particular version of PHP, but that may not work in the next version.

Explicit conversions

PHP offers three different ways for the programmer to manipulate types: conversion functions, type casts (as in the C language), and calling settype() on variables:

- The functions intval(), floatval(), and strval() will convert their arguments to an integer, a float, or a string, respectively. (At this writing, there does not seem to be a boolval() function.)

■ Any expression can be preceded by a type cast (the name of the type in parentheses), which converts the expression result to the desired type.

■ Any variable can be given as a first argument to `settype()`, which will change the type of that variable to the type named in the second string argument.

For example, each of the following approaches will put the correct count of canines (101) into the integer variable `$dog_count` by the end of the code snippet:

Version #1:

```
$dog_count = intval (strval (floatval("101 Dalmatians")));
```

Version #2:

```
$dog_count = (int) (string) (float) "101 Dalmatians";
```

Version #3:

```
settype($dog_count, "float");
settype($dog_count, "string");
settype($dog_count, "int");
```

> **TIP**   Of course, each approach in the example takes an indirect route, converting needlessly to string and float types — it would suffice to convert immediately to the integer type.

Six of the basic type names (integer, float, boolean, string, array, and object) are valid in casts and are valid string arguments to `settype()`. In addition, certain alternate names are valid in casts: (`int`) instead of (`integer`), (`double`) or (`real`) instead of (`float`), and (`bool`) instead of (`boolean`). It is not valid to cast to type resource, and casting to type `NULL` is pointless. (Because the result can only be the value `NULL`, you might as well simply assign instead.)

Conversion examples

Listing 25-1 shows some PHP code that displays various type conversions in an HTML table, with the resulting table shown in Figure 25-1. (This code is not intended as a style example, and it uses several constructs that have not yet been covered — feel free to just look at the output.)

LISTING 25-1

Type conversions

```
$type_examples[0] = 123; // an integer
$type_examples[1] = 3.14159; // a float
$type_examples[2] = "a non-numeric string";
$type_examples[3] = "49.990 (begins with number)";
$type_examples[4] = array(90,80,70);

print("<TABLE BORDER=1><TR>");
print("<TH>Original</TH>");
```

```php
print("<TH>(int)</<TH>");
print("<TH>(float)</<TH>");
print("<TH>(string)</<TH>");
print("<TH>(array)</<TH></TR>");

for ($index = 0; $index < 5; $index++)
  {
    print("<TR><TD>$type_examples[$index]</TD>");
    $converted_var =
        (int) $type_examples[$index];
    print("<TD>$converted_var</TD>");
    $converted_var =
        (float) $type_examples[$index];
    print("<TD>$converted_var</TD>");
    $converted_var =
        (string) $type_examples[$index];
    print("<TD>$converted_var</TD>");
    $converted_var =
        (array) $type_examples[$index];
    print("<TD>$converted_var</TD></TR>");
  }
print("</TABLE>");
```

FIGURE 25-1

Type conversion examples

Original	(int)	(float)	(string)	(array)
123	123	123	123	Array
3.14159	3	3.14159	3.14159	Array
a non-numeric string	0	0	a non-numeric string	Array
49.990 (begins with number)	49	49.99	49.990 (begins with number)	Array
Array	1	1	Array	Array
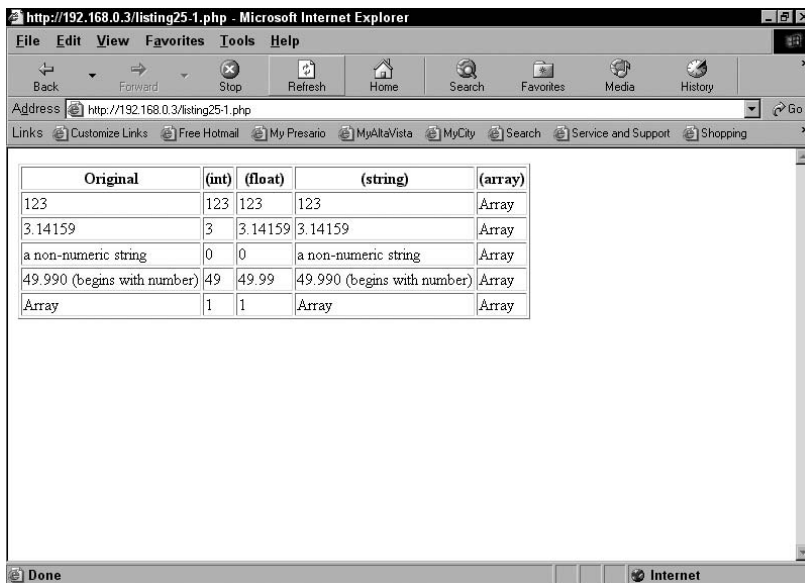
Other useful type conversions

The functions listed in Table 25-2 do not exactly convert types, but they return a different type from their main argument in a useful way.

Other Type Conversion Functions

From\To	Integer	String	Array
Integer		`ord()`	
Float	`ceil()`, `floor()`, `round()`		
String	`chr()`		`explode()`
Array		`implode()`	

The function `ceil()` takes a float and returns the integer greater than or equal to that float. For example:

```
$my_float = 4.7;
$my_int = ceil($my_float); // $my_int is equal to 5
$my_float = -4.7;
$my_int = ceil($my_float); // $my_int is equal to -4
```

The `floor()` function is the opposite of `ceil()`. (We'll drop the intermediate assignment to $my_float now.)

```
$my_int = floor(4.7); // $my_int is equal to 4
$my_int = floor(-4.7); // $my_int is equal to -5
```

The `round()` function takes a float and returns the nearest integer. If the fractional part of the float is exactly one half, the rounding is to the highest absolute number.

```
$my_int = round(4.7); // $my_int is equal to 5
$my_int = round(-4.7); // $my_int is equal to -5
$my_int = round(-4.5); // $my_int is equal to -5
```

TIP   If you're looking for a truncate function (simply dropping the fractional part and, therefore, rounding toward zero), notice that this is the behavior you get simply from typecasting from `float` to `int`.

The function `chr()` takes an integer and returns a one-character string with that ASCII value, whereas `ord()` reverses this, returning the ASCII value of the first character in a string.

Finally, `implode()` and `explode()` allow a certain kind of conversion between strings and arrays. `implode()` creates a string from the array it is given as second argument, separating the elements with the string that is its first argument. For example:

```
$words[0] = "My";
$words[1] = "short";
$words[2] = "sentence.";
$sentence = implode(" ", $words);
print("$sentence<BR>");
```

produces the browser output:

```
My short sentence.
```

`explode()` reverses the process, creating an array from a string:

```
$words = explode(" ", "My short sentence.");
```

Integer overflow

One clever automatic type conversion built into PHP relatively recently is that when integer values *overflow* (that is, they are assigned a value larger than they can hold), they become floats. This makes some sense, because floats can accommodate larger magnitudes than integers can. For example:

```
$toobig = 111;
for ($count = 0; $count < 5; $count++)
  {
    $too_big = 1000 * $too_big;
    print("Is $too_big still an integer?<BR>");
  }
```

produces the following browser output:

```
Is 111000 still an integer?
Is 111000000 still an integer?
Is 111000000000 still an integer?
Is 1.11E+14 still an integer?
Is 1.11E+17 still an integer?
```

The shift you see in this example from literal integers to scientific notation reflects a change of `$too_big`'s type from integer to float. Of course, this may lose some information, because the precision of floats is limited, but it is in keeping with the PHP philosophy of doing the best it can in preference to causing an error.

441

Finding the largest integer

If you need to know the largest integer your PHP will support and, for some reason, you believe that it is not the usual $2^{31} - 1$, here's a handy function (which uses concepts not yet covered):

```
function maxint()
{ /* quick-and-dirty function for PHP int size —
     assumes largest integer is of form 2^n - 1 */
  $to_test = 2;
  while(1)
   {
    $last = $to_test;
    $to_test = 2 * $to_test;
    if (($to_test < $last) || (!is_int($to_test)))
        return($last + ($last - 1));
   }
}
/* sample use */
$maxint = maxint();
print("Maxint is $maxint<BR>");
```

Summary

PHP6 has eight types: integer, float, boolean, NULL, string, array, object, and resource. Five of these are simple types: Integers are whole numbers, floats are floating-point numbers, booleans are true-or-false values, NULL has just one value (NULL), and strings are sequences of characters. *Arrays* are a compound type that holds other PHP values, indexed either by integers or by strings. *Objects* are instances of programmer-defined classes, which can contain both member variables and member functions, and which can inherit functions and data type from other classes. Finally, *resources* are special references to memory allocated from external programs, which memory PHP frees automatically when they are no longer needed.

Only values are typed in PHP — variables have no inherent type other than the value of their most recent assignment. PHP automatically converts value types as demanded by the context in which the value is used. The programmer can also explicitly control types by means of both conversion functions and type casts.

Chapter 26

Learning PHP Advanced Functions

I n Chapter 5 we presented the basic features for user-defined functions in PHP. In this chapter, we move on to some exotic properties of functions, including ways to use variable numbers of arguments, ways to have functions actually modify the variables they are passed, and (cooler still) using functions as data.

Variable Numbers of Arguments

It's often useful to have the number of actual arguments that are passed to a function depend on the situation in which it is called. There are three possible ways to handle this in PHP:

- Define the function with default arguments — any that are missing in the function call will have the default value, and no warning will be printed.

- Use an array argument to hold the values — it is the responsibility of the calling code to package up the array, and the function body must appropriately take it apart.

- Use the variable-argument functions (`func_num_args()`, `func_get_arg()`, and `func_get_args()`).

The following sections address each of these possibilities.

443

Default arguments

To define a function with default arguments, simply turn the formal parameter name into an assignment expression. If the actual call has fewer parameters than the definition has formal parameters, PHP will match actual with formal until the actual parameters are exhausted and then will use the default assignments to fill in the rest.

For example, the following function has all its variables defined with defaults:

```
function tour_guide($city = "Gotham City",
                    $desc = "vast metropolis",
                    $how_many = "dozens",
                    $of_what = "costumed villains")
{
  print("$city is a $desc filled with
         $how_many of $of_what.<BR>");
}
tour_guide();
tour_guide("Chicago");
tour_guide("Chicago", "wonderful city");
tour_guide("Chicago", "wonderful city",
           "teeming millions");
tour_guide("Chicago", "wonderful city",
           "teeming millions",
           "gruff people with hearts of
            gold and hard-luck stories to tell");
```

The browser output is something like this, with the intrasentence line breaks, of course, determined by your browser:

```
Gotham City is a great metropolis filled with dozens of costumed
villains.
Chicago is a great metropolis filled with dozens of costumed
villains.
Chicago is a wonderful city filled with dozens of costumed
villains.
Chicago is a wonderful city filled with teeming millions of
costumed villains.
Chicago is a wonderful city filled with teeming millions of
gruff people with hearts of gold and hard-luck stories to tell.
```

The main limitation of default arguments is that the matching of actual to formal parameters is determined by the ordering of both — it's first-come, first-served. This means that there is absolutely no way to use the default-argument mechanism to tell someone about hard-luck stories in Gotham City.

Arrays as multiple-argument substitutes

If you are dissatisfied with the flexibility of multiple arguments, you can bypass the whole argument-counting issue by using an array as your communication channel.

The following example uses this strategy and, in addition, uses a few tricks like the ternary operator (introduced in Chapter 5) and the associative array (covered in Chapters 8 and 21):

```
function tour_brochure($info_array)
{
$city =
 IsSet($info_array['city']) ?
  $info_array['city'] : "Gotham City";
$desc =
 IsSet($info_array['desc']) ?
  $info_array['desc'] : "great metropolis";
$how_many =
 IsSet($info_array['how_many']) ?
  $info_array['how_many'] : "dozens";
$of_what =
 IsSet($info_array['of_what']) ?
  $info_array['of_what'] : "costumed villains";

print("$city is a $desc filled with
       $how_many of $of_what.<BR>");
}
```

This function checks the single incoming array argument for four different values associated with particular strings. Using the ternary conditional operator ?, local variables are assigned to either the incoming value (if it has been stored in the array) or to our comic book defaults. Now, let's try calling this function with a couple of different arrays:

```
tour_brochure(array()); // empty array
$tour_info =
 array('city' => 'Cozumel',
       'desc' => 'destination getaway',
       'of_what' => 'sandy beaches');
tour_brochure($tour_info);
```

In this example, we call `tour_brochure` first with an empty array (corresponding to no arguments) and then with an array that has three of the four possible associative values stored in it. The browser output we get is:

```
Gotham City is a great metropolis filled with dozens of costumed
 villains.
Cozumel is a destination getaway filled with dozens of sandy
 beaches.
```

In both cases, the `dozens` amount is defaulted, because neither array had anything stored under the `how_many` association.

Multiple arguments in PHP4 and above

Beginning with version 4, PHP offers some functions that can be used inside function bodies to recover the number and values of arguments. They are:

- `func_num_args()`: Takes no arguments and returns the number of arguments that passed to the function it is called from.

- `func_get_arg()`: Takes an integer argument *n* and returns the *n*th argument to the function it is called from. Arguments are numbered starting from zero.

- `func_get_args()`: Takes no arguments and returns an array containing all the arguments in the function it is called from, with array indices starting from zero.

All three of these functions will produce a warning if called outside a function body, and `func_get_arg()` will give a warning if it is called with an index higher than the index of the final argument that was passed.

If your function is going to handle the decoding of arguments using these functions, you can take advantage of the fact that PHP doesn't complain about function calls that have more arguments than formal parameters in the definition. Simply define your function to take no arguments, and then use the functions to catch any that are actually passed.

As an example, consider the following two functions, both of which return an array of the arguments they are given:

```
function args_as_array_1 ()
{
$arg_count = func_num_args();
$counter = 0;
$local_array = array();
while ($counter < $arg_count)
  {
   $local_array[$counter] =
     func_get_arg($counter);
   $counter = $counter + 1;
  }
return($local_array);
}

function args_as_array_2 ()
{
    return(func_get_args());
}
```

The first cumbersome function uses `func_get_arg()` to retrieve the individual arguments and bounds the loop using the result of `func_num_args()`, so that no attempt is made to retrieve more arguments than were actually passed. Each argument is individually stored in an array, which is

then returned. Packaging up the arguments like this is already done for free by `func_get_args()`, so the second version of the function is extremely short.

As another example, let's rewrite our earlier `tour_guide()` function to use the multiple-argument functions instead of default arguments:

```
function tour_guide_2()
{
$num_args = func_num_args();
$city = $num_args > 0 ?
  func_get_arg(0) : "Gotham City";
$desc = $num_args > 1 ?
  func_get_arg(1) : "great metropolis";
$how_many = $num_args > 2 ?
  func_get_arg(2) : "dozens";
$of_what = $num_args > 3 ?
  func_get_arg(3) : "costumed villains";

print("$city is a $desc filled with
      $how_many of $of_what.<BR>");
}
tour_guide_2();
```

This has exactly the same behavior as the default-argument version and is subject to the same limitation. The arguments are passed in by position, so there is no way to replace `"costumed villains"` with something else while leaving `"Gotham City"` as the default.

Unfortunately, no better solution presents itself in PHP6; the array solution is, if somewhat more code-intensive, still the most flexible.

Call-by-value

The default behavior for user-defined functions in PHP is *call-by-value*. This means that when you pass variables to a function call, PHP makes copies of the variable values to pass on to the function. So, whatever the function does, it is not able to change the actual variables that appear in the function call. This behavior can be good or bad. It's a nice reassurance if you only want to use a function for its returned value, but it can also be a source of confusion and frustration if changing the passed variable is actually your goal.

CAUTION One of the most important differences introduced with the previous version of PHP, version 5, and its predecessors, is that object instances are always effectively passed by reference, even if a value of another type would be passed by value. This is the result of the fact that object variables store object handles rather than the objects themselves — the handles actually are copied in a pass-by-value situation, but the underlying objects are not. (We dealt with this issue in more detail in Chapter 20.)

Let's demonstrate call-by-value with a fragile and extremely inefficient implementation of subtraction:

```
function my_subtract ($num1, $num2)
{
  if ($num1 < $num2)
    die("Negative numbers are imaginary");
  $return_result = 0;
  while($num1 > $num2)
    {
      $num1 = $num1 - 1;
      $return_result = $return_result + 1;
    }
  return($return_result);
}
$first_op = 493;
$second_op = 355;
$result1 = my_subtract($first_op, $second_op);
print("result1 is $result1<BR>");
$result2 = my_subtract($first_op, $second_op);
print("result2 is $result2<BR>");
```

Reassuringly, we find that our result is the same both times we perform the same subtraction:

```
result1 is 138
result2 is 138
```

This is true even though my_subtract changes the value of its formal parameter $num1 — that variable only holds a copy of the value that was in the actual parameter $first_op, and so $first_op cannot be affected.

Call-by-reference

PHP offers two different ways to have functions actually modify their arguments: in the function definition and in the function call.

If you want to define a function to operate directly on a passed variable, simply put an ampersand in front of the formal parameter in the definition, like this:

```
function my_subtract_ref (&$num1, &$num2)
{
  if ($num1 < $num2)
    die("Negative numbers are imaginary");
  $return_result = 0;
  while($num1 > $num2)
    {
      $num1 = $num1 - 1;
```

```
            $return_result = $return_result + 1;
        }
    return($return_result);
}
$first_op = 493;
$second_op = 355;
$result1 = my_subtract_ref($first_op, $second_op);
print("result1 is $result1<BR>");
$result2 = my_subtract_ref($first_op, $second_op);
print("result2 is $result2<BR>");
```

Now, if we perform exactly the same subtraction calls as we did the first time, we get the output:

```
result1 is 138
result1 is 0
```

This is because the formal parameter `$num1` refers to the same thing as the actual parameter `$first_op` — changing one means changing the other.

You can also force a function to take arguments by reference by prepending the actual parameters with ampersands (although this is a deprecated capability and may disappear in future PHP versions — in fact, it causes an `E_STRICT` level error in PHP6). That is, we can use our original call-by-value function and get the by-reference behavior, like this:

```
$first_op = 493;
$second_op = 355;
$result1 = my_subtract(&$first_op, &$second_op);
print("result1 is $result1<BR>");
$result2 = my_subtract(&$first_op, &$second_op);
print("result2 is $result2<BR>");
```

producing, once again:

```
result1 is 138
result1 is 0
```

As of PHP4, variable references can be used outside of function calls as well. In general, assigning a variable reference (`&$varname`) to a variable will make the two variables aliases of each other rather than distinct variables with the same value. For example:

```
$name_1 = "Manfred von Richtofen";
$name_2 = "Percy Blakeney";
$alias_1 = $name_1;  // vars have same value
$alias_2 = &$name_2; // vars are the same

$alias_1 = "The Red Baron"; // doesn't change real name
$alias_2 = "The Scarlet Pimpernel"; // anonymous forever

print("$alias_1 is $name_1<BR>");
print("$alias_2 is $name_2<BR>");
```

449

gives the browser output:

```
The Red Baron is Manfred von Richtofen
The Scarlet Pimpernel is The Scarlet Pimpernel
```

CAUTION As we noted in the "Call-by-value" section, in PHP6 it is not necessary to explicitly make references to object instances. Objects are always effectively passed by reference.

Variable function names

One of the neater tricks you can do in PHP is to use variables in place of the names of user-defined functions. That is, rather than typing a literal function name into your code, you type a dollar-sign variable — the function that is actually called at runtime will depend on the string that that variable has been assigned to. In some sense, this allows us to use functions as data. This kind of trick will be familiar to advanced C programmers and to even beginning users of any kind of Lisp language (for example, Scheme or Common Lisp).

For example, the following two function calls are exactly equivalent:

```
function customized_greeting ()
{
  print("You are being greeted in a customized way!<BR>");
}
customized_greeting();
$my_greeting = 'customized_greeting';
$my_greeting();
```

and produce the same output:

```
You are being greeted in a customized way!
You are being greeted in a customized way!
```

Because function names are just strings, they can also be used as arguments to functions or be returned as a function's result.

An extended example

Just for fun, let's see what kinds of trouble we can get into by using some of the more advanced features of functions, including using function names as function arguments.

Listing 26-1 shows an extended example of functions that implement a substitution cipher — a rudimentary kind of cryptography that scrambles messages by substituting one letter of the alphabet for another.

LISTING 26-1

A substitution cipher

```
/* Part 1 - cipher algorithm and utility functions */
function add_1 ($num)
{
  return(($num + 1) % 26);
}

function sub_1 ($num)
{
  return(($num + 25) % 26);
}

function swap_2 ($num)
{
  if ($num % 2 == 0)
    return($num + 1);
  else
    return($num - 1);
}

function swap_26 ($num)
{
  return(25 - $num);
}

function lower_letter($char_string)
{
 return ((ord($char_string) >= ord('a')) &&
        (ord($char_string) <= ord('z')));
}

function upper_letter($char_string)
{
 return ((ord($char_string) >= ord('A')) &&
        (ord($char_string) <= ord('Z')));
}

/* Part 2 - the letter_cipher function */
function letter_cipher ($char_string, $code_func)
{
  if (!(upper_letter($char_string) ||
       lower_letter($char_string)))
    return($char_string);
  if (upper_letter($char_string))
    $base_num = ord('A');
  else
```

```
        $base_num = ord('a');
    $char_num = ord($char_string) -
                    $base_num;
    return(chr($base_num +
                ($code_func($char_num)
                  % 26)));
}

/* Part 3 - the main string_cipher function */
function string_cipher($message, $cipher_func)
{
    $coded_message = "";
    $message_length = strlen($message);
    for ($index = 0;
         $index < $message_length;
         $index++)
      $coded_message .=
        letter_cipher($message[$index], $cipher_func);
    return($coded_message);
}
```

Listing 26-1 is in three parts. In the first part, we define a few functions that do simple math on the numbers from 0 through 25, which will represent the letters A–Z in our cipher codes. Function add_1 simply adds 1 to the number it is given, modulo 26 (which just means that numbers that are 26 and larger "wrap around" to start from zero again). 0 + 1 is 1, 1 + 1 is 2, . . . and 25 + 1 is 0. Sub_1 shifts numbers in the other direction, by adding 25 (which in this modular arithmetic is equivalent to sub-tracting 1). 25 + 25 is 24, 24 + 25 is 23, . . . and 0 + 25 is 25. Swap_2 trades the places of pairs of num-bers (0 to 1, 1 to 0, 2 to 3, 3 to 2, . . .). Swap_26 trades high numbers for low numbers (25 to 0, 0 to 25, 24 to 1, 1 to 24, . . .). Each one of these functions will be the basis of a simple cipher code. Finally, we have a couple of utility functions that test whether a character is an uppercase or lowercase letter.

Part 2 is a single function called letter_cipher(), whose job is to take a math function, like the ones in Part 1, and apply it to encode a single letter. First, it tests whether the string it is handed (which should be a single character) is an alphabetic letter; if not, it returns it as is. If the character is a letter, it is transformed into a number using ord(), and the appropriate letter (a or A) is subtracted to bring the number into the 0–25 range. When it is in that range, we apply the cipher function whose name was passed in as a string, and then we convert the number back into a letter and return it.

Finally, Part 3 is the single string_cipher() function, which takes a string message and a cipher function and returns a new string that is the message encoded via the function. It does this by building a new string, letter by letter, from the message string, where each new letter is the result of applying $cipher_func to the numerical representation of the old letter.

Now let's write some code to try out string_cipher():

```
        $original = "My secret message is ABCDEFG";
        print("Original message is: $original<BR>");
```

```
$coding_array = array('add_1',
                      'sub_1',
                      'swap_2',
                      'swap_26');
for ($count = 0;
     $count < sizeof($coding_array);
     $count++)
{
  $code = $coding_array[$count];
  $coded_message =
     string_cipher($original, $code);
  print("$code encoding is: $coded_message<BR>");
}
```

This testing code takes our four predefined letter-encoding functions, stashes them in an array, and then loops through the array, encoding the $original message and printing out the encoded version. The browser output looks like the following:

```
Original message is: My secret message is ABCDEFG
add_1 encoding is: Nz tfdsfu nfttbhf jt BCDEFGH
sub_1 encoding is: Lx rdbqds ldrrzfd hr ZABCDEF
swap_2 encoding is: Nz tfdqfs nfttbhf jt BADCFEH
swap_26 encoding is: Nb hvxivg nvhhztv rh ZYXWVUT
```

We can take this function-as-data approach one step further and write a function that applies more than one cipher to a message in sequence. This function also uses the variable-argument capability we discussed earlier in the chapter.

```
function chained_code ($message)
{
  /* takes a message, then an arbitrary number of
     cipher-code function names. Returns
     result of applying each code to the previous
     result. */
  $argc = func_num_args();
  $coded = $message;
  for ($count = 1; $count < $argc; $count++)
    {
      $function_name = func_get_arg($count);
      $coded =
         string_cipher($coded,
                        $function_name);
    }
  return($coded);
}
```

The first argument to chained_code() should be a message string, followed by any number of names corresponding to cipher functions. The coded message is the result of applying the first coding

function to the message, then applying the second coding function to the result, and so on. We can test it with various combinations of our predefined letter-coding functions, as follows:

```
$tricky =
  chained_code($original,
               'add_1', 'swap_26',
               'add_1', 'swap_2');
print("Tricky encoded version is $tricky<BR>");

$easy =
  chained_code($original,
               'add_1', 'swap_26',
               'swap_2', 'sub_1',
               'add_1', 'swap_2',
               'swap_26', 'sub_1');
print("Easy encoded version is $easy<BR>");
```

with these results:

```
Tricky encoded version is Ma guwjuh muggysu qg YZWXUVS
Easy encoded version is My secret message is ABCDEFG
```

As you can see, the *tricky* encoding of our message is a combination of our previous codes that doesn't correspond exactly to any of those single coding functions. And the "easy" coding is an even more complicated combination of those functions that produces . . . our original message unchanged! (No, it's not that our cipher code doesn't work — we'll leave it to you to figure out why that particular sequence of coding functions brings us around to our starting message again.)

The moral of our little cryptographic scripting story is that, although the cipher code was mildly complicated, it was made considerably simpler by PHP's support for using function names as function arguments.

Summary

The default behavior for user-defined functions is *call-by-value*, meaning that functions work on copies of their arguments and so cannot modify the original variables in the function call. You can force *call-by-reference* behavior by preceding parameters with &, on either the definition side or the calling side. PHP offers more than one way to let functions take a variable number of arguments. Finally, the functions to be called can be determined at runtime, by substituting a string variable for the literal name of the user-defined function — this allows functions to be treated as data and passed back and forth between other functions.

Chapter 27

Performing Math with PHP

In Chapter 9, we covered the most basic aspects of mathematics in PHP: the numerical types, the basic arithmetic operators, a small set of arithmetic functions, and (because it is so widely used in web scripting) pseudo-random number generation. In this chapter, we round out this coverage by enumerating the built-in mathematical constants; exploring trigonometric, logarithmic, and base conversion functions, and explaining PHP's bc module for arbitrary-precision arithmetic.

Mathematical Constants

When we wrote the first edition of this book (around the release of PHP version 4.0), there was only one documented math constant: `M_PI` (the value of pi as a double). However, many new constants were introduced with PHP v4.0.2. Most of these new constants had to do with pi (or multiples thereof), e (or multiples thereof), or square roots, with a few odd-balls thrown in. The list has since shrunk back down to a slightly smaller number of predefined mathematical constants, for a variety of reasons. Those that remain are listed in Table 27-1. The general naming scheme is `M_<constant-name>`. In cases where the constant is a ratio (x/y), the name is `M_X_Y`, and in cases where there is an operation on a number, the name is `M_OPERNUM` (for example, `M_SQRT2`).

TABLE 27-1

Mathematical Constants

Constant	Description
`M_PI`	Pi
`M_PI_2`	pi/2
`M_PI_4`	pi/4
`M_1_PI`	1/pi
`M_2_PI`	2/pi
`M_2_SQRTPI`	2/sqrt(pi)
`M_E`	the constant e
`M_SQRT2`	sqrt(2)
`M_SQRT1_2`	1/sqrt(2)
`M_LOG2E`	$\log_2(e)$
`M_LOG10E`	$\log_{10}(e)$
`M_LN2`	$\log_e(2)$
`M_LN10`	$\log_e(10)$

Tests on Numbers

PHP offers a handful of functions for doing tests on numbers. Despite PHP's *type looseness*, it's a good idea to employ some of these tests in your code to help anticipate what sorts of results you will get, and how best to handle them.

The first and simplest test is `is_numeric()`. Like most of these tests, `is_numeric` returns a Boolean result, `True` if the supplied parameter is any type of number (signed or unsigned, integer or float) or a mathematical expression that returns a valid number:

```
is_numeric(4) // True
is numeric(4-4) // True
is_numeric(4*4) // True
```

Some caution is warranted, because even if you intend to test a string, a string that could be seen as an algebraic expression by PHP might also return a `True` value:

```
is_numeric(bells/4) // True
```

Finally, remember not to inadvertently put quotation marks around a mathematical expression, as the result will be forced to `False` by the string indicators, even though quoting a simple number or double will not cause this behavior:

```
is_numeric('M_PI * 3') // False
is_numeric('123456') // True
```

In cases like the preceding code snippet, it may be desirable to test with a higher degree of specificity, that is to say, for one of PHP's numeric subtypes, using `is_int()` or `is_float()`. Neither of these tests is substantively different from `is_numeric()` in their mode of operation. They simply test for a particular numeric type. Here are more usage examples and what you might expect from each:

```
is_int(4) // True
is_int(4.2) // False, it's a float or double
is_int('4') // False, this test is stricter than is numeric
is_int(4 * 2) // True, this expression yields an integer
```

You might also occasionally see the test `is_long()`. This test simply maps to the `is_int()` function. The other numeric type in PHP can be tested for using `is_float()`, which is aliased by `is_double()`. Again, usage is not exactly tricky or even complex.

```
is_float(4) // False, but you knew that already
is_float(4.212) // True, but you knew that as well
is_float(4 / 3) // True
is_float(M_PI) // True, maybe you knew that, maybe you didn't
```

Two more tests are slightly more obscure: `is_finite()` and `is_infinite()` test for exactly what their names suggest, although strictly speaking, their range is governed not by actual infinity (how would you test that?) but by the boundaries of the float value that your system allows.

Finally, we have `is_nan()`, which we covered in Chapter 9. You might be tempted to use `is_nan()` to test for any non-numeric value. You'll be unpleasantly surprised if you rely on this functionality. The more appropriate use of this function is to test for an unreasonable or improbable mathematical expression such as `acos(2)`.

Base Conversion

The default base in PHP for reading in or printing out numbers is 10. In addition, you can instruct PHP to read octal numbers in base 8 (by starting the number with a leading 0) or hexadecimal numbers in base 16 (by starting the number with a `0x`).

CAUTION   **For more on read formats of numbers, including octal and hexadecimal notation, see Chapter 54.**

A Glimpse behind the Curtain

How are built-in PHP functions really implemented? This is likely to be of interest only to C programmers and/or those who care about the inner workings of PHP, but we thought that it might be revealing to see why so many PHP functions work just like their C counterparts.

What follows is the actual implementation for the PHP function `ceil`, which is intended to convert a double to the smallest integer that is greater than or equal to it.

```
PHP_FUNCTION(ceil)
{
  zval **value;

  if (ARG_COUNT(ht)!=1||getParametersEx(1,&value)==FAILURE) {
              WRONG_PARAM_COUNT;
  }
    convert_scalar_to_number_ex(value);

    if ((*value)->type == IS_DOUBLE) {
      RETURN_LONG((long)ceil((*value)->value.dval));
        } else if ((*value)->type == IS_LONG) {
          RETURN_LONG((*value)->value.lval);
        }
    RETURN_FALSE;
}
```

Although the capitalized portions (including the `PHP_FUNCTION` declaration) are macros specific to the PHP framework, much of the body of this code is straight C. The code might look dense and confusing at first, but most of the action has to do with PHP's special treatment of types. Here is roughly what is happening, in order:

1. The arguments that `ceil()` was called with are retrieved and counted — if the count is anything other than 1, the function call returns with an error.

2. The single argument is converted to a number if it is a scalar type other than a number — this handles the possibility of string arguments, as in `ceil("5.4")`.

3. Now the numerical argument is tested to see whether it is a PHP long (aka integer) or a PHP double. If it turns out to be a long, the value as a long is returned.

4. The interesting case is if the value is a PHP double. If so, the C double value is extracted, it is run through the C function ceil, the result is converted to a C long, then that value is wrapped up and returned as a PHP long.

In other words, the PHP implementation of `ceil` is simply the C function `ceil`, wrapped up in a lot of type conversion and argument checking. This is the case with many of PHP's functions that have exact analogues in C.

Once numbers are read in, of course, they are represented in binary format in memory, and all the basic arithmetic and mathematical calculations are carried out internally in base 2. PHP also has a number of functions for translating between different bases, which are summarized in Table 27-2.

TABLE 27-2

Base Conversion Functions

Function	Behavior
BinDec()	Takes a single string argument representing a binary (base 2) integer and returns a string representation of that number in base 10.
DecBin()	Like BinDec(), but converts from base 10 to base 2.
OctDec()	Like BinDec(), but converts from base 8 to base 10.
DecOct()	Like BinDec(), but converts from base 10 to base 8.
HexDec()	Like BinDec(), but converts from base 16 to base 10.
DecHex()	Like BinDec(), but converts from base 10 to base 16.
baseconvert()	Takes a string argument (the integer to be converted) and two integer arguments (the original base, and the desired base). Returns a string representing the converted number — digits higher than 9 (from 10 to 35) are represented by the letters a–z. Both the original and desired bases must be in the range 2–36.

All the base-conversion functions are special-purpose, converting from one particular base to another, except for base_convert(), which accepts an arbitrary start base and destination base. Here's an example of base_convert() in action:

```
function display_bases($start_string, $start_base)
{
  for ($new_base = 2; $new_base <= 36; $new_base++)
    {
      $converted =
        base_convert($start_string, $start_base, $new_base);
       print("$start_string in base $start_base
             is $converted in base $new_base<BR>");
    }
}

display_bases("1jj", 20);
```

This code yields the browser output:

```
1jj in base 20 is 1100011111 in base 2
1jj in base 20 is 1002121 in base 3
1jj in base 20 is 30133 in base 4
```

459

```
1jj in base 20 is 11144 in base 5
1jj in base 20 is 3411 in base 6
1jj in base 20 is 2221 in base 7
1jj in base 20 is 1437 in base 8
1jj in base 20 is 1077 in base 9
1jj in base 20 is 799 in base 10
1jj in base 20 is 667 in base 11
1jj in base 20 is 567 in base 12
1jj in base 20 is 496 in base 13
1jj in base 20 is 411 in base 14
1jj in base 20 is 384 in base 15
1jj in base 20 is 31f in base 16
1jj in base 20 is 2d0 in base 17
1jj in base 20 is 287 in base 18
1jj in base 20 is 241 in base 19
1jj in base 20 is 1jj in base 20
1jj in base 20 is 1h1 in base 21
1jj in base 20 is 1e7 in base 22
1jj in base 20 is 1bh in base 23
1jj in base 20 is 197 in base 24
1jj in base 20 is 16o in base 25
1jj in base 20 is 14j in base 26
1jj in base 20 is 12g in base 27
1jj in base 20 is 10f in base 28
1jj in base 20 is rg in base 29
1jj in base 20 is qj in base 30
1jj in base 20 is po in base 31
1jj in base 20 is ov in base 32
1jj in base 20 is o7 in base 33
1jj in base 20 is nh in base 34
1jj in base 20 is mt in base 35
1jj in base 20 is m7 in base 36
```

Notice that although all the base-conversion functions take string arguments and return string values, you can use decimal numerical arguments and rely on PHP's type conversion (but see the cautionary note that follows). In other words, both `DecBin("1234")` and `DecBin(1234)` will yield the same result.

CAUTION **Don't confuse the read formats of numbers with their representations as strings for the purposes of base conversion. For example, although 10 in base 16 is the number 16 in base 10, the expression `HexDec(0x10)` evaluates to the string `"22"`. Why? There are really three conversions happening: when `0x10` is read (converts from hex to internal binary), when the argument is auto-converted (from internal binary number to the decimal string `"16"`), and in the operation of the function (from assumed base 16 to decimal `"22"`). If you want just one conversion, the desired expression is `HexDec("10")`.**

CAUTION **The base conversion functions expect their string arguments to be integers, not floating-point numbers. That means you can't use these functions to convert a binary 10.1 to a decimal 2.5.**

Exponents and Logarithms

PHP includes the standard exponential and logarithmic functions, in both base 10 and base e varieties (shown in Table 27-3).

TABLE 27-3

Exponential Functions

Function	Behavior
pow()	Takes two numerical arguments and returns the first argument raised to the power of the second. The value of pow($x, $y) is x^y.
exp()	Takes a single argument and raises e to that power. The value of exp($x) is e^x.
log()	The "natural log" function. Takes a single argument and returns its base e logarithm. If $e^y = x$, then the value of log($x) is y.
log10()	Takes a single argument and returns its base-10 logarithm. If $10^y = x$, then the value of log10($x) is y.

Unlike with exp() and the base e, there is no single-argument function to raise 10 to a given power, but in its place you can use the two-argument function pow() with 10 as the first argument.

You can verify that exponential and power functions of the same base are inverses of each other, by testing an identity like this:

```
$test_449 = 449.0;
$test_449 = pow(10, exp(log(log10($test_449))));
print("test_449 is $test_449<BR>");
```

which gives the browser output:

```
test_449 is 449
```

Trigonometry

Although explaining the math behind the PHP functions in this chapter is beyond the scope of this book, we've made an exception just this once. (See the sidebar "Trigonometry in One Paragraph." Anyone who doesn't already know trigonometry will, of course, find the sidebar completely impenetrable, but we hope that those who know trig will at least be amused by how short it is.)

PHP offers the standard set of basic trigonometric functions as well as the constant M_PI, an approximation of pi as a double that prints as 3.1415926535898. This constant can be used anywhere you

would use the literal number itself, and it is also interchangeable with the pi() function. (For other constants derived from pi, see the "Mathematical Constants" section at the beginning of this chapter.) Both of the following statements have the same result:

```
$my_pi = M_PI;
$my_pi = pi();
```

The basic trig functions are summarized in Table 27-4.

TABLE 27-4

Trigonometric Functions

Function	Behavior
pi()	Takes no arguments and returns an approximation of pi (3.1415926535898). Can be used interchangeably with the constant M_PI.
Sin()	Takes a numerical argument in radians and returns the sine of the argument as a double.
Cos()	Takes a numerical argument in radians and returns the cosine of the argument as a double.
Tan()	Takes a numerical argument in radians and returns the tangent of the argument as a double.
Asin()	Takes a numerical argument and returns the arcsine of the argument in radians. Inputs must be between –1.0 and 1.0 [inputs outside that range will return a result of NAN (for "not a number")]. Results are in the range -pi / 2 to pi / 2.
Acos()	Takes a numerical argument and returns the arccosine of the argument in radians. Inputs must be between –1.0 and 1.0 [inputs outside that range will return a result of NAN (for "not a number")]. Results are in the range 0 to pi.
Atan()	Takes a numerical argument and returns the arctangent of the argument in radians. Results are in the range -pi / 2 to pi / 2.
Atan2()	A variant of atan() that takes two arguments. Atan($y, $x) is identical to atan($y/$x) when $x is positive, but the quadrant of atan2's result depends on the signs of both $y and $x. Range of the result is from -pi to pi.

Rather than writing down a table of sample function results, let's resort to our usual trick of writing code that automatically displays examples as an HTML table. Listing 27-1 shows both a generalized function for displaying a set of one-argument functions applied to a set of numerical arguments and then the result of using this display function to make trigonometric example tables. The results are displayed in Figure 27-1.

Trigonometry in One Paragraph

Imagine a circle with a radius of 1, centered at 0,0 in the coordinate plane. Start at the right-hand edge (at position (1,0)), and trace a certain distance along the circle counterclockwise. For example, a distance of 2 pi would take you once around the circle and back to your starting point. Clockwise travel counts as a negative distance. For any such distance, the sine function tells you the y-value of the coordinate you arrive at, the cosine function tells you the x-value of that coordinate, and the tangent function is a ratio of the two, from which you can infer the slope of the line tangent to the circle at that point. The functions arccosine, arcsine, and arctangent are in some sense inverses of their corresponding functions — they map back from an x, y, or y/x ratio to the distance of a circular trip that would arrive at that x-coordinate, y-coordinate, or ratio thereof. Because adding a multiple of 2 pi to any distance brings you around to the same point again, these inverse functions might have an infinite number of answers per input, making them ill-defined — instead, they are restricted to a range corresponding to one particular trip around half of the circle and so have well-defined results.

LISTING 27-1

Displaying trigonometric function results

```php
<?php

function display_func_results($func_array, $input_array)
{
/* print a function header */
print("<TABLE BORDER=1><TR><TH>INPUT\\FUNCTION</TH>");
for($y = 0;
        $y < count($func_array);
        $y++)
    print("<TH>$func_array[$y]</TH>");
print("</TR><TR>");
/* print the rest of the table */
for($x = 0;
    $x < count($input_array);
    $x++)
  {
    /* print column entries for inputs */
    print("<TH>".
        sprintf("%.4f", $input_array[$x])
        ."</TH>");
    for($y = 0;
        $y < count($func_array);
        $y++)
    {
      $func_name = $func_array[$y];
      $input = $input_array[$x];
```

```
      print("<TD>");
      printf("%4.4f", $func_name($input));
      print("</TD>");
    }
   print("</TR><TR>");
  }
print("</TR></TABLE>");
}
?>

<HTML>
<HEAD>
<TITLE>Trigonometric Function Examples</TITLE>
</HEAD>
<BODY>

<?php
/* using the function displayer */
print("<H3>Trigonometric function examples</H3>");
display_func_results(array("sin", "cos", "tan"),
                     array(-1.25 * pi(),
                           -1.0 * pi(),
                           -0.75 * pi(),
                           -0.5 * pi(),
                           -0.25 * pi(),
                            0,
                            0.25 * pi(),
                            0.5 * pi(),
                            0.75 * pi(),
                            pi(),
                            1.25 * pi()));

?>
</BODY>
</HTML>
```

CAUTION   The `display_func_results()` **function of Listing 27-1 uses several tricks we've seen in previous chapters: using a string variable as the name of a function to call (covered near the end of Chapter 26) and using the string concatenation operator (`.`) to pull together a print string in the middle of a print statement (covered in Chapter 7).**

Figure 27-1 shows the basic trigonometric functions over an input range of `-5/4 pi` to `5/4 pi` and the basic inverse trigonometric function over inputs from `-1.0` to `1.0`. The very large tangent values are due to denominators that should theoretically be zero but instead differ slightly from zero due to rounding error.

FIGURE 27-1

Trigonometric function examples

Arbitrary Precision (BC)

The integer and double types are fine for most of the mathematical tasks that arise in web scripting, but each instance of these types is stored in a fixed amount of computer memory, and so the size and precision of the numbers these types can represent is inherently limited. Although the exact range of these types may depend on the architecture of your server machine, integers typically range from $-2^{31} - 1$ to $2^{31} - 1$, and doubles can represent about 13 to 14 decimal digits of precision. For tasks that require greater range or precision, PHP offers the *arbitrary-precision* math functions (also known as *BC functions*, from the name of the Unix-based, arbitrary-precision calculating utility).

CAUTION **Especially if you compiled PHP yourself, the arbitrary-precision functions may not have been included in the compilation — you need to have included the flag `--enable-bcmath` at configuration time. To check whether the functions are present, try evaluating `bcadd("1", "1")` — if you get an unbound function error, you will have to reconfigure and recompile PHP.**

Instead of using the fixed-length numerical types, the BC functions have strings as arguments and return values. Because strings in PHP are limited only by available memory, numbers can be as long as you like. The underlying computations are performed in decimal and are done much as you would do them with pen and paper (if you were very fast and very patient). When operating with

integers, the BC functions are exact and use as many digits as needed; when operating with floating-point numbers, computations are done to as many decimal places as you specify. The BC functions are summarized in Table 27-5.

Most of the functions take an optional scale factor (an integer) as a final argument, which determines how many decimal places will be in the result. If such an argument is not supplied, the scale is the default scale, which, in turn, can be set by calling `bcscale()`. The default for the default value (that is, if `bcscale()` has never been called) can also be set in the initialization file `php.ini`.

TABLE 27-5

Arbitrary-Precision (BC) Math Functions

Function	Behavior
bcadd()	Takes two string arguments representing numbers, and an optional integer scale parameter. Returns the sum of the first two arguments as a string, with the number of decimal places in the result determined by the scale parameter. If no scale parameter is supplied, the default scale is used (which is settable by `bcscale()`).
bcsub()	Similar to `bcadd()`, except that it returns the subtraction of the second argument from the first.
bcmul()	Similar to `bcadd()` but returns the product of its arguments.
bcdiv()	Similar to `bcadd()` but returns the result of dividing the first argument by the second.
bcmod()	Returns the modulus (remainder) of the first argument as divided by the second. Because the return type is "integral," no scale argument is taken.
bcpow()	Raises the first argument to the power of the second argument. The number of decimal places in the result is set by the scale factor if supplied.
bcsqrt()	Returns the square root of its argument, with number of decimal places set by the optional scale factor.
bcscale()	Sets the default scale factor for subsequent BC function calls.

An arbitrary-precision example

Here's an example of using the arbitrary-precision functions for exact integer arithmetic. The following code:

```
for ($x = 1; $x < 25; $x++) {
  print("$x raised to the power of $x is " . bcpow($x, $x) . "<BR>");
}
```

will print like this:

```
 1 raised to the power of 1 is 1
 2 raised to the power of 2 is 4
 3 raised to the power of 3 is 27
 4 raised to the power of 4 is 256
 5 raised to the power of 5 is 3125
 6 raised to the power of 6 is 46656
 7 raised to the power of 7 is 823543
 8 raised to the power of 8 is 16777216
 9 raised to the power of 9 is 387420489
10 raised to the power of 10 is 10000000000
11 raised to the power of 11 is 285311670611
12 raised to the power of 12 is 8916100448256
13 raised to the power of 13 is 302875106592253
14 raised to the power of 14 is 11112006825558016
15 raised to the power of 15 is 437893890380859375
16 raised to the power of 16 is 18446744073709551616
17 raised to the power of 17 is 827240261886336764177
18 raised to the power of 18 is 39346408075296537575424
19 raised to the power of 19 is 1978419655660313589123979
20 raised to the power of 20 is 104857600000000000000000000
21 raised to the power of 21 is 5842587018385982521381124421
22 raised to the power of 22 is 341427877364219557396646723584
23 raised to the power of 23 is 20880467999847912034355032910567
24 raised to the power of 24 is 1333735776850284124449081472843776
25 raised to the power of 25 is 88817841970012523233890533447265625
```

If we had used the regular PHP integer type for this computation, the integers would have *overflowed* well before the end, and the rest of the loop would have been calculated in approximate floating point.

Converting code to arbitrary-precision

Let's see what it's like to take an existing piece of mathematical code and retrofit it to use the arbitrary-precision functions.

The following function approximates pi, using the series approximation:

$$\text{sqrt} \ (\ 12 - (12/2^2) + (12/3^2) - (12/4^2) + (12/5^2) - ...)$$

(As we'll see, this series does not converge fast enough for our purposes, but it has the virtue of being a simple formula.)

```
function pi_approx($iterations, $print_frequency)
{
  $squared_approx = 12;
  $next_sign = -1;
  $denom = 2;
```

```
    for ($iter = 0; $iter < $iterations; $iter++)
      {
        $squared_approx += $next_sign * 12/(pow($denom,2));
        $denom++;
        $next_sign = - $next_sign;
        if ($denom % $print_frequency == 0)
          {
            $estimate = sqrt($squared_approx);
            print("$denom iterations: $estimate<BR>");
          }
      }
  }
```

In addition to performing the calculation itself, this code periodically prints its current estimate of pi, so we can see how we are doing. We can call it as follows and then print PHP's value for comparison:

```
pi_approx(10000, 1000);
print("PHP value: " . pi() . "<BR>");
```

The result looks like:

```
1000 iterations: 3.1415936094742
2000 iterations: 3.1415928924416
3000 iterations: 3.1415927597285
4000 iterations: 3.1415927132878
5000 iterations: 3.1415926917946
6000 iterations: 3.14159268012
7000 iterations: 3.141592673081
8000 iterations: 3.1415926685124
9000 iterations: 3.1415926653804
10000 iterations: 3.1415926631401
PHP value: 3.1415926535898
```

Now, not only are we not that close, but we can't hope to be more accurate than PHP's value for pi, because that already uses all the precision available in the double type.

To convert this to an arbitrary-precision version, we must replace all the math functions and operators that need precision with their BC counterparts, like this:

```
function pi_approx_bc($iterations, $print_frequency, $scale)
{
  $squared_approx = "12";
  $next_sign = -1;
  $denom = 2;
  for ($iter = 0; $iter < $iterations; $iter++)
    {
      $squared_approx
        = bcadd(
```

```
                        $squared_approx,
                        bcmul($next_sign,
                            bcdiv(12,
                                bcpow($denom,
                                    2,
                                    $scale),
                                $scale),
                            $scale),
                        $scale);
                $denom++;
                $next_sign = - $next_sign;
                if ($denom % $print_frequency == 0)
                    {
                        $estimate = bcsqrt($squared_approx,$scale);
                        print("$denom iterations: $estimate<BR>");
                    }
            }
    }
```

Notice that although the BC functions want string arguments, we can, as always, use regular numbers in their places and rely on PHP to convert the arguments to strings for us. Also notice that we did not bother making the numerical computations that do not require great precision into BC computations (for instance, we still have $denom++ rather than bcadd($denom, 1). Finally, we added a scale argument to the entire function, which turns the decimal precision of each BC function it calls.

Unfortunately, both your authors and our browsers ran out of patience with this series before it even got to the level of precision of PHP's value. Here are some late results of calling pi_approx_ bc(1250000, 50000, 50):

```
50000 iterations: 3.14159265397177274129723551068347726371297686926596
100000 iterations: 3.14159265368528715924598769254390594927146205337113
150000 iterations: 3.14159265363223483956231649503922204272933217538202
[..]
1150000 iterations: 3.14159265359051530310455255409602580003265549955003
1200000 iterations: 3.14159265359045638461148087403458918944405547147211
1250000 iterations: 3.14159265359040439393304018710072157703501022388304
```

The correct digits in the preceding output are about one digit shy of the PHP value. This is the fault of the series we chose rather than the arbitrary-precision libraries — with a more sophisticated and speedier approximation series, you too can serve up millions of digits of pi to your eager audience.

Somewhat more satisfyingly, evaluating:

```
print("The square root of two is " . bcsqrt(2, 40));
```

gives us many more digits of precision than we could get using doubles:

```
The square root of two is 1.4142135623730950488016887242096980785696
```

469

Summary

Although the primary purpose of PHP is not to do mathematics, it has a pretty comprehensive set of mathematical functions covering basic arithmetic, pseudo-random number generation, base conversion, trigonometry, exponents and logarithms, and a built-in module for doing arbitrary-precision arithmetic.

We covered the numerical types and the most basic functions in Chapter 9, and covered the remaining topics in this chapter. Table 27-6 is a tabular summary of the operators and functions discussed both in Chapter 9 and this chapter.

TABLE 27-6

Summary of PHP Math Operators and Functions

Category	Description
Arithmetic operators	Operators +, -, *, /, % perform basic arithmetic on integers and doubles.
Incrementing operators	The ++ and -- operators change the values of numerical variables, increasing them by one or decreasing them by one (respectively). The value of the postincrement form ($var++) is the same as the variable's value before the change; the value of the preincrement form (++$var) is the variable's value after the change.
Assignment operators	Each arithmetic operator (like +) has a corresponding assignment operator (+=). The expression $count += 5 is equivalent to $count = $count + 5.
Comparison operators	These operators (<, <=, >, >=, ==, !=) compare two numbers and return either true or false. The === operator is true if and only if its arguments are equal and of the same type.
Basic math functions	floor(), ceil(), and round() convert doubles to integers, min() and max() take the minimum and maximum of their numerical arguments, and abs() is the absolute value function.
Base conversion functions	Special-purpose functions (OctDec(), DecOct(), BinDec(), DecBin(), HexDec(), DecHex()) convert between particular pairs of bases, whereas base_convert() translates between arbitrary bases.
Exponential functions	Functions having to do with raising numbers to powers or the inverse: log() (natural log), log10() (base-10 log), exp() (e raised to the power of the argument), and pow() (first argument to the power of the second).
Trigonometric functions	Functions having to do with angular measures: pi() (and the equivalent constant M_PI), sin(), cos(), tan(), acos(), asin(), atan(), and atan2() (a two-argument version of atan()).
Arbitrary-precision (BC) functions	Functions that do arithmetic on arbitrary-length strings representing decimal integers and floating-point numbers: bdadd(), bcsub(), bcmult(), bcdiv(), bcmod(), bcpow(), bcsqrt(). Most of these functions take an optional scale parameter specifying the number of decimal points of precision desired — the default for that parameter is settable using bcscale().

Chapter 28

Securing PHP

"**S**ecurity is not a joking matter," proclaim signs at airports everywhere. The same sign should be posted near your PHP server. Anyone connecting a server to the Internet must take proper security measures or risk loss of data or even money to the keystrokes of malicious crackers.

The mantra of the security-conscious site designer is: *Don't trust the network.* If you're worried about the security of your site, chant this mantra as you code your pages. Any information transmitted to your server via the network — be it a URL, data from an HTML form, or data on some other network port — should be treated as potentially hazardous. This chapter suggests several techniques for *sanitizing* incoming information. You should apply these techniques and spend some time trying to discover other potential hazards and ways to prevent them.

The second rule of thumb for a secure site is: *Minimize the damage.* What if the program you just wrote, which you are *sure* is secure, is actually vulnerable? Just to be on the safe side, limit the damage an intruder can cause after he or she has taken advantage of the vulnerability.

When visitors come to your site, they trust that it contains valid information, that it is not harmful to them or to their computers, and that any information they provide to it is handled properly. Interacting with a site, whether an e-business, recreational, or informational site, involves certain security risks for a visitor. As a site designer, it is your responsibility to protect visitors from these risks. Besides being sure their information is safe on your server, this means you should take measures to safeguard their information while it is in transit from their computers to your server.

But all this should not scare you away from putting your e-business online. The first section of this chapter describes some possible attacks against your server and ways to avoid them. We then discuss cryptographic techniques for protecting your data. At the end of this chapter, we list some web sites that contain up-to-the-minute information on the latest cracker techniques. By watching these sites, you may learn of possible security vulnerabilities before an attacker does and, thereby, avoid disaster.

Possible Attacks

Connecting your server to the Internet is like setting up a storefront on a busy street. You're likely to have quite a few visitors, but if you're not careful, some less than desirable visitors may take advantage of you.

Crackers, script-kiddies, and other fiends

The term *hacker* is commonly used to describe individuals more correctly labeled crackers. Within the computer community, *crackers* are those who, through luck or skill, break into computer systems and cause damage. Hackers are those who can *hack* — read and write efficient (and often obscure) code in many languages. To a programmer, being labeled a hacker is an honor, whereas being labeled a cracker probably means he or she should start reading the Help Wanted section.

As if the term *cracker* were not sufficiently derogatory, young crackers who use tools and scripts they find on the web are called *script-kiddies*. These budding lawbreakers often have little understanding of what they are actually doing. They are usually the culprits behind low-tech attacks such as site defacement. A fairly good indicator of the work of a script-kiddie is the excessive use of misspelling and capitalization, as in W3 R K00L D00Dz.

Site defacement

Often more embarrassing than harmful, site defacements are fairly common because the cracker has an opportunity to publicize his or her exploitation. Site defacements are sometimes left as calling cards by a cracker who entered a system by more complicated means.

It is possible to deface a badly designed web site using only a web browser. Take, for instance, the following program:

```php
<?php if (IsSet($_POST['visitor'])) {
  $visitor = $_POST['visitor'];
  $fp = fopen("database", "a");
  fwrite($fp, "<li>$visitor\n");
  fclose($fp); } ?>
<HTML>
<HEAD></HEAD>
```

```
<BODY>
<H1>Visitors to this site:</H1>
<OL>
<?php $fp = fopen("database", "r");
  print(fread($fp, filesize("database")));
  fclose($fp) ?>
</OL>
<HR>
<FORM><INPUT TYPE="TEXT" NAME="visitor">
<INPUT TYPE="SUBMIT" NAME="submit" VALUE="Sign in!">
</FORM>
</BODY>
</HTML>
```

This program implements a very rudimentary guest book. In reading this code, however, you should feel a bit uneasy. *Don't trust the network*. This program accepts form data that we expect to contain the visitor's name (in the variable $visitor) and stores it in a text file for display to subsequent visitors. For the inputs we expect, there is no trouble.

Now put on your script-kiddie hat for a moment and imagine what would happen if the input contained HTML tags. This simple program would blindly insert those tags into the pages it generates, and other visitors' browsers would interpret them as usual. One particularly malicious tag is the <SCRIPT> tag. A cracker wishing to deface this web page could duplicate the page's appearance on his or her own server (www.example.com) and then sign into the guestbook with the name:

```
<SCRIPT LANGUAGE="JavaScript">
window.location="http://www.example.com/"</SCRIPT>
```

When visitors load the guest book, their browsers receive this tag and immediately begin loading the hacked site. With a little ingenuity, the cracker could then take advantage of the visitors' trust of your site to extract personal information such as passwords or credit card numbers.

The solution to this problem is to sanitize the input data. In this case, we want any characters that have special meaning to a browser to be translated into something harmless. Luckily, PHP provides a way to perform just such a translation. The function htmlspecialchars() converts the characters <, >, ", and & to their representations as HTML entities (such as <). We change the first part of our program to use this new function as follows:

```
<?php if (IsSet($_POST['visitor'])) {
  $visitor = $_POST['visitor'];
  $fp = fopen("database", "a");
  $clean_visitor = htmlspecialchars($visitor);
  fwrite($fp, "<li>$clean_visitor\n");
  fclose($fp); } ?>
```

And we have patched a very significant security hole in our site.

Accessing source code

Even if your PHP source code isn't a trade secret, you should still protect it from exposure to the network. If an intruder can read your source code, then he or she need not experiment to find a weakness. Instead, the intruder can simply analyze the code, looking for common mistakes and other security holes. In general, the more helpful information you provide to potential intruders, the more likely an intrusion. By hiding such tidbits as source code, directory names, or usernames from the network, you can reduce the likelihood of an attack.

CAUTION One handy feature of PHP, error reporting to the browser, is great for development because it helps pinpoint problems — but it can be bad for security, because it can also give directory paths, filenames, usernames, and potentially database names on error. Minimize the risk by turning off error reporting to the browser in production systems, via the `display_errors` directive in `php.ini`. You can still use error reporting to the browser on development systems if you wish, although it's safer to use the `error_log()` function to write error messages to a log.

When PHP is used as a web server module, there is little risk of source code being released by the web server, as any file with the proper extension is parsed by the PHP module. If PHP is installed as a CGI program, however, things are not so simple.

If you cannot run PHP as a server module, the next most secure setup is to run it as an interpreter for CGI scripts, just as you would Perl or Python.

Place all your PHP programs in the `cgi-bin` directory for your server or your account and arrange for the PHP interpreter to be invoked when they are executed. On Unix, this is done by adding a line similar to the following as the first line of every script:

```
#! /usr/local/bin/php
```

To use this setup, you must compile PHP with the `--enable-discard-path` configuration option. This setup has the disadvantage that the URLs for most of your pages contain `/cgi-bin/`.

The next most secure setup is a bit more complicated and is actually counter to the recommendations of CERT, a respected authority on computer security: We place the PHP interpreter itself in the `cgi-bin` directory. It is usually inadvisable to put an interpreter in the `cgi-bin` directory, because the rules for invoking CGI programs would allow any file on the server to be parsed as a program.

PHP is written to operate safely from the `cgi-bin` directory, however, if configured correctly. If you intend to use this setup, first carefully read the security and configuration sections of the PHP manual, as they may contain important information not available as this book went to press.

This setup relies upon the web server to redirect URLs of the form:

```
http://your.server/program.php
```

to URLs of the form:

```
http://your.server/cgi-bin/php/program.php
```

The precise directives that will cause your web server to do this vary. For Apache they are:

```
Action php-script /cgi-bin/php
AddType php-script .php
```

If you are using Apache, be sure to compile PHP with the `--enable-force-cgi-redirect` configuration option. This option utilizes a feature specific to Apache to prevent PHP from executing when invoked by URLs of the second form. Your setup is complete.

If you are using any other server software, you must compile PHP with the `--disable-force-cgi-redirect` configuration option. PHP cannot distinguish the two types of URLs and serves a document of either type. This allows a visitor to view files without regard for web-server-based access restrictions. Assume, for example, that the URL `www.example.com/top/secret/hush.php` has access restrictions placed on it. A cracker could use the URL `www.example.com/cgi-bin/php/top/secret/hush.php` to read the file anyway.

In this case, the web server is giving PHP the path name `/top/secret/hush.php`. PHP determines the location of the program file by prepending the configuration value `doc_root` to the given path name. By default, this value is the same as the web server's document root (the directory corresponding to `www.example.com/`). Setting `doc_root` to another directory will limit PHP to programs in that directory and its subdirectories instead of the entire collection of web server documents. Any visitor may access any of the PHP programs by the method just described, however, without regard for web-server-based access controls. *Be careful!*

Reading arbitrary files

A few common PHP programming mistakes can make it easy for a hacker to read almost any file on the server. Study the following page:

```
<HTML>
<HEAD></HEAD>
<BODY>
<?php if (IsSet($_POST['poem'])) {
  $poem = $_POST['poem'];
  $fp = fopen($poem, "r");
  print (fread($fp, filesize($poem)));
  fclose($fp);
  } ?>
<HR><FORM>Pick a poem:
<SELECT NAME="poem"><OPTION VALUE="jabb.html">Jabberwocky
<OPTION VALUE="graves.html">Cat-Goddesses</SELECT>
<INPUT TYPE="SUBMIT" VALUE="Show Me"></FORM>
</BODY>
</HTML>
```

This simple program displays a number of poems, selectable from a pop-up menu given in the form near the end. Invoke the security mantra: *Don't trust the network.* Clicking Show Me on this page

475

results in URLs such as `poetry.php?poem=graves.html`. A cracker may substitute the filename of some more sensitive file, such as `poetry.php?poem=/etc/passwd`. The program, as given, would dutifully serve up the Unix password file, possibly enabling the cracker to break into a visitor account and do further damage.

The following is an appropriate solution to this problem:

```php
<?php if (IsSet($_POST['poem'])) {
  $poem = $_POST['poem'];
  switch ($poem) {
    case "jabb":
      $poem_file = "jabb.html";
      break;
    case "graves":
      $poem_file = "graves.html";
      break;
    }
  if (IsSet($_POST['poem_file'])) {
    $poem_file = $_POST['poem_file'];
    $fp = fopen($poem_file, "r");
    print (fread($fp, filesize($poem_file)));
    fclose($fp);
    }
  } ?>
```

The advantage of this method is that it explicitly lists the acceptable inputs and gracefully handles unacceptable inputs. If there were more poems to be processed, the `switch` statement could be replaced with a database query, where failure of the query indicates invalid input.

This is *not* a good solution:

```php
<?php if (IsSet($_POST['poem'])) {
  $poem = $_POST['poem'];
  if (!strstr($poem, "/") && !strstr($poem, "\\")) {
    $fp = fopen($poem, "r");
    print (fread($fp, filesize($poem)));
    fclose($fp);
    }
  } ?>
```

The second conditional in this code segment checks for pathname separators in the given filename. This program explicitly describes a set of unacceptable inputs and considers anything else acceptable. It depends on the programmer imagining and checking for every possible undesired input. In this case, the programmer has missed something by making the implicit assumption that no sensitive files are stored in the same directory as the script.

What if a file that should be private escapes your server anyway? There is a chance that some misconfiguration (perhaps by someone else) or an unnoticed security hole will render some or all of your server's files publicly accessible.

PHP allows you to explicitly specify the set of directories in which files can be opened with the configuration value `open_basedir`. See Chapter 29 for more information on the PHP configuration file. This configuration value can be useful to prevent access to entire directories and is a good way to *minimize the damage.*

Many sensitive files, however, must be opened from PHP programs as visitors access the site. A common example is a password file. Access to such a file cannot be blocked with `open_basedir`, but the sensitive information it contains can be encrypted to render it useless to anyone who may steal it.

A password-protected site must verify the password given by a visitor wishing to gain access. One way to do this would be to store a password for safekeeping in encrypted form and then decrypt it when we need to compare it to the user-supplied password. The problem is that if we can decrypt the password, others may be able to decrypt it too. Also, we would have to make sure that no one could see the password after we decrypted it for comparison. Instead, we can use an encryption function that only goes one way and is easy to use for encryption, but that can't be decrypted. Rather than decrypt a stored password and compare the decrypted versions, we *encrypt* the *given* password and compare the encrypted passwords. Unix uses this strategy with its own password file, `/etc/passwd`, and PHP allows programmers to use the same encryption function for their own password files.

The function `crypt(password, salt)` encrypts the given password. The salt adds an extra bit of chance and should be chosen randomly when the password is first recorded. (PHP chooses a random salt if this parameter is omitted.) The function returns the concatenation of the salt value and the encrypted version of the password. The following function will create a new password for a visitor:

```
function new_pw($given) {
  return crypt($given)
}
```

And this function will compare a password given by a visitor with a stored, encrypted password:

```
function verify_pw($given, $stored) {
  $salt = substr($stored, 0, CRYPT_SALT_LENGTH);
  $given_encrypted = crypt($given, $salt);
  return ($stored == $given_encrypted);
}
```

Running arbitrary programs

It's every system administrator's worst nightmare. The server's running more slowly than usual. A look at the running programs on the server reveals that a program entitled `crack` is burning 98 percent of the processor's time. Most likely, this program has been placed here by a cracker who is using it to decrypt (crack) passwords. The administrator logs in to kill the offending program but finds that his password is incorrect. His server has been *root compromised*, and there is no telling how much damage has been done.

Social Engineering

Social engineering is an often overlooked part of cracking. Sometimes it's easier for crackers to extract information (particularly passwords) from human beings than from computers:

Cracker: Hi, John, this is Gary in the IT department. When was the last time you used your company account?

John: Well, I entered a few new purchase orders about an hour ago.

Cracker: Well, John, I'm afraid your account has been compromised. Some of the information in it may have been lost. This could cost the company millions if we don't catch the intruder quickly. We need to open your account and assess the damage immediately. Can you give me your password?

John: Sure, it's . . .

Worse yet, sometimes forgetful visitors note their passwords on scraps of paper in their desks! A determined cracker can easily find a job as a night janitor and look for such notes. Many famous crackers were more notable for their social engineering and research skills than their ability to write code to compromise systems.

In a compromise such as this, an intruder gains interactive access to the server, usually via a Unix shell or MS-DOS command line. Clearly, this is the most difficult type of heist to pull off, but it also bears the greatest reward. Once *inside* a server, the cracker has virtually unlimited power to bring down the server, steal or modify information, or make use of the server's computational power for further wrongdoing. Worse yet, a truly skilled cracker can conceal his or her steps by editing log files and erasing any temporary files he or she has created.

PHP has several program execution functions: `system()`, `exec()`, `popen()`, `passthru()`, and the back-tick (`` ` ``) operator. As an example of the use of one of these functions, the following page returns the Unix `finger` information for a visitor specified through an HTML form:

```
<HTML>
<HEAD></HEAD>
<BODY>
<FORM>Get information on <INPUT TYPE="TEXT" NAME="username">,
<INPUT TYPE="SUBMIT" VALUE="Please"></FORM>
<?php if (IsSet($_POST['username'])) { ?>
  <H1>Results for <?php echo $_POST['username']; ?></H1>
  <pre><?php system("finger " . $_POST['username']); ?></pre>
<?php } ?>
</BODY>
</HTML>
```

The program, as given, takes a username from the HTML form and executes the Unix program `finger` to look up information about that user. You should hear *Don't trust the network* repeating loudly in your head. Unix commands are separated by a semicolon, so anything following a semicolon in the string passed to `system()` is treated as a new command. This new command is executed with all the permissions of the user under which the web server is running.

Under Unix, the command "`rm -rf /`" will delete all files on the server. Imagine the damage if an ill-intentioned visitor typed "`; rm -rf /`" into the form and clicked `Please`.

The best solution to this problem is to filter out everything but valid usernames before invoking `finger`. This requires specific knowledge about username formats on your server, so we do not present an example here. PHP presents a solution that is almost as good. The function `escapeshellcmd()` will sanitize a string for use in a program execution command, rendering harmless any special characters such as the semicolon. We replace the line invoking `system()` in the preceding code snippet with:

```
<pre><?php
    system(escapeshellcmd("finger " . $_POST['username']));
?></pre>
```

Magically, no value the visitor may enter can result in arbitrary programs being executed. This does not, however, prevent the visitor from providing unexpected input to `finger`. Although `finger` does no harm if given incorrect input, other programs may not be so forgiving. If in doubt, err on the side of caution!

To *minimize the damage* of a compromise of this sort, most modern web servers run as a dummy user (often called `nobody` on Unix systems). This user has only the permissions required to run the web server (and any PHP scripts) and read and write the necessary files. But remember, any databases or files that your scripts can modify are modifiable by this user, and thus they are vulnerable if an attacker can run arbitrary programs.

Viruses and other e-critters

Visitors trust software coming from a trusted site. If your site allows visitors to download files uploaded by other visitors, you should warn your visitors to check files for viruses before running them, and you should consider periodically scanning the files on your server for viruses as well. This is a hard problem to solve, particularly with the possibility of embedding viruses in such seemingly harmless files as word processor documents. Indeed, Microsoft was caught in this very bind when it inadvertently released a CD-ROM with a Word document containing the Melissa virus.

CROSS-REF See the section "Site defacement" at the beginning of this chapter for other ways that your visitor may inadvertently receive malicious code.

FYI: Security Web Sites

If you are losing sleep after reading this chapter, fear not. Every administrator and site designer around the world is grappling with the same issues, and there is a strong feeling of solidarity among computer security professionals. Many web sites are devoted to computer security, and almost all of them contain full descriptions of recent security incidents and ways to protect your system from duplicate attacks. Some are designed for security professionals, whereas others have the cracker in mind. Either way, the information they provide is useful and often very interesting.

System Administrators

System administrators, also called *sysadmins*, are the folks who make sure the computers we all use keep on computing and that the Internet keeps on networking. Their jobs are shrouded in mystery: They hold the keys to the mysterious "machine room" where all the critical servers are stored. It's not unusual to see them hurrying into the office at midnight, surely to avert some crisis that could bring the company to its knees.

Sysadmins are also a very cautious lot. They tend to program their servers to report any unusual activity immediately (often to the large-screen alphanumeric pager they carry at all times) and to take swift, decisive action against anything they deem improper or unsafe.

A professor in a Computer Science department once asked his students, as homework, to break into his Linux desktop. To make things a little easier, he gave the encrypted text of his password (see the description of `crypt()` in the section "Reading arbitrary files"). In a testament to the security of the Unix `crypt()` function, none of the students cracked his desktop. Several of his students were denied access to their campus accounts, however, and questioned by university officials because they were running computationally expensive programs named `crack`!

If you aren't your own system administrator, but you are concerned about the security of your site, it is probably a good idea to befriend your local sysadmin. He or she can sometimes suggest ways to make your site more secure and can also be an enormous help in recovering from an incident.

Begin your explorations by checking out these sites.

- **Computer Emergency Response Team (CERT)** (`www.cert.org`): CERT is one of the most popular repositories of official descriptions of security incidents. It publishes advisories on all sorts of security issues, including very clear descriptions of the problem, vulnerable systems, and possible solutions.

- **Security-focus.com** (`www.securityfocus.com`): Security-focus.com provides a great deal of information on all aspects of computer security, from the legal and political to the technical. It also hosts the well-known security mailing list, BugTraq (which can be found under Forums).

- **Insecure.Org** (`http://insecure.org`): Insecure.Org is a fairly well-established site that is not afraid to make cracking tools available and to discuss the nitty-gritty details of many "exploits." This site can be extremely useful if you want to try to break into your own site.

Summary

For any significant web site, security is a crucial part of the site's implementation. You should take extreme care to secure your server from attack and also be sure to protect your visitors' private information from prying eyes. In a time of enormous growth for online businesses, publication of a

story about a major security breach can destroy visitors' confidence in your site, driving them to the competition and possibly leaving your site to evaporate as quickly as it appeared.

In this chapter, we've driven home three basic lessons.

1. *Don't trust the network.* Every byte of data that comes from the Internet should be treated as potentially hazardous. Be as restrictive as possible in defining the inputs you allow. Prefer the solution that lists the acceptable inputs to the one that lists the unacceptable inputs. Be sure that your web server configuration does not allow clients to view your source code or to work around your access restrictions.

2. *Minimize the damage.* Wherever possible, make sure that the damage possible from a particular type of security breach is minimal. Encrypt sensitive data. If you run your own web server, make sure it is running as a dummy user.

3. Finally, *if you run your own server, spend some time breaking into it.* If you're successful, then you've identified a vulnerability that you can patch before an intruder finds it. If you're unsuccessful, you've learned something about your server, and your security precautions have weathered a good test. If you don't run your server, find out who does, and see what he or she can tell you about your site's security.

Chapter 29

Learning PHP Configuration

I n this chapter, we discuss the many configuration options available with PHP, particularly the Unix Apache module version, in some detail. The goal is for you to better understand the tradeoffs of each capability you may enable or disable and the capabilities that may affect each other. We also touch on ways you can measure and improve the performance of your PHP scripts.

Viewing Environment Variables

To see any of the settings discussed in the following section, you have only to use the `phpinfo()` function in a valid PHP script. This function begins with a quick recap of the PHP version, your platform, date of build, and compile-time options; it then moves methodically through your PHP settings. You will also see some information about your web server settings and environment variables.

The output of the `phpinfo()` function is a potential bonanza for crackers, so you shouldn't leave it sitting around on a production server. With the release of 5.2.1, `phpinfo()` includes a `META` tag for robots that causes a web spider honoring the robots protocol not to index a `phpinfo()` file. Nevertheless, because there's nothing that requires adherence to the robots protocol, this should not be relied on as a security measure.

Understanding PHP Configuration

Like most of the best open source software packages, PHP is highly configurable. It's left up to you, the individual PHP user, to find your own balance among the competing virtues of power, flexibility, safety, and ease of use.

Configuration is difficult to describe fully because there are so many possible combinations of options. In some cases, there is an obvious conflict between two configuration directives — you simply have to choose one or the other, end of story. In other cases, you can have both but may need to remember some workarounds. We try to point out as many of these implications as we can, but no one can honestly claim to have tested every possible combination.

The Windows version of PHP now ships with the most popular extensions (for example, MySQL) compiled in and a startling number of shared libraries (`.dlls`) bundled with PHP itself. Many of these libraries have to be built from Unix source, so this effort represents a truly amazing amount of unremunerated, thankless work from the PHP build team.

The truth is that the PHP build for Windows (the so-called manual installation, not the installer version) now offers almost all the functionality of Unix builds with much less effort. Windows users only need to worry about the variables that can be set with the `php.ini` file — not all of which are applicable to Windows versions of PHP anyway. If you only use PHP on Windows, feel free to skip down to the "The php.ini file" section of this chapter, with a glance at the "Apache configuration files" section if you run on Apache.

Unix users have a more specific palette of options. To take full advantage of this power, you need to clearly understand the various means by which you can analyze and control your PHP installation. The three most important on the Unix side are:

- Compile-time options
- Web server configuration files
- The `php.ini` file

A few things can also be controlled with runtime options, system settings, or the presence/absence/ configuration of other software packages.

Compile-time options

During the configure/make process, PHP allows you to specify a number of specific flags. This causes the appropriate extensions to be built into your custom version of the PHP module or binary. None of the information in this section is relevant if you are running a precompiled binary (for example, Windows, Mac OS X, or RPM build).

It's important to understand that most compile-time options are merely necessary preconditions for using a particular function set — but that this capability can still be turned on or off, or important configuration options set, in the `php.ini` file. The compilation step and the configuration file

work together. Think of it this way: You must compile with the flag to use the functionality, but you needn't use the functionality just because you compiled with the flag.

> **TIP**   If you fail to employ the appropriate compile-time option, you get an undefined-function fatal error. This error is almost never seen outside of user-defined functions for any other reason, so it should be considered a red flashing light that you need to check your compilation options. Thankfully, you can retrieve your previous options with `phpinfo()` and then simply add the new features you want, should a recompile ever be necessary.

Most compile-time options are pretty self-explanatory. You merely install the required libraries, build PHP with the `--with-[library][=DIR]` flag and, in some cases, set a configuration option in `php.ini`. In the following sections, we will mention only common cases that require special treatment of some kind.

> **TIP**   Remember that all third-party servers and libraries that you plan to use with PHP must be downloaded and installed *before* you attempt to build PHP. This means the web server, a database server, mail, and LDAP servers, XML, encryption, graphics, and `bcmath` libraries must all be in place before PHP.

--with-apache[=DIR] or --with-apache2=[DIR]

This flag causes PHP to be built as a static Apache module. You must use `--with-apache2` if you've ventured into the newest Apache series. Even though the Apache module version is now by far the most popular build, the PHP developers have chosen to leave the CGI build as the default choice. If you forget this (or the `--with-apxs`) flag when trying to make a static Apache module, you will end up with the CGI version.

You almost certainly want to set the Apache base directory parameter because `make` may default to some unexpected location. Remember that Apache installs in different default directories in the source versus RPM builds — so if you've previously installed an `httpd` via RPM (perhaps as part of a Red Hat Linux installation), you should uninstall the package and leave a clean background for the source build you need now.

A static Apache build will have to be recompiled every time you change PHP versions. Apache server, at this point, changes rather slowly, whereas PHP adds new extensions and releases patches rather frequently, so this may be a significant factor in choosing the `apxs` build instead.

--with-apxs[=DIR] or --with-apxs2[=DIR]

This flag specifies that the PHP module be built as a dynamic Apache module. This saves disk space for Apache, and some people claim the build is easier. The main value of the apxs build is that you will be able to swap PHP modules (while upgrading, for instance) without recompiling Apache. If you upgrade frequently, or if you enjoy trying out experimental builds, this is the best option.

> **CAUTION**   Remember that you can build PHP with *either* the `--with-apache` or `--with-apxs` flags, not both.

--with-[database][=DIR]

All the databases supported by PHP use a similar compile-time flag. The directory need only be specified if it is not the default installation directory. For more information on choosing a database for use with PHP, see Chapter 11. The specific flags and default directories are listed in Table 29-1.

TABLE 29-1

Database Compile-Time Information

Database Name	Default Directory	Flag Syntax
Adabas D*	`/usr/local`	`--with-adabas[=DIR]`
DBase	`bundled`	`--enable-dbase`
Filepro	`bundled`	`--enable-filepro`
IBM DB2	`/home/db2inst1/sqllib`	`--with-ibm-db2[=DIR]`
Informix	no default	`--with-informix[=DIR]`
iODBC*	`/usr/local`	`--with-iodbc[=DIR]`
mSql	`/usr/local/Hughes`	`--with-msql[=DIR]`
MySQL < 4.1	`/usr/local/mysql`	`--with-mysql[=DIR]`
MySQL 4.1 and above	`/usr/local/mysql`	`--with-mysqli[=DIR]`
Oracle	`ORACLE_HOME`	`--with-oci8[=DIR]`
PostgreSQL	`/usr/local/pgsql`	`--with-pgsql[=DIR]`
SAP DB	`/usr/local`	`--with-sapdb[=DIR]`
Solid*	`/usr/local/solid`	`--with-solid[=DIR]`
Sybase	`/home/sybase`	`--with-sybase[=DIR]`
Sybase-CT	`/home/sybase`	`--with-sybase-ct[=DIR]`
SQLite	`Bundled`	`--with-sqlite`

The databases marked with an asterisk use ODBC-based interfaces. These ODBC choices are mutually exclusive — you must limit yourself to a maximum of one.

Each database mandates slightly different configuration options in `php.ini` or other configuration files. Oracle, for example, has its own environment variables that obviate PHP settings. Sybase, Oracle, and some other databases escape single quotation marks with single quotation marks, which requires the `magic_quotes_sybase` option in `php.ini`. MySQL allows you to specify a default

hostname, username, and password — not at all a good idea unless you understand the security implications! Most of these options are standard and self-explanatory, however, and they have little effect on other parts of PHP.

--with-mcrypt[=DIR]

This flag builds in the `mcrypt` library, which includes many of the most popular open cipher algorithms. `mcrypt` is available for download at `http://mcrypt.sourceforge.net`.

There is no documented default directory, although PHP can probably find the one mentioned in the `libmcrypt` documentation. `libmcrypt` must be compiled with the `--disable-posix-threads` option. See Chapter 28 for more information on using PHP's cryptography capabilities.

--with-java[=DIR]

This flag builds Java support into PHP. The `DIR` path should be set to the location of your JDK, and the Java settings in `php.ini` must all be set correctly. This extension cannot be used with a static web server build (for example, `--with-apache`), and this flag will probably not work correctly with Solaris versions of PHP and Java. Please see the Java extension `README` in `/php_[build_directory]/ext/java` for more information.

There is an alternate method of accessing Java from PHP: integrating PHP into a Java servlet environment using a SAPI module. You might want to do this if you use Java extensively, as it is the more efficient method. If you choose the servlet integration method, you do not need this extension. See Chapter 38 for more on using Java with PHP.

--with-xmlrpc

This flag builds Dan Libby's XML-RPC and SOAP implementation into PHP. The XML-RPC package now comes bundled with PHP, so you do not need to specify a directory.

CROSS-REF   **To learn more about XML-based Web services and PHP, see Chapter 41.**

--with-dom[=DIR]

This flag builds with DOM XML support, using the GNOME XML library (a.k.a. `libxml`, `gnome-xml`). The `DIR` path should point to your `libxml` installation; if you don't set this value, it defaults to `/usr`. You can download and learn more about GNOME xml from `www.xmlsoft.org`.

TIP   **Very common shared libraries, such as `libjpeg`, can cause fatal problems at PHP compile time even if you correctly set the directory paths in all the compile-time flags. Common issues include PHP looking for the files in the wrong place, incorrect versions of these libraries already being installed on your machine, or libraries having been built in a form inaccessible to PHP.**

The solution to most of these problems is to upgrade all such shared libraries to the latest version. However, if your client applications are old, this may break them. A possible workaround is to temporarily rename the installed versions of your shared libraries, so they cannot be found by PHP; compile the new versions in different locations; compile PHP using these directory paths, then rename your old versions to their original names. Take good notes if you try this!

--enable-bcmath

This option builds support for arbitrary-precision mathematics from a bundled library. You can set the number of decimal places in `php.ini`.

--enable-calendar

This option builds support for calendar conversion functions (for example, Jewish to Julian) from a bundled library.

--with-config-file-path=DIR

This option allows you to specify the location of your `php.ini` file. You need to use it only if you've deliberately moved it away from the default location, `/usr/local/lib`.

--enable-url-includes

This option allows you to include or require *and execute* files from remote HTTP or FTP servers, like this: `include(http://remotehost/include.php)`. This functionality should be carefully considered, as it has horrible security implications. If you merely want to read in HTML files from other servers, you do *not* need this flag.

--disable-url-fopen-wrapper

This flag turns off the default capability to open files on remote HTTP and FTP servers, like this: `fopen(http://remotehost/include.php)`.

CGI compile-time options

All compile-time options just described are available for the CGI version, except for the module-specific flags (for example, `--with-apache`, `--with-apxs`).

Most users today who use PHP's CGI mode are interested in using it as a standalone binary, similar to Perl, rather than for web development. If this is the case, safe mode is probably beside the point.

--with-exec-dir[=DIR]

Another compile-time option relating to safe mode is `--with-exec-dir`. This option sets the default safe-mode execution directory to `/usr/local/bin`, but that can be changed with the `safe_mode_exec` directive in `php.ini`. Remember you can only run programs from this single directory under safe mode.

--enable-discard-path

If you'd like to place the CGI version of PHP outside the web tree and call it as you would a Perl CGI script (such as with `#!/usr/local/bin/php` as the first line of each script), you need to specify this compile-time flag. You must also make all PHP CGI scripts executable.

--enable-force-cgi-redirect

This flag is a security must for the CGI module. It prevents browser users from calling CGI-bin files directly, thereby bypassing Apache security settings. This is an Apache-specific configuration directive; don't bother trying to enable it if you are running on a different web server or as a standalone binary.

Apache configuration files

If PHP is used with Apache as a module or with CGI, much of PHP's basic file-serving capability is determined by Apache's configuration files. The main ones from recent versions of Apache Server are the `httpd.conf` file for global settings, and the `.htaccess` file for per-directory access settings. Older versions of Apache split up `httpd.conf` into three files (`access.conf`, `httpd.conf`, and `srm.conf`), and some users still prefer this arrangement.

In PHP3, there were specific Apache configuration directives that could substitute for almost every `php.ini` setting. For example, instead of setting `Engine = On` in the first substantive line of `php.ini`, you could put `php3_engine on` in an `.htaccess` file for a similar effect. As the number of PHP configuration directives increased, however, people decided that too many flags were cluttering up Apache's namespace. The naming scheme has been generalized, therefore, to encompass these four basic, configurable directives:

- `php_value name value`: Sets value of variable.
- `php_flag name on|off`: Sets Boolean.
- `php_admin_value name value`: Sets value of variable. Can only be used in main Apache configuration file(s) rather than `.htaccess`.
- `php_admin_flag name on|off`: Sets Boolean. Can be used only in main Apache configuration file(s) rather than `.htaccess`.

An example would be magic quotes for `GET`, `POST`, and `COOKIE` variables. You can use `php_flag` with the name of the variable, like this:

```
php_flag magic_quotes_gpc off
```

If this is all too confusing, don't worry: The new-style Apache configuration directive naming only applies to settings you can change in `php.ini` anyway.

Apache server has a very powerful, but slightly complex, configuration system of its own. Learn more about it at the Apache web site: `www.apache.org`.

The following headings describe settings in `httpd.conf` that affect PHP directly and cannot be set elsewhere.

Timeout

This value sets the default number of seconds before any HTTP request will time out. If you set PHP's `max_execution_time` to longer than this value, PHP will keep grinding away but the user

may see a 404 error. In safe mode, this value will be ignored; you must use the `timeout` value in `php.ini` instead.

DocumentRoot

`DocumentRoot` designates the root directory for all HTTP processes on that server. It looks something like this on Unix:

```
DocumentRoot    "/usr/local/apache_1.3.6/htdocs"
```

It looks like this on Windows:

```
DocumentRoot    "C:/Program Files/Apache/htdocs/"
```

The document root can be almost any directory — it needn't be in the Apache installation directory. You can specify a subdirectory of this as the PHP document root, using the `doc_root` setting in `php.ini`. In this case, HTML files would be served out of the Apache document root and its subdirectories, but PHP would be parsed only in the specified PHP directory and its subdirectories.

AddType

The PHP MIME type needs to be set here for PHP files to be parsed.

Remember that you can associate any file extension with PHP; many administrators set the `.php3` and `.html` types for backward compatibility — but if you wanted to, you could have PHP parse files called `filename.asp` or `filename.jsp`. You can also add multiple types for different versions of PHP. The following are sample `AddType` lines; the first one is the most common for PHP4 and above, but you can add as many of the others as you wish.

```
AddType application/x-httpd-php .php
AddType application/x-httpd-phps .phps
AddType application/x-httpd-php3 .php3 .phtml
AddType application/x-httpd-php .html
```

You can also set this on a per-directory basis with `.htaccess`, simply by adding type lines to `.htaccess` files. PHP files are then parsed only in one directory of your site (for instance, in a forum folder). Alternatively, you can set up a directory with archived versions of files that may have old extensions — so just in that directory, Apache allow files with the `.phtml` extension to be parsed.

Action

You must set this line for the CGI version of PHP with Apache, generally used with Windows. You do *not* need to set this line for the module version of PHP.

```
Action application/x-httpd-php4 "/php/php.exe"
```

490

LoadModule

You must uncomment this line for the Windows apxs module version of Apache with shared object support:

```
LoadModule php4_module modules/php4apache.dll
```

or on Unix flavors:

```
LoadModule php4_module modules/mod_php.so
```

AddModule

You must uncomment this line for the static module version of Apache.

```
AddModule mod_php4.c
```

The php.ini file

The PHP configuration file, `php.ini`, is the final and most immediate way to affect PHP's functionality. Important changes are frequently made in the structure of this file, so if you haven't bothered to really look at every line recently, now may be a good time.

The `php.ini` file is read each time PHP is initialized — in other words, whenever `httpd` is restarted for the module version or with each script execution for the CGI version. If your change isn't showing up, remember to stop and restart `httpd`. If it still isn't showing up, use `phpinfo()` to check the path to `php.ini` (near the top of the file); if necessary, recompile the code with the `--with-config-file-path` flag or just move `php.ini` to wherever PHP expects to find it.

CAUTION **What happens if PHP can't find `php.ini`? Under Windows, right up until the formal release of PHP4, you used to get an "unable to parse configuration file" fatal error. Under Unix and now under Windows as an ISAPI module, interestingly enough, you get no warnings or errors — PHP carries on with default settings, which are the same as if you had not changed any settings in `php.ini-dist`. You need to install `php.ini` only if you want to change the default settings.**

The configuration file is well commented and thorough. Keys are case-sensitive, keyword values are not; whitespace, and lines beginning with semicolons are ignored. Booleans can be represented by `1/0`, `Yes/No`, `On/Off`, or `True/False`. The default values in `php.ini-dist` will result in a reasonable PHP installation that can be tweaked later.

What follows are notes explaining the settings in `php.ini` that are not completely documented in the file or the PHP manual's `configuration.html` page.

short_open_tag = Off

Short open tags look like this: `<?  ?>`. This option *must* be set to `Off` if you want to use XML functions.

disable_functions = [function1, function2, function3 . . . function*n*]

A welcome addition to PHP4 configuration and one perpetuated in PHP6 is the ability to disable selected functions for security reasons. Previously, this necessitated hand-editing the C code from which PHP was made. Filesystem, system, and network functions should probably be the first to go because allowing the capability to write files and alter the system over HTTP is never such a safe idea.

max_execution_time = 30

The function `set_time_limit()` won't work in safe mode, so this is the main way to make a script time out in safe mode. In Windows, you have to abort based on maximum memory consumed rather than time. You can also use the Apache timeout setting to timeout if you use Apache, but that will apply to non-PHP files on the site, too.

error_reporting = E_ALL & ~E_NOTICE

The default value is `E_ALL & ~E_NOTICE`, all errors except notices. Development servers should be set to at least the default; only production servers should even consider a lesser value.

error_prepend_string = ["<font color=ff0000>"]

With its bookend, `error_append_string`, this setting allows you to make error messages a different color than other text, or what have you. We recommend setting the value to `"<blink>"` (and `error_append_string` to `"</blink>"`, of course) for a special treat! The default values result in a red error message. Remember to uncomment these if you want to use them — they're commented out by default.

warn_plus_overloading = Off

This setting issues a warning if the + operator is used with strings, as in a form value.

variables_order = EGPCS

This configuration setting supersedes `gpc_order`. Both are now deprecated along with `register_globals`. It sets the order of the different variables: Environment, `GET`, `POST`, `COOKIE`, and `SERVER` (aka `Built-in`). You can change this order around. Variables will be overwritten successively in left-to-right order, with the rightmost one *winning* the hand every time. This means if you left the default setting and happened to use the same name for an environment variable, a `POST` variable, and a `COOKIE` variable, the `COOKIE` variable would own that name at the end of the process. In real life, this doesn't happen much.

gpc_order = GPC

Deprecated.

auto-prepend-file = [path/to/file]

If a path is specified here, PHP must automatically `include()` it at the beginning of every PHP file. `Include` path restrictions do apply.

auto-append-file = [path/to/file]

If a path is specified here, PHP must automatically `include()` it at the end of every PHP file — unless you escape by using the `exit()` function. `Include` path restrictions do apply.

include_path = [DIR]

If you set this value, you will only be allowed to include or require files from these directories. The `include` directory is generally under your document root; this is mandatory if you're running in safe mode. Set this to `.` in order to include files from the same directory your script is in. Multiple directories are separated by colons: `.:/usr/local/apache/htdocs:/usr/local/lib`.

doc_root = [DIR]

If you're using Apache, you've already set a document root for this server or virtual host in `httpd .conf`. Set this value here if you're using safe mode or if you want to enable PHP only on a portion of your site (for example, only in one subdirectory of your web root).

upload_tmp_dir = [DIR]

Do not uncomment this line unless you understand the implications of HTTP uploads!

session.save-handler = files

See Chapter 24 for details on this setting. Except in rare circumstances, you will not want to change this setting.

ignore_user_abort = [On/Off]

This setting controls what happens if a site visitor clicks the browser's Stop button. The default is `On`, which means that the script continues to run to completion or timeout. If the setting is changed to `Off`, the script will abort. This setting only works in module mode, not CGI.

Improving PHP Performance

There are two schools of thought about web performance. The first is that PHP script performance, theoretical web server speed, chip clock speed, server RAM, and almost everything else is made irrelevant by throughput issues — so why sweat the small stuff? The other is that there's no thrill quite like that of shaving a few microseconds off your script execution time. This section is basically useless for proponents of the former view.

Before you can improve your performance, you have to measure it. We use the time-honored programming performance metric: measuring microseconds. Whip up a little function like this:

```
function exec_time()
{
  $mtime = explode( " ", microtime());
  $msec = (double)$mtime[0];
```

493

```
    $sec = (double)$mtime[1];
    return $sec + $msec;
}
```

This function just reformats microtime output into a double for easier subtraction. Paste or include it at the top of the script you'd like to measure. Now divide the main body of your script into sections and scatter calls to exec_time() at strategic points, like this:

```
<?php
$start_db_call = exec_time();
$db = mysql_select_db("test");
$result = mysql_query("SELECT * FROM user
                       WHERE ID=1");
while ($testrow = mysql_fetch_array($result)) {
  echo $testrow[0];
}
$end_db_call = exec_time();
$runtime = $end_db_call - $start_db_call;
echo "Database call and echo took $runtime seconds";
?>
```

The next time you hit the web page, voilà! A self-timing PHP script, at your service.

CAUTION Using microtime() to measure PHP tells you only what happens between the time PHP begins working on the first measured line of code and the time it finishes working on the last measured line of code. It does not tell you how long your web server is taking to spawn a child process or your CGI to start up, how much latency your server is suffering from, what traffic conditions at your web farm are like, or a lot of other things that affect real-world performance at least as much if not more than actual PHP processing time. To find out that kind of thing, you need measuring tools far beyond PHP. A good start for Apache on Unix is the program called ab (aka Apache Benchmark tool), which ships with Apache.

Now that you know how long the various parts of your script are taking, you can take steps to improve performance. Actually, a little logic should tell you that functions that touch other files or call other daemons should take longer than those that are self-contained within a discrete file. So database calls, include and require statements, objects with inheritance, and XML parsing are just going to take longer than simple arithmetic or echoing a string. But because these advanced functionalities are the best part of PHP, obviously it would be pointless to get rid of them for the sake of squeezing out a few more microseconds.

What you can and should do instead is hunt and destroy gross programming errors that cause unnecessary latency. Infinite loops, you know, are never very stylish. If you can notice a script running slowly with the naked eye, especially on a localhost, it's cause for concern — whip out the microtime and find out where it's going wrong. Pay special attention to known bottlenecks such as: using regex instead of the faster, built-in functions such as explode() in a tight loop; object-oriented programming where it's not needed; bad use of SQL; including multiple instances of the same files; and long loops.

Although it would be better to eliminate all errors in the code itself, you can also help matters by setting the Apache or PHP timeout and max-memory configuration variables as low as possible. Come on — no web page should need a 300-second timeout and you know it. Another configuration setting that may have a good effect on extremely slow scripts is `ignore_user_abort` in `php.ini`.

Recent distributions of PHP have also offered an optimized `php.ini` that sets variables for maximum speed at the possible expense of other virtues. If you choose to use this file, please take the time to understand the effects of its changes, as the changes may affect legacy scripts that expect HTTP_* variables and other globals.

Optimizers and Caches

Until recently, speed-shavers had few options but homemade metrics like the `microtime()` function just described. But now, intriguing tools are beginning to become widely used to increase PHP performance.

One that is available without cost is the Zend Optimizer. This tool makes multiple passes over a PHP script and replaces slower constructs with faster ones that have the same effect. However, the Zend Optimizer is rumored to mostly help inexperienced coders: If you already write tight PHP, it may not be able to add much value.

Another Zend product that promises to affect performance positively is the Zend Accelerator. This product apparently compiles and stores a version of each page in memory, reducing disk reads and redundant compilation and thus speeding Web service. Reliable reports claim that the Accelerator can deliver from two to ten times improvement in number of requests handled. Both the Zend Accelerator and the Zend Optimizer are available at the Zend web site, `www.zend.com`.

There are also optimizing and caching products available without cost. Two popular choices are Nick Lindridge's PHP Accelerator (`www.php-accelerator.co.uk`) and APC (`http://apc.communityconnect.com`). These may be included in PHP 6 or in subsequent point releases, following the initial PHP 6 release. At the time of this writing, the final decision hadn't been made.

Summary

The good thing and the bad thing about PHP configuration are the same: There are a whole heck of a lot of options and more than one way to set many of them. The Unix Apache module is particularly rich in choices, but the development team has labored long and hard to make PHP as customizable as possible.

There are three main ways to configure PHP. The first is via build-time flags, which are only available to those who build from source. Many of these directives are only necessary preconditions, meaning they set default conditions that need to be confirmed or can be reversed elsewhere. The

second is via Apache configuration files (`httpd.conf` and `.htaccess`), which are only available to users of Apache server. The third is via the `php.ini` file, which comes with every PHP distribution.

The `php.ini` file experienced a few significant changes with PHP4, but hopefully has stabilized somewhat with PHP6 One of the most important is the capability to disable functions on an individual basis. Certain features of PHP3 and PHP2 are beginning to be deprecated in this file, such as `register_globals`. And the `php.ini` is no longer an absolute necessity on Windows — versions of PHP now recognize default values even without the file being present in the Windows path.

After you've run PHP for a while, you may wish to tune its performance. Beginning with PHP 5 and continuing into PHP 6, execution times are considerably faster at the same tasks than earlier versions, and in general script execution time isn't the bottleneck to total performance — but you may want to maximize the efficiency of your PHP-enabled server anyway. The main tool available to measure performance is simply echoing `microtime()` at intervals throughout a script. With this simple method, you can try to narrow down and improve the parts of your scripts that are taking the most time. This does not measure anything outside PHP that may affect its performance; for that, you need external tools such as `ab` (Apache Benchmark).

Tools to help speed up PHP are becoming widely available. One of the most intriguing is Zend's Accelerator, which promises at a minimum to double pages served on the same hardware. There are also alternatives available without cost.

Chapter 30

Handing Exceptions with PHP

Until now, programmers have been very creative in figuring out how to deal with error cases within PHP, whether setting and printing error strings or using and abusing the limited error reporting system. Despite its many useful features, PHP has not contained a good system for comprehensively dealing with errors. Fortunately, this changed with PHP5.

Error Handling in PHP

If you are familiar with structured programming languages, such as C and Java, you have probably grown accustomed to the various built-in objects that allow you to handle errors and exceptions. If so, you'll be happy to note that PHP5 now includes, for the first time, an exception-handling object, and that the syntax is very similar to existing languages like Java. In fact, once you learn a bit of syntax you can begin handling errors and exceptions much as you have been with other object-oriented languages.

However, if you have been primarily using PHP, and are unfamiliar with other languages, the idea of exception handling may be new to you. Exception handling is a powerful tool that you will come to appreciate once you understand the concept and put it to good use in your code. This new built-in function will enable you to debug error conditions, recover from unexpected situations, and present a clean interface to your end users without printing errors to the screen.

Errors and exceptions

It is helpful to think of an exception not merely as an error. An exception, as the name implies, is any condition experienced by your program that is

unexpected or not handled within the normal scope of your code. Generally, an exception is not a fatal error that should halt program execution, but a condition that can be detected and dealt with in order to continue properly.

Exceptions, when properly used, can greatly increase the reliability of your application, and cut down on debugging headaches. However, poorly handled or ill-defined exceptions can create more problems than they solve by obscuring the source of an error. Plan on taking the time to properly determine and execute your exception-handling method, and you will be amply rewarded.

Take a look at some sample code that contains some error detection as it might be handled in PHP4 or earlier. We are retrieving a POST variable containing a user's ID, which must be at least nine characters in length and begin with the "usr" prefix. Once the variable is checked for these conditions, we pass it to a function that will verify the existence of the user within the site database.

LISTING 30-1

Error-handling without exceptions

```php
<?php

  // include the file which will validate the user ID
  require_once('includes/usr_functions.php');

  // retrieve the user ID to validate
  $user_id = $_POST['user_id'];

  // set the display message based on whether or not the
  // user ID is valid
  if ( !is_valid_user($user_id)) {
     $msg = "Sorry, $user_id is not a valid user ID.";
  } else {
     $msg = "$user_id is a valid user ID.";
  }


  function is_valid_user($user_id) {

     // return false if the user ID does not begin with "usr"
     $pre_str = "usr";
     if ((strpos($user_id, $pre_str) === false) ||
       (strpos($user_id, $pre_str) != 0)) {
       return false;
     }

     // return false if the user ID is not the proper length
     if ((strlen($user_id) < 9)) {
       return false;
     }
```

```
       if (validate($user_id)) {
         // user ID was found in the database
         return true;
       } else {
         // the specified user ID does not exist in the database
         return false;
       }

}
?>
```

As you can see, any number of conditions might cause the `is_valid_user()` function to return a value of false, only one of which actually pertains to the question of whether the user exists in the database. Using exceptions, you can more easily distinguish among types of errors and deal with them according to the nature of each error.

The Exception class

The new `Exception` class was built into PHP5 and ready for use with any code on a PHP5 or later server. Rather than using Boolean functions as in the preceding example, an instance of Exception can be created or thrown within the code.

Listing 30-2 shows what Listing 30-1 might look like after rewriting to use exception handling.

LISTING 30-2

Error-handling using exceptions

```php
<?php

// include the file which will validate the user ID
require_once('includes/usr_functions.php');

// retrieve the user ID to validate
$user_id = $_POST['user_id'];

try {

   // set the display message based on whether or not
   // user ID is valid
   if (!is_valid_user($user_id)) {
      $msg = "Sorry, $user_id is not a valid user ID.";
   } else {
      $msg = "$user_id is a valid user ID.";

} catch(Exception $ex) {
  // retrieve the message from the exception object
```

499

```php
    $msg = ($ex->getMessage());

}


function is_valid_user ($user_id) {

    // throw an exception if the user ID does not begin
    // with "usr"
    $pre_str = "usr";
    if ((strpos($user_id, $pre_str) === false) ||
        (strpos($user_id, $pre_str) != 0)) {
        throw new
        Exception('$user_id does not contain the proper prefix.');
    }

    // throw an exception if the user ID is not the proper length
    if ((strlen($user_id) < 9)) {
        throw new Exception('$user_id is less than the
            required length');
    }

    if (validate($user_id)) {
        // user ID was found in the database
        return true;
    }
    else
        // the specified user ID does not exist in the database
        return false;
    }

}

?>
```

Don't be thrown by the use of the word *throw*. In this case, it is used to create a new Exception object. Now errors that have nothing to do with normal flow are handled as separate exceptions rather than mingling with the rest of the application.

The try/catch block

Exceptions are caught and handled using a try/catch control construct. You will want to include any code that may generate an error or exception within the try() construction. Whenever any exception is thrown by the code, the try() block execution is terminated, and the remaining code within the try() construction is not executed. The catch() block is then consulted to find the proper

type of exception, and the exception is then dealt with according to the code within that particular catch block. We now have different conditions based upon the type of exception that was thrown, rather than one general, nonspecific error.

Throwing an exception

One of the nice things about using exceptions is the ability to display as much — or as little — information as you need. There are several methods available for use with an Exception object, which you can use to create your own error messages or to deal with conditions accordingly.

The following code shows how you might throw and immediately catch a generic exception, and then take apart the Exception object to recover the message, the error code, and the originating file and line number.

```php
<?php

try {

  throw new Exception('Syntax error');

} catch(Exception $ex) {

    // the input string passed to the object
    $msg = ($ex->getMessage());
    // customizable error code
    $code = ($ex->getCode());
    // name of the file that threw the exception
    $file = ($ex->getFile());
    // line number containing the exception
    $line = ($ex->getLine());

  echo "Error no. $code: $msg in file $file on line $line";
}

?>
```

Here, a standard, PHP style error message is displayed. However, you as the programmer can display anything that you like or can even change application behavior based on the specific error.

Note that, although in this example the code that throws the exception is the only thing in the block, we could have had arbitrarily complex code that calls a function defined in a different file, which throws an exception only some of the time. Whenever an exception is thrown, control will revert to the catch block associated with the try.

Multiple catch blocks can be used to deal successfully with more than one type of exception. We'll look at an example in the following section.

Defining your own Exception subclasses

PHP also allows you to define your own classes that inherit from the `Exception` class. Now you no longer have to rely on the `getMessage()` function for information on the specific type of error that was generated. Subclasses can be defined as in the following example:

```php
<?php

class CustomException extends Exception {
    public function __construct($message) {
        parent::Exception($message);
    }
}

?>
```

Let's look at the example covered in Listing 30-3 and consider using custom exceptions. People signing in using this code may forget to include the "usr" prefix in their username, so if it's missing you might want to try adding it and validating again rather than immediately halting the program.

LISTING 30-3

Recovering using custom exceptions

```php
<?php

// include the file which will validate the user ID
require_once('includes/usr_functions.php');

// retrieve the user ID to validate
$user_id = $_POST['user_id'];

try {

   // set the display message based on whether or not
   // user ID is valid
   if (!is_valid_user($user_id)) {
      $msg = "Sorry, $user_id is not a valid user ID.";
   } else {
      $msg = "$user_id is a valid user ID.";
   }

} catch(PrefixException $ex) {

  // if prefix is missing, try again with proper prefix
  $user_pre = "usr" . $user_pre;
  if (!is_valid_user($user_pre)) {
    // second attempt has failed, retrieve message
    $msg = ($ex->getMessage());
```

```php
   }

}

//define custom exception classes

class PrefixException extends Exception {
   public function __construct($message) {
      parent::Exception($message);
   }
}

class LengthException extends Exception {
   function __construct($message) {
      parent::Exception($message);
   }
}

echo $msg;

function is_valid_user ($user_id) {

   // throw an exception if the user ID does not begin
   // with "usr"
   $pre_str = "usr";
   if ((strpos($user_id, $pre_str) === false) ||
      (strpos($user_id, $pre_str) != 0)) {
      throw new
PrefixException('$user_id does not contain the proper prefix.');
   }

   // throw an exception if the user ID is not the proper length
   if ((strlen($user_id) < 9)) {
       throw new
  LengthException('$user_id is less than the required length');
   }

   if (validate($user_id)) {
      // user ID was found in the database
      return true;
   }
   else
      // the specified user ID does not exist in the database
      return false;
   }

}

?>
```

Note that we attempted to recover from the *missing prefix* error condition. You can easily deal with individual types of errors now that they have been defined separately.

Limitations of Exceptions in PHP

The Exception object is completely new in PHP5 and as such is still in the rough stages of development. As of this writing, PHP does not support the use of `finally()` or `throws()` methods as do Java and other languages. Also, unlike other languages, native PHP errors — including errors, which are normally printed to the client-side browser — are not yet mapped to exceptions. Because of this, for example, a SQL statement error within a try/catch block will not automatically throw an exception that can be caught and dealt with. This handy functionality will most likely be included in a future version of PHP, so it's worth mentioning and keeping an eye out for. Some of these errors can be dealt with using techniques described in the next section.

Other Methods of Error Handling

If you're still using PHP4 or an older version, or are not comfortable dealing with classes and objects, there are several error-handling functions that have been available in PHP for some time, including native PHP errors, defining an error handler, and triggering a user error.

Native PHP errors

PHP generates several types of errors, depending on severity. Here are three common types. You can find more information at `www.php.net/error_reporting`.

- **Notice:** These errors are not serious and do not create a serious problem. By default they are suppressed, unless the logging level is changed in the `php.ini` file.

- **Warning:** Failed code has created an error, but does not terminate execution. Usually the error is displayed, but the script continues to run. (See Figure 30-1.)

- **Fatal error:** A serious error condition has rendered the script unable to run. A fatal error terminates the script. (See Figure 30-2.)

Each type of error is also represented by a constant that can be referred to within your code: `E_USER_NOTICE`, `E_USER_WARNING`, and `E_USER_ERROR`. The error-reporting level can be manually defined within a script, as in these examples:

```
//report only fatal errors
error_reporting(E_USER_ERROR);

//report warnings and fatal errors
error_reporting(E_USER_WARNING | E_USER_ERROR);

//report all errors, including notices and E_STRICT
error_reporting(E_ALL);
```

FIGURE 30-1

Native PHP Warning allows the page to finish rendering
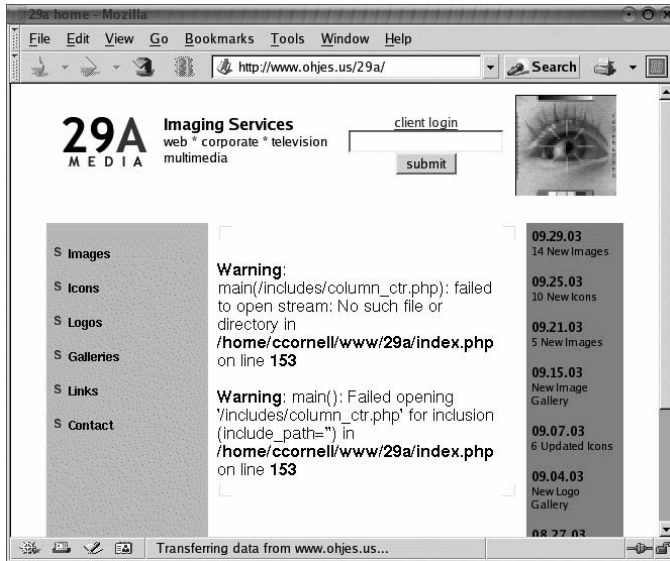
FIGURE 30-2

Native PHP Fatal error terminates execution of the page
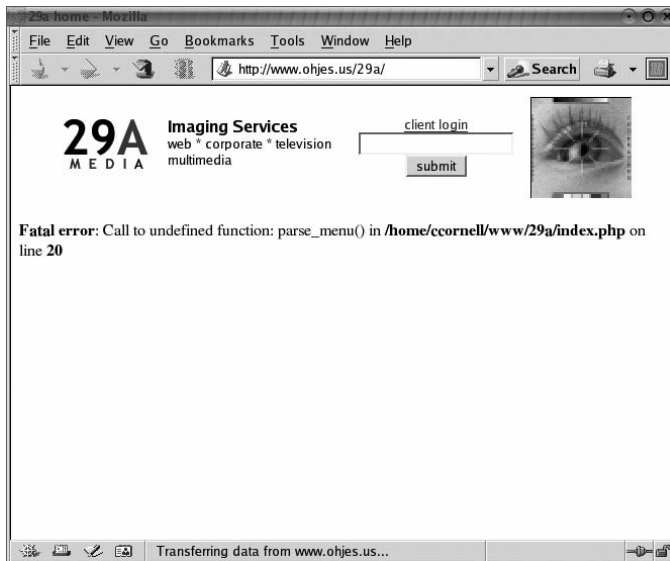
CAUTION **Suppressing error reporting to avoid printed errors can lead to hair loss during the debugging process! You will instead want to deal with an error handler, for the most part.**

Because notices never make it to the client, and don't impair functionality, you're almost always safe in disregarding them for error-handling purposes. Conversely, the custom error handler cannot handle fatal errors; PHP considers them serious enough to terminate the script, no questions asked. So the usefulness of the custom error-handling function is generally limited to warnings. The primary use of this function is to avoid printing *program-ese* error messages for the end user and disrupting the flow of the application.

Defining an error handler

There's an important question to ask at this point: What information do you want displayed to the user when an error occurs? Usually, it's not important, or even preferable, to display details of the inner workings of your application to an end user; not to mention that errors look ugly on a web page. By creating a function that designs a custom error message, then setting that function as the default error handler, you can avoid the awkward and unprofessional display of errors to a user.

First, let's create a function and determine what information we would like to provide. We will need to accept as input parameters the error type, message, filename, and line number.

```php
<?php
   function error_msg($err_type, $err_msg, $err_file, $err_line)
   (
      echo "<div class='errorMsg'>";
      echo "<b>Error:</b>";
      echo "<p>";
      echo "We're sorry, but an error has occurred " .
         "in this page. ";
      echo "Please access the <a href='/help.html'>Help" .
         "</a> page, ";
      echo "or try again later.";
      echo "</div>";
      echo "<div class='finePrint'>";
      echo "Error type: $err_type: $err_msg in $err_file " .
         "at line $err_line";
      echo "</div>";
   }
?>
```

We have elected, in this case, to provide information on the specific error and where it occurred. Depending on the situation, we may want to provide very little information to the user other than the fact that an error has occurred and provide information on what to try next.

Now that we have defined our custom error handler, we simply need to refer to the function within the code using the set_error_handler() function:

```php
set_error_handler("error_msg");
```

FIGURE 30-3

A custom error handler preserves the interface

With this code in place, all errors that are enabled by the error-reporting level will direct through our custom function, with the unfortunate exception of a fatal error. See Figure 30-3 for an example of such an error being displayed.

Triggering a user error

PHP4 can be used to trigger a user error, which is roughly equivalent to throwing an exception in PHP5. An error of any type can be thrown by passing an error message and an optional error-level constant:

```php
<?php

  // trigger a user error if the user id is not valid
  if ( !is_valid_user($user_id)) {
     trigger_error("Invalid user ID", E_USER_WARNING);
  }

?>
```

The triggered error is best used in conjunction with a custom error handler. Once this error is thrown, your defined error handler will be used to provide a formatted error message to the user. Note that the error-level constant defaults to E_USER_NOTICE, which is ignored unless you have specifically set error reporting otherwise.

Logging and Debugging

As mentioned earlier, exception handling and error reporting can turn your debugging efforts into a nightmare if not executed with care. An error that is difficult to track down can become almost impossible when it has been suppressed or rerouted carelessly. With a bit of planning, however, exceptions and error handlers can greatly simplify maintenance of an application.

Previously, we have examined the process of dealing with errors primarily as a means of avoiding disruption for the user. The same process can also be applied to enable logging and debugging for the programmer. Within your catch() control block, include a call to a built-in function such as error_log(), with any relevant information that might aid in the debugging process. Note that you might want to explore the use of the getTraceAsString() function which can be helpful for logging.

```
} catch(Exception $ex) {

    // the input string passed to the object
    $msg = ($ex->getMessage());
    // customizable error code
    $code = ($ex->getCode());
    // name of the file that threw the exception
    $file = ($ex->getFile());
    // line number containing the exception
    $line = ($ex->getLine());

    // write to error log
    $log_msg = "Error $code in $file at line $line: $msg : " .
       time();
    error_log ($log_msg, 3, "/var/tmp/php_error.log");

    //print to screen
  echo "Error no. $code: $msg in file $file on line $line";

}
```

The process is similar when using a custom error handler (as is the aforementioned getTraceAsString()):

```
<?php
    function error_msg($type, $msg, $file, $line) (
```

```
                // write to error log
                $log_msg = "Error $type in $file at line $line: $msg : " .
                        time();
                $log_path = "/var/tmp/php_error.log");
                error_log ($log_msg, 3, $log_path);

                //print to screen
                echo "Error type: $err_type: $err_msg in $err_file " .
                   "at line $err_line";

            }
    ?>
```

error_log() accepts one of four integers as the second parameter, which sets the message type in conjunction with the third, or location, parameter:

> 0: uses the operating system's system logging mechanism
>
> 1: sends the error to a specified e-mail address (extra headers may be added as a fourth parameter)
>
> 2: sends the error through PHP's debugging connection (remote debugging must be enabled)
>
> 3: error message is appended to a destination error log file

Summary

Error handling continues to become easier and more uniform with the ongoing development of PHP. The Exception class provides a means of separating error conditions, or exceptions, from the flow of the application. Using the try/catch block and custom-defined Exception subclasses, errors can be intercepted and even recovered.

Previous versions of PHP also provide a measure of error handling and reporting. User errors can be triggered as an alternative to throwing an exception, and custom error handlers can provide a better user experience as well as useful debugging information. Logging and debugging are crucial to successful exception and error handling.

Chapter 31

Debugging PHP Programs

Debugging — finding and eliminating errors — is part of software development. As a PHP programmer, you should be aware of all the tools available to you as you seek to eliminate malfunctioning elements in your software systems.

There are many such tools, not least because PHP applications usually rely on the capabilities of several servers (such as an HTTP server and a database management server), each of which typically comes equipped with its own logging and reporting capabilities with which it keeps its users hip to what's happening. Plus, PHP has a considerable error-reporting facility of its own (you can choose to have error messages printed alongside normal output or logged to a file for more discreet analysis). The language also has a number of functions with which you can have your programs generate custom error reports, and at the very least you can use conditional print statements to monitor the activity of programs (and the values of variables within them) as they execute.

On top of the built-in error-reporting capabilities of PHP and its supporting technologies, PHP programmers now have access to the sorts of debugging tools that programmers working with other languages have had for years. Chief among these is the Zend debugging environment, which allows you to monitor variable values, set breakpoints, and step through programs at any pace you like. Though the Zend debugging environment won't be covered in this chapter, you can find out more information about it at www.zend.com.

This chapter aims to introduce you to the tools and techniques available to you as you work to perfect your PHP software.

511

General Troubleshooting Strategies

The two basic elements of a debugging effort are figuring out what's wrong, and then fixing it (without breaking something else as a side effect of your solution). It doesn't matter whether you're diagnosing a PHP program, a telephone switch, an electronic circuit, or a Buick — certain principles apply regardless. Bear these ideas in mind as you try to figure out what ails your software.

Change one thing at a time

It's a basic rule of experimentation: You can't be sure what caused a given effect if there are multiple variables. Make a single change, then examine the output and see if the unwanted behavior is fixed. If not, try one more change (possibly changing the first one back to the way it was).

Try to isolate the problem

If you can narrow the problem down to a single library or function, you've made significant progress in locating the cause. Use special `var_dump()`, `echo()`, and `print_r()` calls to output trace information frequently. This will allow you to see when troublesome changes are taking place and when variables stop holding the values you think they're holding.

You can also use a visual debugger (like Zend Studio) to monitor programs and their members as they execute.

Simplify, then build up

It sounds obvious, but if you're having trouble with a given function or feature, cut it out (either literally or with comments) and make sure that everything runs without it. Then replace dynamic data with static data (replace a database query with simple variable assignment statements). Get it working right under simple conditions, and add complexity in stages, testing all the way to see when errors appear.

Check the obvious

We've all heard the story about the call to tech support in which the customer complains that he can't see his mouse pointer move, and after lots of diagnostics it turns out that the machine isn't plugged in. It's probably apocryphal, but in any case make sure your that web server is working properly on its own, and that a basic "Hello, World" script renders properly. You can also add `phpinfo()` to the end of a simple test script to get a lot of information about your PHP interpreter's version and environment details.

Speaking of version, make sure that you're not trying to do something that requires `register_globals` (the infamous setting in `php.ini`) to be on. That setting is set to no by default, as of PHP4.2, and it's tripped up more than one programmer.

Document your solution

It's extraordinarily common: You struggle with an error condition for hours (or longer) and finally reach a solution. At that point, don't immediately head out to celebrate. Take a minute to document what happened and what the solution was. That way, you'll be ready when the same problem pops up again — and it will.

After fixing, retest

It's not unusual to fix a problem and in doing so break something else. That's why it's important to retest your system beyond the scope of the bug you were originally after. This also points out why it's important to isolate bugs as much as possible — it limits the scope of retesting you have to do.

A Menagerie of Bugs

A number of different kinds of bugs plague programmers. Some bugs are both simple in nature and easily found (as is the case with syntax errors and spelling mistakes). Others are significantly more difficult to catch, which is why this chapter is here.

Compile-time bugs

PHP is a compiled language — it's compiled just before it executes, so the compilation isn't as obvious as it is in C or Java.

A compile-time bug is obvious to the Zend Engine, which does the compiling. The compiler will raise an objection, often with a line number, and you can go fix the problem. Examples of compile-time errors are mistyped variable names, forgotten semicolons, and mismatched parentheses.

Runtime bugs

A runtime bug doesn't appear until after your program is under way, and may result from some outside condition, such as unexpected input from a user or unanticipated behavior by a database. These have to be tested for, as they usually won't make themselves evident to programmers under all conditions.

Logical bugs

Logical bugs are perhaps the most difficult of all to spot and can be very difficult to fix if they result from an error in thinking.

Say that you wanted to launch a space probe and have it enter orbit around Mars. However, because your navigation algorithm didn't allow for metric input from those pesky Europeans, your space probe crashed into the Martian surface. The software did exactly as it was told, which, strictly speaking, was to drive the rocket into Mars. That's a logical error.

513

The point: Make sure that your programs not only generate output, but generate the correct output. Get out the calculator and make sure that the program's results are right, or compare its results to values known to be good. And don't use PHP to program spacecraft, just to be safe.

CROSS-REF **For a guide to the most common symptoms of the most common compile-time and run-time bugs, see Chapter 10.**

Using Web Server Logs

Because most PHP programs result in some sort of HTML page, which is in turn served by an HTTP server such as Apache or Microsoft Internet Information Server (IIS), it is possible for errors to be introduced by the web server software. For that reason, it is important to be familiar with the way in which your web server manages error reporting and logging and to know how to access and interpret the logs you need.

Apache

The Apache HTTP Server maintains two log files in plain-text format. They are:

- `Apache/logs/access.log`: Notes every HTTP request for a file, including its date, time, and result (success or failure, as indicated by a numeric status code). The access log also records the IP address from which each request came.

- `Apache/logs/error.log`: Records error conditions only.

The Common Log Format

By default, entries in the Apache error.log file use the standardized Common Log Format. Entries in this format each correspond to a single instance of request/response activity (requests and responses are, after all, what HTTP servers handle). For example, one line might correspond to a request for an HTML page (and its subsequent service by Apache). The next line might correspond to the (automatic) request for and service of a JPEG file embedded in that HTML document.

In any case, Common Log Format entries look like this (in a single line):

```
192.168.100.1 - david [10/Nov/2003:18:00:30 -1100]
"GET /index.html HTTP/1.0" 200 6590
```

The most important elements of that line are:

- `192.168.100.1`: The IP address of the client making the HTTP request

- `david`: The username of the authenticated user making the request

- `[10/Nov/2003:18:00:30 -1100]`: The date, time, and UTC offset of the request

- `GET`: The nature of the HTTP request: `GET` or `POST`

- `index.html`: The requested file
- `HTTP/1.0`: Version of the HTTP protocol used for the request
- `200`: Response code describing the result of the request (more on this later in this section)
- `6590`: The number of bytes served out in HTTP response corresponding to this request

You'll find a more complete treatment of Apache log files, including the more obscure elements of the Common Log Format, at `http://httpd.apache.org/docs/logs.html#errorlog`.

HTTP response codes

Though there are many HTTP response codes (the most famous being the "404 Not Found" error), they exhibit a pattern that aids rapid decoding. In a nutshell:

- 200-series codes indicate success
- 300-series codes indicate a redirection
- 400-series codes indicate a client-side error (like a request for a nonexistent document)
- 500-series codes indicate a server-side error

You'll find a full list of HTTP response codes at `www.w3.org/Protocols/rfc2616/rfc2616-sec10.html`.

Monitoring Apache logs with tail

Under Unix (including Linux), you usually have access to the GNU text utility suite. When it's time to monitor log files, one of the most useful of these tools is tail.

In its default behavior, tail will return the last (that is, most recent) 10 lines of a specified file. You can use it like this:

```
tail access.log
```

and get 10 lines of Common Log Format output (assuming that 10 loggable events have taken place).

More usefully, though, you can use tail in its follow (`--follow`) mode. In follow mode, tail returns the 10 newest lines of a specified file, then goes into an infinite loop in which it watches for changes in the file and displays them when they happen. It's a simple way to monitor log files, and lots of administrators dedicate several terminals to the purpose of running tail `--f` sessions on various log files. The syntax is simple:

```
tail --follow=name --retry error.log
```

That results in a constantly updated display of the contents of error.log. By specifying `--follow=name` and `--retry`, the command guarantees that tail watches the file itself, not the file descriptor.

IIS

The Microsoft HTTP server handles logging differently. Rather than logging to a file, IIS records its status and error-reporting information so that it is available for examination in the Event Viewer, which is one of the Administrative Tools on a Windows 2000 or XP system.

You'll find IIS errors in the System Log portion of the Event Viewer window, with a source name of W3SVC.

Microsoft offers advice on troubleshooting IIS errors at `http://technet.microsoft.com/en-us/library/cc739055.aspx?ppud=4`.

PHP Error Reporting and Logging

PHP can itself be a tremendous help in spotting errors. Straight out of the box, PHP will report error messages with output — right into the browser window, complete with line numbers. This is as far as most people get with PHP's debugging aids, but it's important to know about the details of configuring error-reporting behavior in order to get the most out of it.

While PHP will show you the line number on which it has detected an error, you have to be aware that that is not always the line to which you should go in order to make a repair. A forgotten closing quotation mark or neglected semicolon sometimes is not picked up by the interpreter until several lines later, so you should be prepared to go back a bit to look for syntax errors of that kind.

Error reporting

When the PHP interpreter places an error message in a program's output (most often resulting in the error message being displayed in a browser window), it's engaging in *error reporting.* Error reporting is a useful diagnostic tool that's turned on by default but that should be disabled on any PHP interpreter associated with a production server.

Error reporting is turned on and off in `php.ini`. The key value is `display_errors`. If you want errors to be rendered as part of your output, this line should appear in `php.ini`:

```
display_errors=On
```

If you do not want errors to be displayed (and you shouldn't want them displayed on any publicly accessible machine), the line should read like this:

```
display_errors=Off
```

If left on in a production server environment, error reporting can result in important details of your software being inadvertently displayed to users. For example, an unexpected condition could cause the name of a variable or a database table to appear in an unsecured browser window. An attacker could use this information to exploit the production server. Figure 31-1 shows an error reported as part of the regular output to a browser window.

Error reporting in browser output

Error logging

Similar in function to error reporting, error logging causes error events to be recorded to a text file, rather than to the screen. It's a more secure option, and because log files should be kept in a directory with limited access, it's the error-recording technique that's preferred for production HTTP servers.

As is the case with error reporting, error logging is turned on and off in `php.ini`. To turn it on, use this option:

```
log_errors=On
```

Alternately, use this:

```
log_errors=Off
```

By default, error logging is disabled in `php.ini`.

CROSS-REF For more detail on error reporting and logging, see Chapter 30.

Choosing which errors to report or log

Whether you choose to use error reporting (on screen) or error logging (to a file), you can specify which errors are considered serious enough to record. In `php.ini`, the `error_reporting` value defines your logging preference. By default, `error_reporting` is set like this:

```
error_reporting=E_ALL & ~E_NOTICE
```

That setting specifies that all errors and warnings are to be reported (a fact denoted by `E_ALL`), and (&) that runtime notices are not to be reported (denoted by `~E_NOTICE` — the ~ indicates `NOT`). Other possible values are included and documented in the numerous comment lines of `php.ini` itself.

The level of reporting defined by `error_reporting` affects the behavior of error logging (as enabled by `log_errors=On`) and error reporting (as enabled by `display_errors=On`) or both simultaneously if both are enabled.

Error-Reporting Functions

PHP, benevolent language that it is, comes equipped with a variety of functions programmers can use to help locate problems and generally report on aspects of their programs' status. These range from ordinary output-generating statements — `print()`, `echo`, and the like — used in contexts that reveal details of variable values, to specialized functions that write to operating systems' logging mechanisms.

This section introduces some PHP functions you can use to spot problems and report on your programs' condition.

Diagnostic print statements

The simplest troubleshooting technique involves placing echo and print statements in your code at key locations, so that the output contains information about the progress of execution through various functions and the values of key variables at different points. This is sort of a poor man's debugger — you can trace variables during execution and see if (and if so, when) they change to some unexpected value.

Here's a simple program that uses `echo` statements for tracing purposes:

```
<html>

<head>
<title>Test</title>
</head>

<?php

function innerFunction($value) {

  echo "<BR>In innerFunction() now...";
  $returnValue = $value . " things<BR>";
  echo '<BR>$returnValue = ';
  echo $returnValue;
  return $returnValue;
}

function outerFunction() {

  echo "<BR>In outerFunction() now...";
  $returnValue = "many";
  echo '<BR>$returnValue = ';
  echo $returnValue;

  return innerFunction($returnValue);
}
```

```
echo "<P>The time has come, the Walrus said, to talk of ";
echo outerFunction() . ".";

?>
</html>
```

Using var_dump()

The usual printing functions are handy, but more specialized ones can prove more useful for debugging purposes.

Chief among these is var_dump(), an extraordinarily clever function that, among other things, will automatically render the contents of an array in a way that's comprehensible to a human reader.

Recall that this code:

```
$stateCaps =
  array( 'New South Wales' => 'Sydney',
         'Victoria' => 'Melbourne',
         'South Australia' => 'Adelaide');
echo $stateCaps;
```

will result in some pretty useless output. It will say simply:

```
Array
```

Not too handy. In contrast, the same array definition, followed by this line:

```
var_dump($stateCaps);
```

results in much more useful output:

```
array(3) { [u"New South Wales"]=> unicode(6) "Sydney" [u"Victoria"]=>
unicode(9) "Melbourne" [u"South Australia"]=> unicode(8) "Adelaide" }
```

It's immediately obvious to the person doing the debugging what the contents (keys and values) of the array are.

Using syslog()

PHP provides a function, syslog(), with which you can write directly into the log of the operating system running your PHP environment. It's a handy function, useful if you want to log all system problems to a standard location or if you want to alert a system administrator who might not be directly involved in PHP development.

Simply put, syslog() allows you to specify the degree of severity associated with the event to be logged and to specify a message describing it. Those values are then written out as an aid to diagnostics.

519

This code illustrates all possible syslog() severity options:

```php
<?php

$logOptions =array(LOG_DEBUG,
                   LOG_INFO,
                   LOG_NOTICE,
                   LOG_WARNING,
                   LOG_ERR,
                   LOG_CRIT,
                   LOG_ALERT,
                   LOG_EMERG);

$exclamations = array('Look!',
                      'Take note!',
                      'Hey!',
                      'Uh-oh!',
                      'Oops!',
                      'Oh No!',
                      'Look out!',
                      'AIYEEEEEE!');

foreach ($logOptions as $key => $value) {

syslog($value, $exclamations[$key]);

}

?>
```

That code results in eight errors being written to the system log.

In a Unix system, PHP syslog() is functionally the same as syslog(3) — refer to its man page (man 3 syslog) — and on a Microsoft Windows system, PHP syslog() writes to the Event Log (specifically, to its Application Log portion).

The error-defining codes, in order of increasing severity, are:

- LOG_DEBUG
- LOG_INFO
- LOG_NOTICE
- LOG_WARNING
- LOG_ERR
- LOG_CRIT
- LOG_ALERT
- LOG_EMERG

In Microsoft Windows, the first three of those (LOG_DEBUG through LOG_NOTICE) are considered informational; the fourth and fifth are considered warnings, and the final three are noted in the Event Viewer as Alerts. All of them are shown with a source value of `c-client`, which corresponds to an ancillary process of the Apache server.

The Event Viewer shown in Figure 31-2 reflects the results of the preceding listing.

FIGURE 31-2

Errors at various severity levels

Logging to a custom location

Under Linux, you can use this procedure to write log messages to a file of your choosing. Modify your `/etc/syslog.conf` file to include this line:

```
local0.debug    /var/log/php.log
```

Then restart the syslog daemon:

```
/etc/init.d/syslog restart
```

With that done, you can refer to `LOG_LOCAL0` as part of an `openlog()` call, as is done here:

```php
<?php
define_syslog_variables();
openlog("CustomLog", LOG_PID, LOG_LOCAL0);
$errorMessage = "Aiyeee! Dying now.";
syslog(LOG_EMERG, $errorMessage);
closelog();
?>
```

The `LOG_PID` argument that supplements `openlog()` states that the process ID of the offending thread should be recorded in the log file with the other details.

Using error_log()

You can use error_log() to send an error message almost anywhere , including to an electronic mail address. It's an easy and convenient way to report on unexpected conditions that crop up in your PHP software, yet few developers bother to use it.

The basic syntax for error_log() is:

```
error_log(message, type [,destination])
```

In that syntax, type has one of four possible values:

> 0: The message is handled according to the setting of error_log in php.ini.
>
> 1: The message is sent by SMTP electronic mail to the address specified by the destination parameter.
>
> 2: The message is referred to a remote debugger.
>
> 3: The message is added to the end of the file specified by the destination parameter.

The following code shows several possible uses for error_log():

```php
<?php

// Writes message as specified by error_log in php.ini.
error_log("Goodness me!", 0);

// Writes message to e-mail address.
error_log("This is not spam.", 1, "webmaster@wiley.com");

// Writes message to a remote debugger:
error_log("Problem!", 3);

// Writes message to a file:
error_log("Save me!", 4, "/log/php.log");

?>
```

In php.ini, error_log is normally not set to anything — it's commented out by default. If you want to use the 0 option to do anything, you'll have to modify the error_log setting to something like this:

```
error_log = syslog
```

which causes error_log() to be the equivalent of syslog(3) on Linux (and thus also the PHP syslog() function) and to write events to the Event Viewer under Microsoft Windows. Alternately, the error_log value can equal a filename:

```
error_log = c:\logs\php.log
```

Summary

This chapter attempted to show you where bugs come from and how to catch them when they find their way into your PHP programs. You saw that there are a number of resources available to you as you attempt to track down problems — many of them native to the PHP language or otherwise freely available.

First among your bug-catching tools is the bug-reporting capability of PHP itself. If you configure `php.ini` correctly — and this chapter showed you how — you can get precise reports of the line numbers on which the interpreter is running into trouble. Depending on your security requirements, you can have the troubleshooting information conveniently displayed with the rest of the output (typically in the browser window), recorded in a log file, or both. You also can keep an eye on the HTTP logs maintained by your web server. These will help you monitor GET request data and spot requests for nonexistent files.

Additionally, you should use PHP language constructs that make your programs more self-diagnostic and troubleshooter-friendly. Functions like `openlog()` and `syslog()` will record event information when problems occur and can really help with tracking down problems. In an even simpler strategy, you can use carefully placed `print()` and (especially) `print_r()` statements to reveal what's going on in your code as it executes.

Troubleshooting PHP is not everyone's favorite activity, but with the right tools and (more important) the correct attitude, it can be fun.

Chapter 32

Learning PHP Style

This chapter is about the major points of PHP style and how it can enhance the functionality, maintainability, and attractiveness of your code. This discussion is intended to help new PHP developers make the main stylistic decisions, most of which are common to all programming languages.

We also hope this chapter may help new PHPers decipher other people's code. It can be very alarming to someone just learning scripting to read three different tutorials, which appear totally incongruent but lack any explanation of the discrepancies. The information in this chapter will help you tease out the functionally important bits of code from the mere stylistic quirks and thus gain a better understanding of what you're seeing.

The Uses of Style

The primary goal of a program is, of course, functionality. After all, if your PHP script chokes, who's even going to care how good it looks? Error messages are never all that stylin'. But there is a vast difference between simply whipping up something that will work and writing well-formed code that can be clearly understood by others.

PHP programmers confront all the same style issues that other programmers do, including:

■ **Readability:** Sure, you understood it when you wrote it, but what about the next person who reads it? What if the next person is you?

■ **Maintainability:** What happens when your health-advice site finally makes that conversion from Fahrenheit to Celsius? (A wrong answer: Replacing 790 occurrences of the string 98.6 in your source code.)

■ **Robustness:** Your web site works fine when it's getting the inputs you expect. What about when it gets the inputs you don't expect?

■ **Conciseness and efficiency:** Fast code is better than slow code, and (others things equal) code using fewer keystrokes is better than code with more keystrokes (but other things are almost never equal).

This chapter will give a quick overview of some strategies for achieving these goals in PHP, before moving on to some code organization issues that are unique to PHP.

Readability

Before a PHP script can aspire to be maintainable or elegant, it has to be human-readable. The human eye likes clear patterns, logical organization, and meaningful repetition. It also helps to have the most significant word or character at the beginning of a line, instead of buried in the middle.

If you develop HTML mostly through use of a WYSIWYG tool, your notions of legibility may be very odd indeed. These programs are notorious for writing badly structured graphics-oriented HTML, filled with invisible GIFs and absolute sizing and other little horrors.

Trying to add PHP directly to HTML files like this is an exercise in frustration, like trying to dance with someone who has no rhythm. If you insist on doing so, remember: It's not PHP's fault, so please direct your abuse to the other vendor. However, in lieu of yet another moralistic Unix-centric anti-WYSIWYG lecture, we'll now try to make a concrete suggestion or two for those who can't totally avoid such tools.

Probably the single most effective step you can take to increase legibility is to run all machine-produced HTML through a utility that will make it more human-readable. It doesn't take very long at all and will improve matters substantially. A good one is `HTML Tidy`, freely available from `tidy.sourceforge.net`.

This utility will also clean up common errors in your HTML source, such as missing end tags. Furthermore, it has some (admittedly limited at this point) capability to cope with PHP, if you've used the standard `<?php ... ?>` tags — so you can also try cleaning up those "I'm in such a hurry, so just this once I'll save a Microsoft Office document in HTML format and then stick in a couple of PHP tags' situations."

A somewhat more labor-intensive approach that gives you finer control is to run the code through an HTML validator. This is a utility (many are web-delivered) that lists all the specific points at which a page is not in compliance with HTML standards. However, unlike `HTML Tidy`, it does not actually rewrite the source code; you can choose to make changes on a point-by-point basis.

Comments

Putting comments in code is just like flossing your teeth: important for health and hygiene, the object of many good intentions, all too often skipped "just this once," and long regretted later if undone.

The problem is that there's no immediate glory to be had from commenting — all the benefits are longer term and diffuse. Let's face it: You rarely hear hackers oohing and aahing over the beautiful commenting of the guy in the next cubicle, and few web sites' go-live dates are allowed to slip so that the programmers can put the finishing touches on their comments. Commenting comes into its own later, when your team leader quits (in the middle of a major site redesign) to join a neo-Luddite community, and the rest of you are sitting around scratching your heads and thinking "Huh?" in unison, as you desperately try to write up some documentation in time for the scheduled release. So what kinds of things should you comment? We feel you *must* explain:

- Anything with future "what the heck was I thinking?" potential (usually due to extreme cleverness or extreme ugliness)
- Anything you suspect might be a temporary expedient
- Anything that will lead to dire consequences if tampered with by someone less experienced.

Things that ideally should be noted include:

- The date the file was originally created and the name of the creator
- The date the file was most recently altered, the name of the alterer, and possibly an explanation of the rationale behind the alteration
- Any other files or programs that depend on the existence of this file
- The intended purpose of the file and of its constituent parts
- Things you might want to mention in documentation you're planning to write later
- The reason you want to save something that isn't being used (alternate versions, archive copies, and so on), conditions under which it might become okay to throw it away, or your plans for what to do with it

Obviously, you're in a better position than we to decide whether these items are strictly necessary. If you're using PHP for a very small, purely personal site, maybe commenting would be superfluous, but the bigger and more complex the site, the more you need to annotate your own work. In theory, it's possible to overcomment, but, in practice, few programmers are guilty of this offense.

> **TIP** As we detailed in Chapter 4, there are several styles of PHP comments. Remember that none of these will be visible from the client machine, even when HTML source is viewed. If you want client-readable hidden text, you must use HTML comments.

PHPDoc

For very large and complex programs, code-embedded comments are not sufficient. You want separate documentation that someone can read without delving into the code itself. The problem with documentation like this, however, is not just how anyone gets the time to write it (since that is often low priority), but how it stays in sync with the code (which is even lower priority). When starting a new day job, we have more than once confronted a very common choice: Should we get to know the code by reading the current code itself or by consulting some very nicely written documentation of the code as it was two years ago?

527

One approach, used with some languages, is to employ a tool that produces documentation by extracting specially formatted, embedded comments from the code. For example, if you have followed a given commenting convention, you can point the `javadoc` tool at your Java code and it will extract class and method comments into a set of HTML pages documenting the API. This is not a magic solution for the problem of keeping docs in sync with code. (It will break down, for example, if people begin writing new methods by copying old methods, and leaving the original comments in place.) But at least developers have to write only one description of a given method rather than two.

There is an analogous `phpdoc` tool that uses PHP (naturally) to scan PHP code for special comments, producing HTML output. If you are doing a large-team project, though, especially one making heavy use of object-oriented PHP, you might find `phpdoc` to be helpful. For more on `phpdoc`, see `www.phpdoc.org`.

File and variable names

Some people act like thinking up variable names is equivalent to being forced to write an epic poem — they go into a kind of writer's block and become creatively incapacitated. For instance, we once had an intern who was apparently unable to think up a single name or even a halfway decent scheme for doing so. This person's habit was to name every new file according to simple sequential order: `file16.html`, `file17.html`, `file18.html`, and so forth. Each variable on a web page was called `var1`, `var2`, and so on. This story would be a lot funnier if it had happened to someone else.

Because PHP generally requires a lot more variables than HTML, you need a robust naming scheme for all occasions. The following sections include a few tips.

Long versus short

Longer is generally better because it's more informative. You can break up long names with underscores or capitalization if necessary.

Even though most filesystems technically allow for long filenames, the results are not pretty when viewed as icons — so GUI users may be consciously or unconsciously averse to using long filenames. Icon labels are usually quite short and, thus, naturally lend themselves to very concise filenames. Try giving a file a long, complex name (like `PoachedPeachesRecipe.php`) and putting it on your desktop — the result is just viscerally displeasing.

CAUTION **Most GUI-oriented filesystems allow and even encourage filenames with spaces in them (for example, `My Document.doc`). Unix systems do in theory, but in practice it's not such a good idea. PHP will try to cope gracefully with such filenames, but it may not be able to do so in all situations.**

One benefit of using PHP for dynamic content generation is that you can use shorter filenames that will be expanded and differentiated by `GET`-style query strings. For example, a static site might use this style of filename to uniquely identify each page:

```
FeatureHitchcockBirds.html
MiniseriesIrvinSpy.html
```

A dynamic site, on the other hand might identify the same pages like this:

```
feature.php?ID=1
miniseries.php?ID=2
```

In this situation, you can have the best of both worlds: short filename plus unique identifier. (The only downside of this is that some search engines still discriminate against pages with dynamic arguments, under the theory that the contents are likely to change with every page view and, therefore, won't be worth indexing.)

PHP sets no particular limit on the length of variable names. So feel free to invent lengthy but informative variables like `$AddressOfClientCompanyInSasketchewan`. Hey, it's your script — we're just living in it. You only need to be careful if you plan to use a lot of long-name variables as part of a `GET`-method form.

Underscores versus camelcaps

There are two typical ways to break up long variable and file names in Unix. Underscores look like this:

```
$name_of_favorite_beer
```

whereas *camelcaps* look like this, with the internal capital letters giving the name a humped profile:

```
$NameOfFavoriteBeer
```

It's a purely personal preference, which style you use (unless you have agreed on a particular style with your colleagues). PHP itself uses underscores (`$PHP_SELF`), but this usage by PHP could be construed as an argument in favor of either scheme for PHP programmers. Just remember you can't use dashes and should be careful with dots.

CAUTION   **Unix filenames are case sensitive all the time. Filenames in other OSes, such as Windows, are not case sensitive. If you might be in a position to move PHP files between OSes, be careful.**

The main thing to strive for here is consistency. It's frustrating to spend a lot of time trying to figure out why `$My_Number` was never assigned, only to find out that it's because you called it `$MyNumber` when you assigned it.

Reassigning variables

Situations arise in which you deliberately want to keep using the same variable name over and over rather than coming up with new names. This happens when you need to be certain only one variable of a particular type will be valid at any given time. For instance, you might want to be sure there can be no confusion about which of two database queries will be used for an operation, which you can ensure by using the same name for both (for example, `$query`). PHP will overwrite the former with the latter, and your variable will always be minty fresh.

Uniformity of style

Although we have talked in a very free-and-easy way about how all these stylistic choices are up to you, there are situations where it is actually good to have a consensus on what code should look like and then enforce that. This is particularly true when many programmers will contribute code to a project. The reasons that are usually advanced for a uniform style are:

- It makes it easier to read code from multiple programmers, because you don't have to get used to a new indenting or layout style every time you see new code.
- It makes life easier for version control software (like CVS). If I change a code file that you created and my editor changes the indentation, there will be a lot of apparent but spurious differences.

The closest thing PHP has to a consensus style is the coding standard developed by the maintainers of the PEAR project. See Appendix E for a discussion of both PEAR and this coding style.

Maintainability

Many seasoned programming veterans, especially those who are also managers, tout the importance of maintainability above that of any other virtue.

The problem is, of course, that maintainability is in direct conflict with all the other goals — especially speed. When Internet Time gets into the ring with Hypothetical Future Code Maintenance by Someone Probably Not Myself, everyone knows how the story is going to end. Still, the main mental mantras of maintainability are worth keeping in mind:

- The things that are most likely to be changed should be the easiest to find.
- Changing those things should not have unpredictable effects.
- Each change should have to be made in only one place.

Avoid magic numbers

A *magic number* is a numerical value that might someday have to be changed but is buried deep in code, often in multiple places. Imagine, for example, these lines of code found in your bank's hypothetical PHP-based web site:

```
print("The interest rate on your CD can be as high as
        5.5%!<BR>");
$sample_gains = 5000 * 1.055;
print("After a year, a \$5000 investment could grow to
        \$$sample_gains!<BR>");
```

Now, when times get tighter and the rate goes down to 5.0 percent, someone has to find and change every instance of the rate. So, someone does a text search for 5.5, which misses the 1.055 in the second line here, and now your bank is engaging in false advertising.

For simple sites, a better alternative can be as easy as using an $interest_rate variable, which is assigned very visibly at the top of a script — a change in rate means a change only to that assignment statement. More complex sites might produce their pages as function calls, with variables like $interest_rate being passed in as an actual parameter. Finally, some sites will go so far as to have all their content imported from a database, so that no piece of information has to ever be changed directly in code.

Functions

Having tried to maintain a complex site using a web-scripting language that did not support functions, we can say from our own experience that functions are crucial to maintenance. The art of procedural abstraction via functions needs a book in itself, but here's some brief advice:

- Always look for opportunities to bundle naked PHP code into a function, especially in cases where it might be reused.

- Try to keep function definitions short — if a definition gets too long, break it up into multiple functions.

- Always load all your function definitions before any code that calls any functions.

Include files

One of the great benefits of dynamic web page generation over static HTML is the opportunity to fight redundancy. Anyone who has ever managed a static site of any size knows how much of each file is boilerplate — and even editing a single character on each page isn't a picnic if your site has 200 pages.

PHP makes it very easy to drop anything into your scripts, from one character to a whole separate program, by using the built-in include or require functions. The syntax is simply:

```php
<?php include("filename.ext"); ?>
```

You can also use a variable filename, like this:

```php
<?php
$LastName = "Park";
include("$LastName.inc");
?>
```

which will result in the contents of the file Park.inc being spliced in at the location of the include statement.

You can use any extension you want for the included file. Popular choices include `.txt`, `.inc`, and even `.html` to remind yourself that the file will show up in HTML mode.

A few things to remember:

- PHP will drop the entire text of the file into your PHP script *in HTML mode* (as explained in Chapter 4). If the included file is itself meant to be parsed as PHP, you must use valid PHP tags at the beginning and end. If any part of the file is meant to be parsed as PHP, you must use valid PHP tags around that section.

- Recall the difference between `include/require` and `include_once/require_once`. In general, if what you are including or requiring is a set of function or class definitions, you should use the `once` variant. If it is straight PHP or HTML, then which variant you load depends on whether you would ever want that literal block to repeat in your output; if not, then the `once` version is probably still what you want.

- `include` can also be used to assemble complex web pages from text files instead of from a database. In some cases, this can even be faster — usually when the included data is just a sizable text file(s). However, after you go to the trouble to make a database connection for any reason, it's probably just as fast to store your big chunks of text there, too.

Object wrappers

Although we haven't covered PHP's object system in detail yet, it's worth noting that consistent use of objects can make code more maintainable, much as functions do. For example, some developers of database-enabled PHP sites are disciplined enough to wrap up all of their database-specific functionality in the methods of an object, so that the rest of their code doesn't even know what kind of database is supporting the site. In theory, then, if they decide to move from a mySQL database to an Oracle database, only the object-level code will have to be changed.

CROSS-REF   **For details on PHP's support for object-oriented programming, see Chapter 20.**

Consider using version control

For large multiprogrammer projects in industrial settings, version control isn't "something to consider" — it's a must. Similarly, large decentralized open-source projects could not survive without *CVS* (Concurrent Versions System) or *SVN* (Subversion). Even if you are working by yourself or with one other person on a hobby project, using version control can free you to do more experimentation, secure in the knowledge that you can get to the older versions of your code if something goes awry.

See `www.cvshome.org` for more information on CVS. SourceForge also offers free web-based CVS project hosting for open source projects (`www.sourceforge.net`).

Robustness

The two commandments of robustness are:

- Code should detect unexpected situations and respond gracefully rather than dying.
- If code must die, better that it die informatively.

Writing robust code is at first a difficult task of imagination, where the programmer tries to think ahead to all the things that might go wrong and to cover those cases. The ideal situation is for that habit of mind to become a habit of code, so that the coder has a standard set of tests that wrap around the standard potential problems. Although most of the robustness issues in PHP are the same as in any language, there are two kinds of situations to cover that are more specific to PHP: problems with an external service, and problems having to do with variable type.

Unavailability of service

PHP is in part a "glue" language, offering a single environment where a variety of different code libraries and external services can be invoked. Any given PHP page might open a file, connect to a database, query an LDAP server, send an HTTP header, or send mail via an SMTP server. The important habit to develop is covering cases where for some external reason a service is unavailable, or times out, or behaves oddly, or gets interrupted in the middle.

Often, such services have error states that can be retrieved and printed if the only option left is to informatively die. For example, a reasonable construct for making a connection to a mySQL database is:

```
$connection = mysql_connect([arguments]) or
              die("Connection failed:  $php_errormsg<BR>");
```

This is preferable to the weird and unexpected errors you would see if your code went happily ahead assuming that it had a live database connection. An alternative that is better from a security point of view is:

```
$connection = mysql_connect(...) or error_log($php_errormsg);
if (!$connection){
   die("Connection Failed");
}
```

because this will avoid displaying interesting facts about your PHP and database configuration to the user's browser.

Even better style is to use the exception-handling facility introduced in PHP5, rather than simply failing with die(). Exceptions can be thrown whenever a problematic condition is encountered, and recovered from, at a single point in the code (if you so choose). If it is possible to recover from the problem, exceptions make it easy to structure your code to support that. If the script must die

anyway, exceptions make it easy to propagate the negative information and display it at the right time, rather than just aborting execution.

CROSS-REF   Exceptions and exception handling are covered in detail in Chapter 30.

Unexpected variable types

Although the type-looseness of PHP is for the most part a good thing, it leaves a little bit of uncertainty for the programmer about exactly what type a variable or value will turn out to be. Unless you come to know all the type conversion rules very well, it can be surprising to have code that is accustomed to strings suddenly run across a value that is a number, all because some PHP construct decided that any string composed only of numerical digits must really be a number at heart. One interesting robustness check is to use a text editor to search your code for $ (thereby finding every variable) and ask yourself for each one what would happen if the type turned out to be surprisingly different.

Efficiency and Conciseness

Efficiency and conciseness are not the same thing. Efficient code runs using a small amount of execution time or computer memory, while concise code accomplishes a given task in a small number of lines or keystrokes. In this section, we give some quick tips toward writing efficient and concise PHP code, along with our extremely opinionated commentary about in what senses these goals are worth striving for.

Efficiency: only the algorithm matters

There was a time when computer memory and computer cycles were so precious that it was worth a lot of effort to boil down your code to the smallest number of resulting machine instructions possible. This is still true in certain areas of software development (kernel programming, graphics libraries), but for most development tasks, saving a few instructions or a few K is not worth backing off on any other goal. This is especially true for web scripting, where there is always going to be some overhead of purely Internet-related execution delay. If it takes half a second for a user to fetch your page, regardless of how your page is produced, then an extra five milliseconds on the server side will be lost in the noise.

With that said, there's one variety of efficiency that matters and will probably always matter: the broad algorithm or approach that your code uses for a task. For example, if your code locates a name in a database by querying the database for all names and then doing a string comparison for each name to see if it's the one you want, you'll soon find out how much efficiency can matter.

Efficiency optimization tips

Here are some quick mantras to repeat as you code.

Don't reinvent the wheel

It's usually a bad idea to write code that duplicates a language-level facility, unless it's for purposes of fun or education. For example, any programmer worth his or her salt should write sorting routines at some point in their education, but no programmer should have to keep writing them (unless it is actually in their job description). Most high-level programming languages offer some kind of sorting capability (either as part of the language or in a library), and it's very likely that the programmer who wrote them did a better job than you will. PHP is no exception here — the array type supports several types of sorting, and most of the databases supported by PHP have sorting options built into the query language. Either of these options will be faster and more reliable than what you get by rolling your own.

Discover the bottleneck

Although it's good to try to use efficient algorithms from the beginning, it's often not worth doing other kinds of optimization until you find out that too much of some resource is being used. At that point, you want to tighten things up, and you'll get the most reward for your effort if you focus on the piggiest parts of your code. Most code follows the 90/10 rule: 90 percent of the time is spent in 10 percent of the code, and you want to locate that 10 percent.

One technique that programmers often use to locate that 10 percent is called *profiling*. A profiler is a utility that tracks code as it runs, noting the time spent in every function call, and producing a neat summary of the results. Unfortunately, at this writing, there is no good general profiling utility for PHP (although you may want to look at XDEBUG with Kcachegrind). So the best bet for now is the poor man's profiling technique: printing the value of the function call `microtime()` in various places in your script to see where the time is going. If the 90/10 rule is in effect, the time sink will usually be glaringly obvious.

Focus on database queries

Although we cover database efficiency in more detail in Chapter 18, you should be aware that database queries are usually the biggest time sink for PHP sites that have database backends. Especially if your database-enabled site doesn't do a lot of other computationally intensive work, your first suspicious glance should be at the queries, and your next task should be to try to identify a query that is particularly time-consuming. After you've identified a guilty query, there are a host of techniques available to speed that query up, many of which don't have anything to do with PHP.

CROSS-REF   **For details on optimizing database-enabled PHP code, see Chapter 18.**

Focus on the innermost loop

Let's say that you have a page with embedded looping constructs, like the following:

```php
for ($x = 0; $x < 100; $x++)
{
  do_X();
  for ($y = 0; $y < 100; $y++)
    {
      do_XY();
```

```
        for ($z = 0; $z < 100; $z++)
          {
            do_XYZ();
          }
      }

  }
```

Unless you have a really good reason to think otherwise, your optimization focus should be on the function do_XYZ() (which will execute 1,000,000 times) rather than on the other two functions (10,000 times and 100 times).

Conciseness: the downside

Before we get into how to write more concise code, let us say that we think conciseness is an over-rated virtue, for the following reasons.

Conciseness rarely implies efficiency

Although it's true that somewhere in the guts of the PHP engine, the characters of the code you write are being consumed one by one (and so, in theory, more code takes more time), in practice, the Zend-based parsing engine of PHP is so zippy that the number of characters just doesn't matter. Ditto for the time or space consumed in extra variable assignments or the overhead of extra function calls.

You should generally strive to reuse code whenever possible. This means thinking through the design of your code and using functions where it makes sense rather than repeating code.

Conciseness trades off with readability

Remember that every keystroke you omit might be the keystroke that would have let someone figure out what the heck you were thinking when you wrote the code. For example, take a look at the following admirably concise function:

```
function sieve($n) {
  for ($i = 2; $i <= sqrt($n); $i++)
    for ($j = $i, $ind = $i * $j; $ind <= $n;
         $j++, $ind = $i * $j)
      $carray[$ind] = 1;
  for ($i = $n, $plist = array(); $i > 1; $i--)
    if (!$carray[$i]) array_push($plist,$i);
  return($plist);
 }
```

Obviously, this implements the Sieve of Erasthones, and $plist is a list of all the prime numbers less than $n. Obviously.

So why do programmers strive for conciseness? The first reason is that it saves them time (but only at the time of actual code writing). The second reason (and we're only half-joking) is that they're afraid some other programmer (probably one trained in C) will come along later, laugh at them, and point out that their code could have been written in only half the space.

Conciseness tips

If you must write code that fits in less space, try some of the following techniques.

Use return values and side effects at the same time

It's a very common trick to exploit the fact that the value of an assignment is the value assigned, as in the following pseudocode:

```
while ($next = GetNextOne())
    DoSomethingWith($next);
```

where `GetNextOne()` is some function that returns useful values in sequence and then returns a false value when it runs out of them. When a false value is returned, `$next` is false, and the `while` loop terminates.

Use incrementing and assignment operators

The incrementing operators (`++` and `--`) shorten statements that involve adding or subtracting one from a variable, and the combined assignment operators (`+=`, `*=`, `.=`, and so on) make certain kinds of assignments more concise.

CROSS-REF **The incrementing operators and the arithmetic assignment operators are covered in Chapter 9, and the combined string assignment operator (`.=`) is covered in Chapter 7.**

Often these operators are used in combination with the previous trick, as in:

```
while ($count--)
    DoSomethingWith($count);
```

which (assuming that `$count` starts as a positive integer) would call its function for the very last time on the value 1.

Reuse functions

This is one case where conciseness is good, because functions are good. If you can identify any stretches of code that get duplicated in your pages, try to replace each one with a call to a single function that packages up that code. Your code will be shorter by the amount of the duplication and also easier to maintain.

There's nothing wrong with Boolean

Beginning programmers often have an odd distrust of Boolean values, not realizing that they can be passed around as freely as any other kind of value. This leads to code that wastes a lot of space, like the following:

```
function DivisibleByBad($num1, $num2)
{
    if ($num1 % $num2 == 0)
       return(TRUE);
    else
       return(FALSE);
}
/* using the function */
if (DivisibleByBad(9, 3))
  $divisible_result = TRUE;
else
  $divisible_result = FALSE;
if ($divisible_result == TRUE)
  print("It's divisible!<BR>");
else
  if ($divisible_result == FALSE)
     print("It's not divisible!<BR>");
```

A more concise version would look like:

```
function DivisibleByBetter($num1, $num2)
{
    return ($num1 % $num2 == 0);
}
/* using the function */
if (DivisibleByBetter(9,3))
  print("It's divisible!<BR>");
else
  print("It's not divisible!<BR>");
```

You could obviously take this one step further and get rid of the function itself, like this:

```
if (9 % 3 == 0)
  print("It's divisible!<BR>");
else
  print("It's not divisible!<BR>");
```

But (once again) *Functions Are Good* — an explanatory function name is a little piece of documentation in itself, and any function you write gives you a chance to reuse it later, which, in turn, makes your code more maintainable.

Use short-circuiting Boolean expressions

Certain kinds of Boolean tests aren't safe to apply until you've done other tests. It's tempting to deal with this by insulating the problematic tests with `if` constructs. For example, imagine that you want to print the ratio of two variables that are bound to integers, but only if they are bound to integers and only if the ratio is greater than two. Also, you want to avoid a division-by-zero warning. You might overcautiously write:

```
if (IsSet($x))
  {
    if (IsSet($y))
    {
      if (Is_Integer($x))
      {
        if (Is_Integer($y))
        {
          if ($y != 0)
            {
              if ($x / $y > 2)
                  print("Ratio is " . ($x / $y));
            }
        }
      }
    }
  }
```

You can be equally overcautious and still type a little less, as in:

```
if (IsSet($x) && IsSet($y) && Is_Integer($x) &&
    Is_Integer($y) && $y != 0 && $x / $y > 2)
  print("Ratio is " . ($x / $y));
```

The tests will be applied in left-to-right order, and if any test fails, the tests to the right of it will not be evaluated.

HTML Mode or PHP Mode?

There's a spectrum of ways to combine PHP and HTML, functionally all pretty much the same. Your choice will depend on extrinsic factors such as your particular team's workflow.

The easiest way to demonstrate all this is to simply write the same script in minimal PHP, maximal PHP, and medium PHP styles. We will also include a version using the *heredoc* construct (discussed in Chapter 7). Remember, these are equally correct and return much the same result. The stylistic decision is just a matter of preference and consistency, and (sometimes) slight differences in functionality.

Minimal PHP

The code in Listing 32-1 shows a simple self-submitting form, which returns some simple information about days of the week. First, it tells you what day it is today and then gives you a chance to ask it what day it will be in a few days (using a form with a pull-down list). The typical text output looks like this:

```
4 days from this moment it will be Thursday
Today is Sunday
Please choose a number of days, and we'll tell you what day it will be
that many days from now
```

It is written using a *minimal PHP* style, meaning simply that, as much as possible, the code is pure HTML, dropping into PHP mode only when dynamic data must be displayed (such as the current day of the week or an answer that depends on submitted data).

LISTING 32-1

calendar.php - An example with minimal PHP

```
<HTML><HEAD TITLE="Calendar Server"></HEAD><BODY>
<H2>Welcome to the Calendar Server</H2>

<?php if (IsSet($_POST['DAYS'])) { ?>
  <P> <?php echo $_POST['DAYS'];?>
  days from this moment it will be
  <?php $date = getdate(time() +
                    ($_POST['DAYS'] * 86400));
       echo($date['weekday']);
}?>

<P>Today is <?php $date = getdate();
                echo($date['weekday']); ?>
<P>Please choose a number of days, and we'll
tell you what day it will be that many days from now:

<FORM METHOD=POST ACTION="calendar.php" >
<SELECT NAME=DAYS>
 <OPTION VALUE=1>1<OPTION VALUE=2>2<OPTION VALUE=3>3
 <OPTION VALUE=4>4<OPTION VALUE=5>5<OPTION VALUE=6>6
</SELECT>
<INPUT TYPE=SUBMIT NAME=SUBMIT VALUE=SUBMIT>
</FORM>
</BODY></HTML>
```

This code takes the minimal style to an extreme — note the funny business with the `if` statement near the top, where some straight HTML text (`days from this moment`) is included conditionally, based on the results of the PHP statement that precedes it.

540

As we have said, this is a matter of taste, but we don't like this version very much. It's a bit hard to see exactly what is being produced by the PHP snippets that are spliced into the page.

Maximal PHP

At the opposite extreme, consider the code in Listing 32-2. This version is in PHP mode all the time and simply prints all the HTML it needs to as it goes.

LISTING 32-2

calendar.php An example with maximum PHP

```php
<?php
print("<HTML><HEAD TITLE=\"Calendar Server\"></HEAD><BODY>");
print("<H2>Welcome to the Calendar Server</H2>");

if (IsSet($_POST['DAYS'])) {
  print("<P>" .$_POST['DAYS'] .
        " days from this moment it will be ");
  $date = getdate(time() +
          ($_POST['DAYS'] * 86400));
  print($date['weekday']);
}

$date = getdate();
$day_of_week = $date['weekday'];
print("<P>Today is $day_of_week");
print("<P>Please choose a number of days, and we'll
tell you what day it will be that many days from now:");

print("<FORM METHOD=POST ACTION=\"calendar.php\" >");
print("<SELECT NAME=DAYS>");
for ($i = 1; $i < 7; $i++) {
 print("<OPTION VALUE=$i>$i");
}
print("</SELECT>");
print("<INPUT TYPE=SUBMIT NAME=SUBMIT VALUE=SUBMIT>");
print("</FORM></BODY></HTML>");
?>
```

Another term for *maximal PHP* might be *CGI-style*, because we are not taking advantage of the HTML-embeddedness of PHP at all, and might as well be writing a CGI script in C or Perl (and who wants to do that?). Again, a matter of taste, but it's a little bit hard to visualize the structure of the HTML page that will result from running this.

Medium PHP

An intermediate version that better exploits functions is shown in Listing 32-3. It spends about half its text in PHP mode, defining functions, before dropping back to a minimal style for the rest of the script.

LISTING 32-3

calendar.php - An example with medium PHP

```php
<?php
function maybe_print_answer_date () {
  $seconds_in_day = 60 * 60 * 24;
  if (IsSet($_POST['DAYS'])) {
    print("<P>" .$_POST['DAYS'] .
          " days from this moment it will be ");
    $date = getdate(time() +
          ($_POST['DAYS'] * $seconds_in_day));
    print($date['weekday']);
  }
}

function print_day_options () {
 for ($i = 1; $i < 7; $i++) {
   print("<OPTION VALUE=$i>$i");
 }
}

function get_day_of_week($time) {
  $date = getdate($time);
  return($date['weekday']);
}

?>
<HTML><HEAD TITLE="Calendar Server"></HEAD><BODY>
<H2>Welcome to the Calendar Server</H2>
<?php maybe_print_answer_date(); ?>
<P>Today is <?php echo get_day_of_week(time());?>
<P>Please choose a number of days, and we'll
tell you what day it will be that many days from now:

<FORM METHOD=POST ACTION="calendar.php">
<SELECT NAME=DAYS>
  <?php print_day_options(); ?>
</SELECT>
<INPUT TYPE=SUBMIT NAME=SUBMIT VALUE=SUBMIT>
</FORM>
</BODY></HTML>
```

Maybe the medium style is a little bit more verbose, but to our eyes, it's also a little easier to read and modify than the minimal and maximal styles.

The heredoc style

Listing 32-4 shows a rewrite of the medium style, using the heredoc syntax for constructing strings.

LISTING 32-4

Heredoc style

```php
<?php
function answer_string ($days) {
  $seconds_in_day = 60 * 60 * 24;
  $return_string = "";
  $day_string =
    day_of_week_string(time() + $days *
                       $seconds_in_day);

  $return_string .=
    $_POST['DAYS'] .
    " days from this moment it will be " .
     $day_string;
  return($return_string);

}

function day_of_week_string($time) {
  $date = getdate($time);
  return($date['weekday']);
}

function calendar_form_string () {
  $option_string = "";
  for ($i = 1; $i < 7; $i++) {
    $option_string .= "<OPTION VALUE=$i>$i";
  }
  $self_string = "thisfile.php"; $return_string=<<<EOT
<FORM METHOD=POST ACTION="$self_string" >
<SELECT NAME=DAYS>$option_string</SELECT>
<INPUT TYPE=SUBMIT NAME=SUBMIT VALUE=SUBMIT>
</FORM>
EOT;
  return($return_string);
}

// set up string variables
$answer_string = IsSet($_POST['DAYS']) ?
```

```
                 answer_string($_POST['DAYS']) :
                 "";
$day_of_week = day_of_week_string(time());
$form_string = calendar_form_string();

// set up page string
$page_string=<<<EOT
<HTML><HEAD TITLE="Calendar Server"></HEAD><BODY>
<H2>Welcome to the Calendar Server</H2>
<P>$answer_string
<P>Today is $day_of_week
<P>Please choose a number of days, and we'll
tell you what day it will be that many days from now:
$form_string
</BODY></HTML>
EOT;

echo($page_string);
?>
```

Tim Perdue of SourceForge has made the stylistic argument that it's a bad idea for functions to print output to the browser. The heredoc example follows this advice and uses functions only for calculations and for building strings that are returned. The heredoc construct is used twice, first to build a string corresponding to the form that will be displayed and then (using the form string) to build the string corresponding to the entire page. The last act of the script is to echo out the string it has constructed.

This is the most verbose of the four, but the heredoc syntax has the advantage (over the maximal style) that we never have to do any escaping of quotes. Another advantage (over the minimal style) is that we can simply include variables in our page template without dropping out of HTML mode to do so. It is probably the best version of the four at separating logic from page structure.

Separating Code from Design

Many of the topics in this chapter have obvious implications for the separation of code and design. Here are a few additional techniques we should mention.

Functions

As you can see from our Medium PHP example in Listing 32-3, using self-defined functions can be a very flexible and powerful formatting tool, as well as one of the things that make PHP better than a tag-based scripting language.

Cascading style sheets in PHP

As you doubtless already know, there are four generally accepted ways to apply styles to your web pages:

- By applying CSS formatting to individual tags.
- By using `<STYLE>` tags (optionally inside a pair of HTML comment tags).
- By using `<LINK>` tags.
- By using `@import`.

In PHP, you could also use the `include` function to apply styles in a nonstandard way, although it's not clear how much of a gain this would be. For instance, you could `include` a text file containing everything between the `<STYLE>` tags, instead of linking to an external style sheet.

We should also mention the anti-style sheet, a practice almost as long-deprecated as it is common: using outdated HTML tags such as `FONT`, `BGCOLOR`, and `ALINK`. Although you shouldn't do it at all, PHP can help you do it more efficiently if for some reason you must. This usage, for instance:

```
<FONT FACE="<?php include("fontlist.txt"); ?>" SIZE=+2>
```

would at least allow the poor, overworked web developer to change the fonts throughout the whole site with a single edit of the text file. Not that we can condone this kind of thing! Only slightly less kludgy would be this alternative:

```
<P STYLE="font-family: <?php include("fontlist.txt"); ?>">Text here</P>
```

Templates and page consistency

As you can now imagine, PHP allows a wide variety of approaches to site design, which you can fit to your particular style and the organization of the people who work on the site. If your techies can't talk to your artists, you may want to set things up so that they never touch the same files; if you're a tech artist, you may express yourself by the very intermingling of code and graphics. If your site has a large number of pages or is very content rich, you may find (as we have) that it's helpful to choose a particular kind of file organization or template, and stick to it across the site. One simplified example follows, which is similar to templates we have used on www.mysteryguide.com and www.sciencebookguide.com.

```
<?php
/* load general functions */
include("general-functions.inc");
/* load functions specific to this page */
```

```
include("renaissance-functions.inc");
/* page-wide variables */
$PageTitle = "Painters of the Renaissance";
$db_connection = make_database_connection();
?>

<HTML>
<HEAD>
<TITLE>
   <?php print("$PageTitle"); ?>
</TITLE>
</HEAD>
<BODY>
<H3>
    <?php print("$PageTitle"); ?>
</H3>
<TABLE>
<TR><TD>
   <?php print_left_side($db_connection); ?>
</TD><TD>
   <?php print_right_side($db_connection); ?>
</TD></TR>
</TABLE>
   <?php print_footer($db_connection); ?>
</BODY><HTML>
```

In this example, every page loads the same file of sitewide utility functions, then loads a file of functions specific to that page, then defines variables that will be global for the page, and finally intersperses PHP commands in some boilerplate HTML. The content is in columns, and the actual content displayed depends on the particular page's functions, which always have the same names, but with definitions varying for each page. Changing what's displayed in the columns means either changing the per-page functions or (more likely) modifying the database contents. It would be possible for a nonprogrammer to do some limited design on this page by operating directly on the HTML and being careful to leave the PHP alone.

The preceding example is just one simplified possibility from a range of ways to divide up the labor of displaying a PHP page. Another that we like even better is the heredoc technique that we discuss in the section, "The heredoc style," earlier in the chapter. Your particular strategy will depend on the type of site, the size of the site, and the styles of the people contributing.

Finally, note that all these strategies really just adopt a convention about separating logic and display in PHP. If you need an even stronger distinction, there are PHP-based templating systems available that further insulate the display people from the innards of program logic. One example is YATS (Yet Another Template System), available at `http://yats.sourceforge.net`.

Summary

Most of the elements of PHP style are desirable in any programming language. You want to write readable code, with appropriately abstracted functions, consistent indentation, and explanatory comments. You want to stay away from magic numbers, "cloned" code repetition, overuse of global variables, and cryptically clever tricks. Your program should work on the inputs you expect, do something reasonable with inputs you didn't expect, and have the grace to die informatively in situations you *really* didn't expect.

Some of the PHP-specific style issues have to do with organizing file inclusions, how intimately you mix your PHP with your HTML, and more generally the separation of code from design. A wide range of styles are okay here, but you should strive for page-level and sitewide consistency.

Part IV

Other Databases

Chapter 33

Connecting PHP and PostgreSQL

You might find that MySQL or even simple text files meet all your data storage and retrieval requirements. Nothing about a simple flat data structure in a small quantity demands a relational database model. However, as we mentioned earlier in this book, you *do* have choices when it comes to databases. In the next chapter, we'll look at a commercial offering, Oracle. In this chapter, we'll look at what is possibly the granddaddy of the free/open source database alternatives, PostgreSQL (pronounced "post-gress-q-l" or sometimes "Postgrey").

Why Choose PostgreSQL?

This is the part where the open source purists start waving their hands in the air and yelling with uncontrolled excitement! And the excitement is easy to understand — PostgreSQL is a true open source database, made available under the simple and portable BSD license. You can read the almost vanishingly short text of the license at `www.postgres.org/licence.html`.

Are you back yet? See, we told you it was short. So reason number one is not so much the license itself as the freedom from an 85 page EULA laced with sneaky provisions that nobody alive really understands, and that, in many cases, you can't even see until you get the box open. By then it's too late — all the money's gone. Which brings us to reason number two: PostgreSQL is free. We don't mean "free on Mondays, Wednesdays, Fridays, and the vernal equinox," nor do we mean "free to the right people," nor even the more conventional and arguably understandable "free for noncommercial use." PostgreSQL is completely and totally free (unless the developers change their minds).

551

The term free applies to more than just the cost. As with the GNU General Public License, you can alter, repackage, and redistribute PostgreSQL as a standalone product or with your own applications. Arguably better than the GPL for businesses, using and distributing PostgreSQL will not "infect" all your code with the copyleft.

Finally, PostgreSQL supports some nifty special features and elements of the ANSI SQL92 and SQL99 standards that simply aren't available or fully developed in other databases, as well as the ability to work with object and hierarchical data.

Of course, there are some disadvantages as well. First, consider PostgreSQL's ability to work with object and hierarchical data. Wait a minute; didn't we just sell that one as a feature? To the PostgreSQL lovers in the audience, it may seem odd to sell these "features" as disadvantages. Don't flame us yet. We're just taking a moment to expound on the Keep It Simple Stupid (KISS) philosophy. You don't need a sledgehammer to drive in a picture nail, and you don't need an object relational database to store addresses and phone numbers.

Some folks will debate this point, but simply put, PostgreSQL is not as fast as MySQL under low-load circumstances. At some point, PostgreSQL's more robust design will offer performance advantages in really large data sets, but on average, MySQL is just that little bit perkier for lots of reads. Again, if you have simpler data storage and retrieval needs, there's no need for you to go swimming around in the PostgreSQL waters.

Finally, PostgreSQL is more complicated. Permissions management, for one, is not as clear-cut as it is with MySQL. PostgreSQL also offers some features that may cause the novice's eyes to glaze over — schemas and stored procedures, while useful, aren't strictly necessary and you may find them to be clutter. Some people work much better when there are no nonessentials on their desks; the same principle may hold true here.

All these things said, on balance PostgreSQL is a great tool — a superlative tool even — for many jobs. Its user base may be smaller than MySQL's, but its devotees tend to be very loyal. We can't cover it very comprehensively and still stay within the focus of this book, but if you've gotten this far, weighed all the benefits and disadvantages, and you choose PostgreSQL, the rest of this chapter should set you off in the right direction.

Why Object-Relational Anyway?

Object relational databases (ORDBMS) are a relatively new class of product compared to the relational database model that was developed in the early 1970s. In addition to implementing the relational model discussed in Chapter 13, ORDBMS borrow from pure object databases, which excel at handling media objects, spatial, and series style data. An ORDBMS implements object properties on the components of a relational database in order to have the benefit of both worlds. This, in turn, facilitates interaction with the object features of PHP. Because PHP's object model is significantly more powerful since version 5, there's never been a better time to use these two tools together.

This means that choosing PostgreSQL potentially offers the developer greater extensibility. It's fairly easy to define and add custom data types, operators, functions, and indexing methods. This has far reaching implications for fine-tuning performance; especially when working with particularly complicated data structures.

At the data structure level, PostgreSQL tables and objects can benefit from either static or dynamic inheritance. That is to say, a child object created from its parent can be made to adopt characteristics identical to those of the parent — either just once when it is created, or perpetually (as changes are made to the parent, they are passed on to the children).

This chapter assumes that you have PostgreSQL installed already. If not, please consult the PostGreSQL web site at `http://www.postgresql.org` for more information. The web site contains the most current information on installation and versions for several platforms.

But is it a database yet?

Not quite. Like MySQL, Postgres features some command-line utilities in addition to a unified interactive client for working with PostgreSQL databases. In this section, we'll use some of the command-line utilities stored in the Postgres binary directory to get started. Let's start by inspecting what's already there by using the following command, which you need to run as the Postgres user (or a user with authority), typically user "postgres" on most newly installed systems.

```
psql -l
```

It should return something like the following.

```
        List of databases
    Name    |  Owner  | Encoding
------------+---------+-----------
 template0  | pgsql   | SQL_ASCII
 template1  | pgsql   | SQL_ASCII
(2 rows)
```

Because a database is a mandatory argument to the interactive client and we don't want to work directly on one of our template databases, we'll start by creating a single database that we'll use for the rest of this chapter.

```
createdb sample
```

That's all there is to creating a new database, but of course, this is just a blank slate at this point — a minimalist canvas based on the template database and on which we will paint a structure to our exacting specifications.

Now we can use the interactive client to work with our database:

```
psql sample
```

This will open a copy of `sample` and drop us at a prompt along with some instructions. The prompt inside the interactive client is different from the shell prompt, which can help you distinguish where you are exactly. This is especially useful on a Linux system where multiple shells open at one time is an everyday occurrence. The prompt will take the form `<databasename>=#`, such that our prompt will look like:

```
sample=#
```

If you issue the `\?` command from this prompt, you'll get a very long list of everything you can do from this prompt, exclusive of SQL specifics, which of course are also supported here. We can't offer exhaustive coverage here, but a few key commands are worth exploring:

- The `\h` command lists available help for all of the supported SQL constructs such as `SELECT`, `DELETE`, `GRANT`, and so on.

- The `\d` command, along with one of its accepted parameters, will display information about your database or specific objects in it. Of specific interest are `\dt`, which will list all tables in the current database, and `\d <tablename>`, which will show the structure of the specified table.

- `\H` turns on HTML output, which is handy for exporting data quickly to the web. In conjunction with the `\T` command you can customize the html output somewhat, and you can use `\o` to send it all to a file.

Incidentally, you can call all of these options on the command line when starting `psql`. Simply substitute a hyphen for the slash, and `psql` will execute the commands in sequence and then exit. For even greater utility, you can group these commands in a text file and read them in from the command line.

Down to Real Work

Let's build a simple structure inside our sample database. This example is necessarily abbreviated, but it is designed to give you a quick but useful familiarity with the Postgres and its SQL syntax. If you have already exited from our previous example, get back in to the sample database:

```
psql sample
```

Let's start by defining a simple table to hold the names of some cartoons we really like:

```
CREATE TABLE cartoons(id serial, cartoon varchar(30));
```

The following command allows us to check our result:

```
\d cartoons
```

```
            Table "public.cartoons"
```

```
 Column |         Type         | Modifiers
--------+----------------------+-----------
 id     | integer              |
 cartoon | character varying(30) |
```

Just so we can say we did something relational, we'll create a second table to hold the names of some of the characters in these cartoons.

```
CREATE TABLE characters(id int4, character varchar(15));
```

Postgres offers an astonishing 47 data types, so obviously our example barely even scratches the surface. To be clear, some of these types are really other types with increasing data specificity built in. (You may already be familiar with this concept, commonly called an *input mask* in other database tools.)

Now we'll use some SQL to insert a record into the cartoons table. Note that because the serial type is just an integer with an auto-increment flag attached, we do not need to specify anything for ID:

```
INSERT INTO cartoons (cartoon) values('Adventures of Scoobiro');
```

The absence of an error message suggests that our efforts have met with success, but it's a good idea to check this with a SELECT statement anyway, first to be sure and second because, after a number of records have been entered, we may mentally lose track of where the auto-increment value stands. We'll need the number to define a relationship with our characters table.

```
SELECT * from cartoons;
 id |  cartoon
----+------------
  1 | Scooby
(1 row)
```

Of course, we'd be in pretty bad shape if we managed to mess that one up. Let's also put in a character or two into the characters table.

```
INSERT INTO characters(id, character) VALUES(1, 'Shaggy');
INSERT INTO characters(id, character) VALUES(2, 'Daphne');
```

So we've built a database, created a couple of tables, and added a couple of records, just to make sure things are going well. Before we make a web application out of this, however, we need to create a minimally privileged user — one who can access the tables in our sample database, write to them, and read from them, but not modify them or investigate any other aspect of the system.

In MySQL, users that did not already exist were implicitly created by the GRANT command. While Postgres also uses the GRANT syntax, we must explicitly create a user first:

```
CREATE USER cartoonfan PASSWORD 'secretword';
```

After creating the user, we must give cartoonfan some privileges:

```
GRANT SELECT, INSERT, UPDATE, DELETE on cartoons to cartoonfan;
```

```
GRANT SELECT, INSERT, UPDATE, DELETE
on characters to cartoonfan;
GRANT SELECT, INSERT, UPDATE, DELETE
on cartoons_id_seq to cartoonfan;
```

Note that we must issue this command at the table level. For security reasons, the wildcard character does not function in this context. The last command given is necessary for the `serial` field type we selected in the cartoons table.

PHP and PostgreSQL

Table 33-1 itemizes the PHP functions for working with PostgreSQL databases. There are many more functions than we can possibly elaborate on here. Many of them will make sense only after you have gained more familiarity with Postgres. Also, many of the function names for Postgres changed with PHP version 4.2. Because this is a PHP6 book, we're going to concentrate on the new names rather than the old.

TABLE 33-1

Common PostgreSQL Functions in PHP

Function	Behavior
`pg_connect()` and `pg_pconnect()`	Takes a single connection string as an argument, enumerating connection parameters such as host, database, port, user and password. `pg_pconnect()` is the persistent version. Returns a connection resource. See the listings below for usage examples.
`pg_query()`	This is the standard pass-through mechanism for sending basic SQL to the server. In earlier versions, it was called `pg_exec()`, but this name has been deprecated. Although optional, `pg_query()` likes to see a connection resource, followed by a comma before the actual SQL.
`pg_fetch_row()` `pg_fetch_assoc()` `pg_fetch_array()` `pg_fetch_object()` `pg_fetch_result()` `pg_fetch_all()`	Each of these functions takes at least a query-result resource as an argument and returns varying results depending on the function chosen and how it is called. Each of these except `pg_fetch_all()` used to require a counter argument if you wished to iterate through the returned rows. This argument is still available but is not necessary. These functions differ primarily in the results they return, which in the same order as they are listed are: (1) a numerically indexed array starting at an offset of 0, (2) an associative array with field names as indices, (3) returns both a numeric and an associative array, (4) returns the rows and values in object notation, (5) returns a specific row and column offset, and (6) returns a multidimensional array of the entire result set.
`pg_affected_rows()`	Returns the number of tuples (rows) affected by an `INSERT`, `UPDATE`, or `DELETE` query.

Function	Behavior
pg_free_result()	Frees the memory used by a query result.
pg_num_fields()	Returns the number of fields in a query result. Use with SELECT statements.
pg_num_rows()	Returns the number of rows in a query result. Use with SELECT statements.
pg_close()	Closes the PostgreSQL Connection. Takes a connection resource as an argument.

The Cartoons Database

We'd like to be able to add some cartoons and characters to our database using a handy web-entry system. We've deliberately oversimplified our example to get through the key concepts quickly, so it should be fairly easy to put together a system that allows us to achieve this task. Listing 33-1 shows a welcoming page to our Cartoons database.

LISTING 33-1

index.php

```
<html>
<head>
<title>Cartoons Database</title>
</head>

<body>

<h1>Cartoons and Characters Database</h1>

<p>Welcome to the cartoons and characters database. Existing
entries are provided below. Use the provided functions to get
more details and to edit, add or delete entries.</p>

<?php
$connect_parameters = "host=localhost dbname=sample
user=cartoonfan password=secretword";
if ($link = pg_connect($connect_parameters)) {
  $sSql = "select * from cartoons";
  $sResult = pg_query($link, $sSql);
  if (pg_num_rows($sResult) > 0) {
    print("<table border=\"1\">");
    print("<tr><th>ID</th><th>Cartoon</th>
          <th>Characters</th></tr>");
    while ($sRow = pg_fetch_object($sResult)) {
      print("<tr><th>$sRow->id</th>
```

```
              <td>$sRow->cartoon</td>");
       $tSql = "select * from characters where
                 id = '$sRow->id'";
       $tResult = pg_query($tSql);
       print("<td>");
       while ($tRow = pg_fetch_object($tResult)) {
         print("$tRow->character ");
       }
       print("</td></tr>");
     }
     print("</table>");
   } else {
     print("<p>There are not currently any records in the
            cartoon database.</p>");
   }
} else {
   print("<p>Connection to the cartoons database has
failed</p>");
}
?>

</body>
</html>
```

Notice how different our connect parameters are from what we'd use with MySQL. They aren't even comma separated! This function is even insensitive to the order in which they are supplied — you simply have to use the appropriate parameter name and supply the corresponding value. Additional recognized parameters for this function are `options`, `tty`, and `port`.

The rest of this script reads much like a similar script would for MySQL. A query is defined, we get the results of that as an object, issue a second query to get the character data, and use the ubiquitous print statement to put it all into a nice HTML table.

This is fine as far as it goes, but we don't yet have a way to insert, edit, and delete records. In Listing 33-2, we create a new form for the purpose of inserting records.

LISTING 33-2

insert.php

```
<html>
<head>
<title>Cartoons Database</title>
</head>

<body>
```

```php
<h1>Cartoons and Characters Database</h1>
<?php
if ($_POST['action'] == "Insert") {
  $connect_parameters = "host=localhost dbname=sample
  user=cartoonfan password=secretword";
  $link = pg_connect($connect_parameters);
  $escCartoon = pg_escape_string($_POST['cartoon']);
  $iSql = "insert into cartoons(cartoon)
          values('{$escCartoon}')";
  if (pg_query($link, $iSql)) {
    $jSql = "select currval('cartoons_id_seq') as oid";
    $jResult = pg_query($jSql);
    $j_id = pg_fetch_result($jResult, 0, 'oid');
    $characters_array = explode( "\n", $_POST['characters']);
    for($i=0;$i<count($characters_array);$i++) {
      $char = trim($characters_array[$i]);
      $escChar = pg_escape_string($char);
      $cSql = "insert into characters(id, character)
              values($j_id, '{$escChar}')";
      pg_query($cSql);
    }
    print("<p>Your submission was successfully inserted.
          You can submit another, if you wish</p>");
  } else {
    print("<p>We were unable to insert the records as
          submitted. You can try again, if you wish</p>");
  }
} else {
print("<p>Welcome to the cartoons and characters database.
Enter the");
print("name of your favorite cartoon below, and choose
submit.</p>");
}

?>

<form action="insert.php" method="post">
<p>Enter the name of a favorite cartoon<br>
<input type="text" name="cartoon"></p>
<p>Enter the name of some characters from the cartoon.
(You can enter more later). Use a hard return to
separate each name.<br>
<textarea cols="15" rows="8" name="characters">
</textarea></p>
<input type="submit" name="action" value="Insert">
</form>
<p><a href="index.php">Return to the main page.</a></p>
</body>
</html>
```

This script is doing a lot, so let's review it. Note first the separation of the PHP and HTML elements. All of the conditional code appears at the top, and the conditional display requirements are set up such that we don't have to weave in and out of PHP to get the job done. We're going to display a form even if a submission has just been made so that the user can submit entries one right after the other without an intermediate step. The first conditional, at the top of the page, checks to see if the page is being called from the form. If not, we print a simple instruction set. If the page is the result of a form submission, the fun begins.

We start with a connection to our database. This item is not tested as we just connected on our index page, so we will, perhaps perilously, assume a valid connection. The next step is to generate the SQL for the cartoon table from the POST data and insert a new record into the database. The next bit is designed to get the insert ID of the inserted cartoon. Actually, PHP offers pg_last_oid() for this purpose, but as of this writing, it did not yield the desired result and is currently deemed unreliable.

Once we have the correct ID, we do some manipulation on the contents of the textarea that contains our characters. Refer to Chapter 21 on arrays for more on explode(). Basically, we're splitting the field each place we find a line break and popping the resulting elements into an array. Next we call trim() to get rid of the superfluous space character left behind by explode(). This is done because Windows systems submit an additional \r wherever a line break occurs. If we explode on this as a separator, explode will fail when the form is submitted from Linux, for example, because the character is not there.

Finally, we iterate through the resulting array, putting each character into the characters table with its own insert query, and we return a message of success or failure.

When we went back to test our results, we found a few problems with index.php and so we present a revised version in Listing 33-3. We could have just changed the original listing, but this so nicely illustrates the debugging process that we've included it this way to point out the improvements.

LISTING 33-3

index.php (improved)

```
<html>
<head>
<title>Cartoons Database</title>
</head>

<body>

<h1>Cartoons and Characters Database</h1>
<?php
if ($_POST['action'] == "Insert") {
  $connect_parameters = "host=localhost dbname=sample
user=cartoonfan password=secretword";
  $link = pg_connect($connect_parameters);
```

```php
    $escCartoon = pg_escape_string($_POST['cartoon']);
    $iSql = "insert into cartoons(cartoon)
              values('{$escCartoon}')";
    if (pg_query($link, $iSql)) {
      $jSql = "select currval('cartoons_id_seq') as oid";
      $jResult = pg_query($jSql);
      $j_id = pg_fetch_result($jResult, 0, 'oid');
      $characters_array = explode( "\n", $_POST['characters']);
      for($i=0;$i<count($characters_array);$i++) {
        $char = trim($characters_array[$i]);
        $escChar = pg_escape_string($char);
        $cSql = "insert into characters(id, character)
                  values($j_id, '{$escChar}')";
        pg_query($cSql);
      }
      print("<p>Your submission was successfully inserted.
      You can submit another, if you wish</p>");
    } else {
      print("<p>We were unable to insert the records as submitted.
              You can try again, if you wish</p>");
    }
} else {
  print("<p>Welcome to the cartoons and characters database.
Enter the");
  print("name of your favorite cartoon below, and choose
submit.</p>");
}

?>

<form action="insert.php" method="post">
<p>Enter the name of a favorite cartoon<br>
<input type="text" name="cartoon"></p>
<p>Enter the name of some characters from the cartoon.
(You can enter more later). Use a hard return to
separate each name.<br>
<textarea cols="15" rows="8" name="characters">
</textarea></p>
<input type="submit" name="action" value="Insert">
</form>
<p><a href="index.php">Return to the main page.</a></p>
</body>
</html>
```

First, we added a link to our now finished insert form. Second, we are now passing the action parameter via our submit button, purely for the sake of orderliness. We also encountered a problem with our characters display. When we submitted a character with a space in its name (okay, we admit, we were trying to submit *Wonder Woman*), it was unclear where one character ends and the

561

next one begins. Using a little regex and string concatenation, we've now caused this to display as a comma-separated list.

Now we need a form for editing records. Listing 33-4 is what we've come up with.

LISTING 33-4

edit.php

```php
<html>
<head>
<title>Cartoons Database</title>
</head>
<body>
<h1>Cartoons and Characters Database</h1>
<?php
$connect_parameters = "host=localhost dbname=sample
user=cartoonfan password=secretword";
$link = pg_connect($connect_parameters);
if ($_POST['action'] == "Update") {
$escCartoon = pg_escape_string($_POST['cartoon']);
$escId = pg_escape_string($_POST['f']);
$sSql = "update cartoons set cartoon = '{$escCartoon}'
        where id = '{$escId}";
  if (pg_query($sSql)) {
    $dSql = "delete from characters where id = '{$escId}'";
    pg_query($dSql);
    $characters_array = explode( "\n", $_POST['characters']);
    for($i=0;$i<count($characters_array);$i++) {
      $char = trim($characters_array[$i]);
      $escChar = pg_escape_string($char);
      if($char <> '') {
        $cSql = "insert into characters
                (id, character)
                values({$escId}, '{$escChar}')";
        pg_query($cSql);
      }
    }
    print("<p>Your edits were successfully posted.</p>");
  } else {
    print("<p>Update of record $_POST[f] failed.</p>");
  }
  print("<p><a href=\"index.php\">Return to the main
        page.</a></p>");
} else {
  $escId = pg_escape_string($_GET['f']);
  $sSql = "select * from cartoons where id = {$escId}";
  $sResult = pg_query($sSql);
```

```
  $sRow = pg_fetch_object($sResult);
  print("<form action=\"edit.php\" method=\"post\">");
  print("<p>Edit the name of a favorite cartoon<br>");
  print("<input type=\"hidden\" name=\"f\"
         value=\"$_GET[f]\">");
  print("<input type=\"text\" name=\"cartoon\"
         value=\"$sRow->cartoon\"></p>");
  print("<p>Edit the name of some characters from
         the cartoon. ");
  print("(You can enter more later). Use a hard return to ");
  print("separate each name.<br>");
  print("<textarea cols=\"15\" rows=\"8\"
         name=\"characters\">");
  $escId = pg_escape_string($_GET['f']);
  $cSql = "select * from characters where id = {$escId}";
  $cResult = pg_query($cSql);
  while ($cRow = pg_fetch_object($cResult)) {
    print("$cRow->character\r\n");
  }
  print("</textarea></p>");
  print("<input type=\"submit\" name=\"action\"
         value=\"Update\">");
  print("</form>");
 print("<p><a href=\"index.php\">Return to the main
         page.</a></p>");
}
?>
</body>
</html>
```

Like `insert.php`, `edit.php` is a recursive action form. The form post is sent to the same script and the action taken depends, essentially, on the contents of a hidden variable, `$action`. In the absence of this variable, we just retrieve the records from the cartoons and character tables and drop them back in the form much as they appeared in the original insert form. To parse the characters back into their original positions, note that we have essentially reversed the explode function we used in `insert.php`, but we'll shift it back into forward when we go to post the updates.

One thing you might find strange is that before we post the updates to the characters table, we delete all the entries. This is because we haven't created a unique key for each character, so there isn't a way to conveniently refer to an individual record unambiguously. Yes, this is a design flaw. In a bigger project, it would be a substantial design flaw. But there are some advantages to the way we've done this. A form to update both characters and cartoons in the same action is much easier to do in this scenario. It would be quite a bit more code intensive, though certainly feasible, to add a serial ID to the characters table, retrieve all three fields, set up a multidimensional array in our form, retrieve it, and process it in our script such that multiple records are updated in a single operation. You get the idea. Sometimes it's okay to opt for simplicity.

Finally, we've made a few changes to index.php, which we're going to show you in Listing 33-5. All we've done is added a way to get at the edit functions and a simple routine for deleting a record, parent and children, in one operation.

LISTING 33-5

index.php (final)

```
<html>
<head>
<title>Cartoons Database</title>
</head>

<body>

<h1>Cartoons and Characters Database</h1>

<p>Welcome to the cartoons and characters database. Existing
entries are provided below. Use the provided functions to get
more details and to edit, add or delete entries.</p>

<?php
$connect_parameters = "host=localhost dbname=sample
user=cartoonfan password=secretword";
if ($link = pg_connect($connect_parameters)) {
  if($_GET['action'] == "d") {
    $escId = pg_escape_string($_GET['f']);
    $dSql = "delete from characters where id = '{$escId}'";
    pg_query($dSql);
    $dSql = "delete from cartoons where id = '{$escId}'";
    pg_query($dSql);
  }
  $sSql = "select * from cartoons";
  $sResult = pg_query($link, $sSql);
  if (pg_num_rows($sResult) > 0) {
    print("<table border=\"1\">");
    print("<tr><th>ID</th><th>Cartoon</th>
          <th>Characters</th><th></th></tr>");
    while ($sRow = pg_fetch_object($sResult)) {
      print("<tr><th>$sRow->id</th>
            <td>$sRow->cartoon</td>");
      $tSql = "select * from characters where id = '$sRow->id'";
      $tResult = pg_query($tSql);
      print("<td>");
      $character_string = "";
      while ($tRow = pg_fetch_object($tResult)) {
        $character_string .= "$tRow->character, ";
      }
```

```
     $new_character_string = ereg_replace("(, )$", "",
                                         $character_string);
     print("$new_character_string</td>");
     print("<td><a href=\"edit.php?f=$sRow->id\">Edit</a> |");
     print(" <a href=\"index.php?f=$sRow->id&action=d\">
                       Delete</a></td></tr>");
   }
   print("</table>");
  } else {
  print("<p>There are not currently any records in the
        cartoon database.</p>");
  }
  print("<p><a href=\"insert.php\">Add a Record</a></p>");
} else {
  print("<p>Connection to the cartoons database has
failed</p>");
}
?>

</body>
</html>
```

Summary

PostgreSQL is an interesting and powerful database tool. Although we did not comprehensively cover all of its utility here, we have shown you enough basics to get started with it. Check the PHP documentation at `www.php.net/pgsql` for a comprehensive listing of PostgreSQL functions.

Chapter 34

Using PEAR DB with PHP

PEAR DB is one of several wrappers around PHP's database extensions that seek to generalize the concept of a database connection. PEAR DB is fully object oriented. With PEAR DB implemented in your software, you'll find it easier to allow your application's users to select from several databases (the list of supported databases includes all the majors, and plenty of second-string players as well). You may also find it easier than before to handle errors and react to unexpected occurrences that take place during connections and queries.

This chapter aims to explain what PEAR DB is all about. It's not hard, but it is important that you understand PEAR DB in theory and practice, because there are substantial pros and cons to its use. We start with a general discussion of the benefits and costs of a database abstraction layer; then we explain the basics of using the PEAR DB package.

For reference, examples, and further explanations of concepts, have a look at the PEAR DB web site at `http://pear.php.net/package/DB`.

CROSS-REF For general information on PEAR and instructions on how to locate and install PEAR libraries, see Appendix E.

IN THIS CHAPTER

Pear DB concepts

Pear DB functions

567

Pear DB Concepts

There are a number of concepts in PEAR DB that you must understand to work with the class effectively. These include:

- Data Source Names (DSNs)
- Connection
- Query
- Row retrieval
- Disconnection

The following subsections explain each of these items in turn.

Data Source Names (DSNs)

A Data Source Name (DSN) is simply a text string that describes where a database is and how to access it. A DSN is very much like a URL for a database. Because databases almost always have user-level access control, DSNs specify usernames and passwords. Note that DSNs specify the *database* to which your program is connecting, not the *table* within the database.

In order to form a DSN, you have to come up with a number of values:

- The type of database you're connecting to (FrontBase, MySQL, Oracle, or whatever you fancy)
- The hostname or IP address of the machine running the database server
- The database name
- The username you want the software to use
- The corresponding password

Very often, you'll need to access these values in order to create DSNs in several different but related PHP programs. For that reason, it often makes sense to assign the literal values to variables in a special program that's imported elsewhere. Such a program might look like this (call it `dbSpecs.php`):

```php
<?php
$phptype = 'mysql';
$dbHost = 'spock';
$database = 'inventory';
$username = 'phpUser';
$password = 'sesame';
?>
```

The only odd bit there is the value for $phptype. Its value is a standard string that corresponds to a specific database (MySQL, in this example). The related sidebar shows all valid options for the database identifier.

Valid Database Identifiers for DSNs

Here's a list of all the database management servers supported by PEAR DB, complete with the strings you should use to identify them in DSNs.

- FrontBase (`fbsql`)
- InterBase (`ibase`)
- Informix (`ifx`)
- Mini SQL (`msql`)
- Microsoft SQL Server (`mssql`)
- MySQL (`mysql`)
- Oracle 7/8/8i (`oci8`)
- ODBC (`odbc`)
- PostgreSQL (`pgsql`)
- SyBase (`Sybase`)

Having defined your variables in a simple library, when you want to create a DSN, you can import `dbSpecs.php` and have access to its values. The advantage is that if the values in `dbSpecs.php` ever change, you need to modify them in just one place. The import statement is simple:

```
require_once('dbSpecs.php');
```

With that imported, you can use the values defined in `dbSpecs.php` to create a DSN. Remember, though, that you need to register the variables if you create your DSN inside a function (this is a standard characteristic of PHP variable scoping):

```
global $phptype;
global $hostspec;
global $database;
global $username;
global $password;
```

Once you have the values available, stringing them together into a properly formatted DSN is pretty simple. The standard format looks like this:

```
$dsn = "$phptype://$username:$password@$hostspec/$database";
```

569

There are more obscure options available for use in defining DSNs. They're described in detail at
`http://pear.php.net/manual/en/package.database.db.intro-dsn.php`.

TIP   As always, we recommend storing database connection variables outside the web tree for greater security.

Connection

When you have a valid DSN, it's a simple matter to tell PEAR DB to establish a connection to the
database the DSN describes. The line of code you need looks like this:

```
$db = DB::connect($dsn);
```

If the connection attempt succeeds, everything's great — $db contains an object representing a con-
nection to the database, and you can run queries against that object (among other amusing and educa-
tional activities). Because you're a good programmer, though, you should allow for errors. Here's how:

```
if (DB::isError($db)) {

  die($db->getMessage());
}
```

The isError() method returns true if the database object — $db in this case — represents a con-
nectivity error, and false if not (meaning, if the connection succeeded). In this case, the code aborts
if the database connection wasn't successfully established.

Query

Building on the successful database connection, you'll typically want to run a query against the
database. That's accomplished with Structured Query Language (SQL), which is described in detail
in Chapter 13.

The customary procedure is to write your SQL query as a string and stuff that string into a variable
called $sql or $query:

```
$sql = "SELECT cityName FROM cities ORDER BY cityName";
```

Then use that variable as a parameter for the query() method of your database object (in other words,
of the variable that represents the connection to the database), and assign the results to another variable:

```
$result = $db->query($sql);
```

Once again, allow for the possibility of error:

```
if (DB::isError($result))
  $errorMessage = $result->getMessage();
  die($errorMessage);
}
```

Row retrieval

As a result of the `query()` operation, $result contains a result set. A result set is some number of rows (possibly zero).

To extract values from the result set, use a `while` loop like this one:

```
$i=0;
while ($row = $result->fetchRow()) {
  $returnArray[$i] = $row[0];
  ++$i;
}
```

This loop exists to use the `fetchRow()` function against every row in the result object, thus extracting it. Because we know the result set has only one column (the SQL statement requested only `cityName`), we can take the first element of every row array ($`row[0]` — the only element) and put it into another array, $`returnArray`. Presumably, we'll do something useful with $`returnArray` later, but that's not anything to do with PEAR DB directly.

Disconnection

When a database connection has served its purpose, use the `disconnect()` method of the database object to free the resources associated with the connection:

```
$db->disconnect();
```

It's important to remember to do this; otherwise, the connection remains active until it times out (a long time). This means that the database server, in an active environment, could become over-whelmed with zombie connections.

A complete example

It may be beneficial to have a look at a complete function that performs database access operations by way of the PEAR DB class.

The role of this function, which we'll call `getItems.php`, is to extract all items from a Microsoft SQL Server database called `INVENTORY_item`. To return all items, we need a `SELECT` statement that draws all columns out of the `INVENTORY _item` table. Because `INVENTORY _item` has no foreign keys (an assumption made for the purposes of this illustration), extracting its data involves only sending a straightforward `SELECT` statement to the database server via a PEAR DB connection.

Let's examine `getItems.php` line by line to see how this is done.

```
<?php
require_once ('DB.php');
require_once ('dbSpecs.php');
```

First, we must import the PEAR DB classes and dbSpecs.php, which contains information about the database server and security credentials for it (it's listed earlier in this chapter).

```
function getItems()
{
  global $phptype;
  global $hostspec;
  global $database;
  global $username;
  global $password;
```

In the function, the five global variables (from dbSpecs.php) must be declared for them to be accessible.

```
$dsn = "$phptype://$username:$password@$hostspec/$database";
$db = DB::connect($dsn);
if (DB::isError($db)) {
  die($db->getMessage());
}
```

Using the PEAR DB procedure discussed earlier in this chapter, the program connects to the database server. The program checks for an error condition and aborts if one is found to exist as a result of the connection attempt.

The program then defines the SQL statement to be run against the database to which a connection has been established:

```
$sql = "select id, desc, weight, packageQty, unit, supplierID,
cost from INVENTORY_item";
```

That's the SQL query that is to be sent to the database server. Note that we specify the columns, even though we want all of them. That way we know what order the columns will be in when results come back.

```
$result = $db->query($sql);
if (DB::isError($result)) {
 $errorMessage = $result->getMessage();
 die($errorMessage);
}
```

The program sends the query to the database and checks to see if an error message comes back.

```
$returnArray = array();
while ($row = $result->fetchRow()) {
  $id = $row[0];
  $desc = $row[1];
  $weight = $row[2];
  $packageQty = $row[3];
  $unit = $row[4];
```

```
      $supplierID = $row[5];
      $cost= $row[6];
      $returnArray[] =
        array('id' => $id, 'desc' => $desc, 'weight' => $weight,
              'packageQty' => $packageQty, 'unit' => $unit,
              'supplierID' => $supplierID, 'cost' => $cost);
  }
  $db->disconnect();
  return $returnArray;
  }
  ?>
```

The remainder of the program involves setting up an array called `$returnArray`. It is filled with a series of subarrays, making it a two-dimensional array. The subarrays are associative arrays; their keys correspond to column names in the database, and their values come from each row of `$result`.

PEAR DB Functions

The PEAR DB class is fairly extensive, with members far more numerous than the widely used ones covered already in this chapter's examples. Most of the other members are specialized, and as such come in handy only under particular circumstances.

This section summarizes some of the most useful members of the PEAR DB class but is not comprehensive. Be sure to refer to `http://pear.php.net/manual/en/package.database.php` for the official list and documentation.

Members of the DB class

The `DB` class itself is the main PEAR DB class, and is used to represent a connection to a database (or an attempt to create one). Methods include:

- `DB::connect()`: Uses a DSN to connect to a database

- `DB::isWarning()`: Returns true if a connection attempt yielded a warning

- `DB::isError()`: Returns true if a connection attempt yielded an error

Members of the DB_Common class

The methods of the `DB_Common` class may be invoked on a database connection for such purposes as executing queries and getting information from the database. Methods include:

- `DB_Common::affectedRows()`: Returns the number of rows affected by a query.

- `DB_Common::disconnect()`: Disconnects a database connection.

- ■ `DB_Common::getAll()`: Returns all rows returned by a query.

- ■ `DB_Common::getAssoc()`: Returns all rows as an associative array.

- ■ `DB_Common::getCol()`: Returns all rows in a specified column.

- ■ `DB_Common::getOne()`: Returns the value in the first column of the first row.

- ■ `DB_Common::getRow()`: Returns the first row.

- ■ `DB_Common::nextId()`: Allows you to exercise extra control over the establishment of unique `id` values, as in a primary key column.

- ■ `DB_Common::query()`: Sends an SQL query string.

Members of the DB_Result class

The members of the `DB_Result` class may be invoked on a result set, which is what exists after a query is run on a database.

- ■ `DB_Result::fetchInto()`: Extracts a row into a specified variable

- ■ `DB_Result::fetchRow()`: Extracts the next row

- ■ `DB_Result::free()`: Destroys the result set

- ■ `DB_Result::numCols()`: Returns the number of columns in a result set

- ■ `DB_Result::numRows()`: Returns the number of rows in a result set

Summary

This chapter introduced PEAR DB, the class with which PHP makes database connectivity more generic. You saw that PEAR DB is designed to make the task of switching from one database server to another extremely easy and that it generally succeeds in this design goal. In many cases, switching connectivity from one database server to another requires the modification of only one word in a whole database-enabled program.

The process of connecting to a database via PEAR DB involves establishing a DSN, which is essentially a URL for a database. The DSN specifies the hostname of the database server, as well as the database name, and the username and password required to gain access to the database. With a valid DSN, it's possible to establish a connection, run a query, and extract rows of values from the query's results before disconnecting. PEAR DB also includes methods — which you should take care to use — for detecting error and warning conditions.

Much more information on Pear DB can be found at its official web site `http://pear.php.net/package/DB`.

Chapter 35

An Overview of Oracle

Oracle databases are extremely powerful, reliable, and fast for certain kinds of queries. They are also a testament to the power of marketing, as they enjoy a mystique unmatched by any other data storage product. In this chapter, we try to cut through the hype and give you a practical foundation in using Oracle with PHP.

TIP   Much of the information in this chapter is also applicable to IBM DB2, which PHP connects to using the ODBC interface, InterBase, and similar products. Much of Oracle's functionality is also duplicated by PostgreSQL, although this open source database unfortunately has a considerably different API than any other database. The point of this chapter is not to push Oracle over another product; it is to introduce a market-leading commercial database to those who want to learn how to use it. Feel free to mentally substitute *DB2* or *PostgreSQL* for every instance of *Oracle* in the following section.

When Do You Need Oracle?

We have never had a prospective employer or client for a PHP job who did not at least mention the possibility of using Oracle. Ironically, the one organization that really needed this type of functionality was very slow to adopt it, whereas everyone else didn't need Oracle at all but wanted to architect current development around the theoretical possibility that they might need it later. In an anecdotal way, this experience testifies to the niche Oracle Corporation has managed to carve out in the minds of the entire software industry.

Through a powerful marketing machine (and, of course, a fine product), Oracle has managed to make its name synonymous with size, scale, and (by implication) success. Therefore, every businessperson who dreams of making it big in any business having to do with data has a fantasy that someday he or she will need and be able to afford an Oracle installation. When the humble PHP developer says, "You don't need Oracle," this is too often taken as equivalent to saying, "You're never going to amount to anything" — in other words, a big splash of cold water right in the face of the entrepreneur or manager. This does not tend to endear the humble PHP developer to the person writing the checks, even if your motivation is to save money and trouble for the boss. Sometimes reality cannot compete with fantasy, even in the supposedly hard-headed world of business.

So how do you realistically decide whether you need Oracle (or a workalike) or not? Certain well-understood factors are tripwires. If you are in one of these situations, your pondering is done. If not, the chances are very high that your needs could be met perfectly well — perhaps even better — by a cheaper, easier-to-use database.

Money

If you keep track of money or anything that can be converted to money (credit card charges, equities, airline miles, royalty payments, vacation time), you need the transactional model. Done, end of story, move on to the next thing. Not only are these things important to track end to end, but people tend to get very annoyed when you fail to do this correctly. You do not want to be matching up failed charge attempts or stock purchases programmatically.

The only exception is if you do not handle the financial part of the transaction yourself. It's fairly common these days for even pretty large web sites to route their fulfillment transactions through some third party. If you are willing to send a customer off to your fulfillment partner, and then take its word on whether the transaction was completed successfully or not, you can offload the transactional database requirements onto your partner.

Other rivalrous resources

If you track other kinds of rivalrous resources on a large scale — airline tickets, concert tickets, inventory — you very likely need a transactional database. Note that the operative term is *large scale*. This feature applies only if you are in imminent danger of having *colliding writes*, meaning that the chances are good that, between the time you check whether a resource is available and the time you write it into another table, someone else may have claimed it. In other words, this stricture applies to computer-scale time. If you run a web site that takes a couple hundred new registrations a day, you are not going to run into this kind of problem except as a rare fluke.

Huge data sets

Again, *huge* is the key. Millions of rows is not huge. If you will never realistically get past tens of millions of rows, you don't need Oracle for sheer scalability.

Lots of big formulaic writes or data munging

Oracle's stored procedures can increase speed immensely in situations where you have the same kind of big write or data processing happening all the time. Stored procedures amount to moving part of the code into the database itself. Processes that use these stored procedures will get done faster than processes that don't, because you've already told the database exactly what to do in a certain situation. Databases without stored procedures require all instructions to come from the PHP program, which is far slower in loops than simply handing off the input data and letting the database run with it.

For instance, if you run an e-commerce site that just takes inputs from one kind of order form and nothing else, you'll get increased overall speed from a stored procedure. The database will not have to finish one step and then wait for the program to tell it the next step or to feed it the next bit of data. PHP will simply open up the connection, shovel some data down the pipe as fast as it can, kick off the stored procedure, and possibly wait for a response.

Stored procedures obviously add value in cases where there is massive data munging happening on the database side. For instance, if you have a big data warehouse that takes a described data set and performs lots of analytical operations on it, then returns the results of those operations, you will very likely get much improved performance from a stored procedure.

Remember that stored procedures don't help you in cases where there is variability in the process. For instance, if you make ad hoc queries on a data warehouse all the time, those queries will be extremely slow compared to a query that you make every day using a stored procedure. Also, you have to use stored procedures pretty often to choose a database on this basis. A whole infrastructure needs to be in place for you to derive benefits from stored procedures, and you need to decide whether it's worth the very real costs to build this infrastructure.

Triggers

Triggers are just what they sound like: The database keeps track of state; when some kind of triggering event happens, it kicks off one or more stored procedures. You need triggers if you need to respond to certain well-defined data events in real time and are willing to give up global performance to do it. If your changes do not have to occur instantly — if, for instance, they can happen once an hour — you are probably better off without triggers as you can schedule stored procedures to run periodically. On the other hand, if you have to detect changes every second or so, triggers will be much more efficient than cronjobs or other alternatives. This is a pretty rare feature for a web architecture, so think hard about whether you need triggers.

Legal liability

Finally, there are nontechnical reasons why you might need Oracle — one of which is legal liability. Imagine that you run a business that is in some kind of personal information space — credit reports or medical records, perhaps. If somehow two records get mixed up, you would face significant legal liability for this most private type of information. One line of defense is to assert that you followed industry standard practices to safeguard data integrity, including use of the industry-leading database system. The tens of thousands of dollars you spend on an Oracle installation may seem cheap in this circumstance.

Bottom line: two-year outlook

Even if you don't need any of these features now, it behooves you to think whether you might need them sometime in the foreseeable future. But don't try to build a web architecture for the ages — new things happen so fast on the Internet that cathedral-builders are made fools of daily.

In general it seems like a two-year window is about right unless your needs are changing exceptionally quickly or slowly. If you don't see any of the preceding conditions looming on your horizon within two years or so, don't worry too much about switching to Oracle.

On the other hand, if you meet any of the preceding criteria, don't waste time — you need Oracle or one of its functional competitors.

Oracle and Web Architecture

Having an Oracle database on your backend implies certain things about your architecture and your web development team. Do not think you can finesse these issues. If you're not ready to accept them, you're not ready for Oracle.

Specialized team members

The minute you have an Oracle installation, you need at a minimum a DBA and a PL/SQL programmer. Do not think that it's at all reasonable to expect web developers or binary programmers to take over any of the tasks of these positions, as they often are expected to do with simpler data stores. Database tasks such as installation, tuning, refreshing, setting indexes, maintaining hot backups, and writing complex stored procedures are very highly specialized skills, which require specialized training and probably certification.

One of us once worked in a situation where (mercifully briefly) a team of programmers tried to get a database installed and running with help from a couple of contractors rather than full-time Oracle professionals. This was an indescribably miserable experience for everyone concerned, with data corruption and bizarre server issues on a daily basis. The minute a real DBA and PL/SQL programmer came in, the situation improved immeasurably — 10-second queries became half-second queries, and so on. There's a reason that these professionals make the big bucks, and we can assure you that you don't want to learn why the hard way.

Shared development databases

Due to the cost and complexity of the database, you will almost certainly need to limit the number of separate development instances in the office. This means that if all your developers are working against one development database that gets corrupted or otherwise goes down, the whole team will be out of commission until the problem can be resolved. On the other hand, MySQL and SQL Server make it easy for every developer to run a local instance if necessary.

And don't even *think* about cheating on your licenses — an Oracle software license audit is nothing you want to experience.

Limited schema changes

Lighter-weight databases make it possible to make schema changes on the fly — even in production. In Oracle, this is not a good idea. It's more difficult to add new database-driven features quickly because of the time and planning necessary to successfully create schema changes. With Oracle, you really need to design the entire schema in advance (as you're supposed to do but never quite take the time to complete).

Tools (or lack thereof)

Be prepared to invest significantly in tools. Almost no tools are available on the market to help you with Oracle from a PHP platform specifically. More generally, Oracle is not interested in data sets smaller than a few tens of thousands of records, and many of the tools available are not really worth running on data sets smaller than that.

Replication and failover

Some database servers, like MySQL, work best in clusters with one master handling writes while a bunch of slaves serve up reads. This means that, for maximum efficiency, MySQL code should be written to take advantage of multiple connections on a single page — each MySQL function specifying a particular resource handle, for instance.

Oracle, however, works best in a single monolithic instance. Remember that eBay grew to be one of the biggest sites on the web, and suffered many costly site failures, before even beginning to move away from a monolithic database instance. This tends to create a single point of failure model that affects all aspects of deployment — after all, what's the point in having massive redundancy in web servers when all of them could be struck down by a single hardware problem on the database server? Most important for PHP developers, writing code for Oracle will not require you to juggle multiple connections on a per-page basis, as there is little to be gained by this practice.

Data caching

One piece of good news is that Oracle can be configured to cache data. This means that if you make the same query over and over for a while — for instance, you validate the user against the database on every page load — the database will only actually perform the query the first time. On subsequent requests, it will just serve up the result from some cache memory, for a much faster result. This means that you will not have to implement a custom data-caching scheme of your own, which is a very expensive and tricky task.

However, we should mention that lighter-weight databases are rapidly catching up to Oracle's data caching capabilities for small queries. Microsoft SQL Server is reported to have excellent data caching, and MySQL implemented it for the first time in version 4.0.1.

Using OCI8 Functions

CAUTION   The remainder of this chapter explains how to use PHP's Oracle functions. If you have no familiarity with Oracle at all, and especially if you do not have extensive experience with other SQL databases, you will probably also need to consult an Oracle reference such as *Oracle 11g For Dummies* (Zeis and Ruel, Wiley, 2009).

PHP has two Oracle extensions: Oracle and OCI8. The Oracle extension is deprecated for versions of Oracle after 7 and should be avoided if at all possible. The OCI8 extension is literally dozens to hundreds of times faster at most queries, and also allows much better handling of cursors. We will not be describing the Oracle extension at all; use OCI8 from the beginning.

TIP   Although Oracle is currently in version 11, the OCI8 extension is still the one you want. Apparently the PHP team has simply decided not to change the name with every update.

If you walk through a couple of typical Oracle queries step by step, you can see that the procedure is a bit different from that for MySQL or SQL Server:

```
$name = str_replace("'", "''", $name);
$query = "SELECT product_id FROM product
        WHERE product_name = '$name'";
$stmt = OCIParse($conn, $query);
OCIExecute($stmt, OCI_DEFAULT);
$err_array = OCIError($stmt);
if ($err_array) {
   $err_message = $err_array['message'];
   $$error_str = $err_message;
   OCIFreeStatement($stmt);
} else {
   OCIFetchStatement($stmt, $res, OCI_RETURN_NULLS);
   OCIFreeStatement($stmt);
}
$product_id = $res['PRODUCT_ID'][0];
```

Right away you will notice differences such as string escaping, parsing and executing, memory management, and fetching data sets. We explain these differences in more detail in the sections that follow. Please refer to the preceding code block for all references that do not have their own code examples.

Escaping strings

Remember that Oracle uses Sybase-style string escaping — in other words, it escapes single quotation marks with a single quotation mark, not with a backslash. This means that you will have to manually escape every string, since `magic_quotes` has the wrong effect, or you have to set Sybase-style magic quotes in your `php.ini` file.

Parsing and executing

In most other databases, you send a SQL query over the pipe, and it returns either a value or an error message. In Oracle, you have an intermediate step in which you must parse the query for correctness. If your SQL is bad, the query will never be sent to the database. An invalid SQL query will return an Oracle error code via `OCIError()`.

In most cases, the query will be fine and you'll move right to executing the query. This is your opportunity to specify a mode for your query, either `OCI_COMMIT_ON_SUCCESS` or `OCI_DEFAULT`, which does not auto-commit. We generally do not auto-commit because we like to know when a rollback is necessary, but this is largely a matter of preference. Of course, for `SELECT` statements it doesn't matter which mode you choose because no commitment or rollback is happening anyway. See the section on transactionality later in this chapter.

Error reporting

Oracle error reporting is also unique. For one thing, you can specify whether you're asking for a *global*-, *connection*-, or *statement*-level error. A global error is a failure to get a connection. A connection error is an invalid SQL statement that chokes at the parse stage — remember, it isn't a statement yet, so it can't be a statement-level error. A statement error involves a problem with a properly parsed statement.

The product of `OCIError()` is an associative array, where `code` is the error code and `message` is a text string.

Memory management

In Oracle, you are expected to do some memory management manually. In particular, it's a good idea to free statement memory when you're done with the statement handle, and cursor memory when you're done with the cursors. Theoretically, all the memory will be reclaimed at the end of every script, but if your scripts are big you might want to free memory as you go.

Ask for nulls

If you want null values in your data set, you have to ask for them in the fetch function. Otherwise, Oracle will not return the field name as part of its associative array, and thus an array member you expect to see will not exist. Many other databases will automatically return the null values.

Fetching entire data sets

Oracle has four fetching functions: `OCIResult`, `OCIFetch`, `OCIFetchInto`, and `OCIFetchStatement`. The first three correspond fairly straightforwardly to the single-column-, row-, and array-fetching functions enjoyed by all other PHP database extensions — but the last is unique to Oracle. It fetches the entire result set into one big array. This can mean less looping and faster access to your data. Of course,

you should make sure that you are returning a reasonably small result set; otherwise, you can bring PHP to its knees.

All caps

Oracle column names, which become associative array indices, are in all capital letters. You can ask for the field in lowercase, but by the time it comes back from the database it will be in all caps, as shown in the following code:

```
$query = "SELECT product_name, modified, created
        FROM product
         WHERE product_id = 1";
$stmt = OCIParse($conn, $query);
OCIExecute($stmt, OCI_DEFAULT);
$err_array = OCIError($stmt);
if ($err_array) {
   $err_message = $err_array['message'];
   $error_str = $err_message;
   OCIFreeStatement($stmt);

} else {
   OCIFetchStatement($stmt, $res)
   OCIFreeStatement($stmt);
   $product_name = $res['PRODUCT_NAME'][0];
   $modified = $res['MODIFIED'][0];
   $created = $res['CREATED'][0];
}
```

Transactionality

Oracle's famous transactionality implies that all INSERT, UPDATE, and DELETE statements must be committed or rolled back before they will really be stored in the database. You can commit automatically during the OCIExecute() step, but if you really want an entire string of queries to succeed or fail together — the essence of transactionality — it's better to commit by hand when all of them are complete.

```
$query = "DELETE FROM product
        WHERE product_name = '$product_name'";
$stmt = OCIParse($conn, $query);
OCIExecute($stmt, OCI_DEFAULT);
$err_array = OCIError($stmt);
if ($err_array) {
   $err_message = $err_array['message'];
   $error_str = $err_message;
   OCIFreeStatement($stmt);
   OCIRollback($conn);
} else {
   OCIFreeStatement($stmt);
}
```

```
$query = "INSERT INTO product (product_name, modified, created)
         VALUES ('$product_name', SYSDATE, SYSDATE)";
$stmt = OCIParse($conn, $query);
OCIExecute($stmt, OCI_DEFAULT);
$err_array = OCIError($stmt);
if ($err_array) {
   $err_message = $err_array['message'];
   $error_str = $err_message;
   OCIFreeStatement($stmt);
   OCIRollback($conn);
   exit;
} else {
   OCIFreeStatement($stmt);
}
OCICommit($conn);
```

The net effect of this code block will be to ensure that there is only, at most, one row in the database for each product name. If either the DELETE or the INSERT fails, the state of the database will not be changed. You should never have a situation where you have zero or two rows with the same product name. Note the second argument to OCIExecute() is OCI_DEFAULT, which means that statements will not be auto-committed and must, therefore, be committed by hand.

CAUTION It's very important to remember that commits and rollbacks occur on a *connection*, not on a statement. Everything since the last commit or rollback will be entered into the database when you call OCICommit() or OCIRollback(). You may call it as many times during a script as you like, but generally all the parts of a transaction are committed or rolled back together.

Stored procedures and cursors

Stored procedures are programs that execute on the database. They move some of the programming into the database layer rather than the PHP layer. PHP merely sends data to the function and handles any returned values.

Because you are connecting to a particular compiled program on the database server, you need to establish a more specific kind of connection to the stored procedure. This connection is called a *cursor*, and the process of designating variables for use by the stored procedure is called *binding to a cursor*. Basically you are taking a PHP variable and transforming it into a variable on the Oracle side, or creating a PHP variable in which to store data returned from Oracle. Cursors must be executed separately from ordinary Oracle statements.

The code block that follows shows a simple example of stored procedure being called from PHP. In this case, we are executing the stored procedure named get_categories(), which takes no inputs and returns one output, OUT1, which we are binding to the PHP variable name $cursor1.

```
// Call stored procedure get_categories
$request = "begin DEV.get_categories(:OUT1); end;";
$cursor1 = OCINewCursor($conn);
$stmt = OCIParse($conn, $request);
OCIBindByName($stmt, ":OUT1", &$cursor1, -1, OCI_B_CURSOR);
```

583

```
OCIExecute($stmt);
OCIExecute($cursor1);
$err_array = OCIError($conn);
if ($err_array) {
    $err_message = $err_array['message'];
    echo $err_message;
    OCIFreeCursor($cursor1);
    OCIFreeStatement($stmt);
    OCILogoff($conn);
    exit;
}
while (OCIFetchInto($cursor1, &$cat_array)) {
    $opt_str .= "<OPTION VALUE=\""
                .$cat_array[0]
                ."\">"
                .$cat_array[1]
                ."</OPTION>\n";
}
OCIFreeCursor($cursor1);
OCIFreeStatement($stmt);
OCILogoff($conn);
```

It's possible to kick off a stored procedure or Oracle function and walk away, but more often you will be waiting for a result. This result will come on one or more cursors. These cursors, and possibly any other variables that come back to PHP, need to be created by `OCINewCursor()` and bound to PHP variables using `OCIBindByName()`. They are not immediately available to PHP otherwise. After that, they can be treated much like statements — their contents can be fetched, and their memory must be freed.

Project: Point Editor

The *product point editor* is a very simple Oracle tool that allows you to edit the data associated with a single item in a product catalog (for instance, for an e-commerce site). Not only does it fetch information that all products share, such as product name and SKU, but it automatically sweeps up variable product attribute data for editing. This is data associated with each item, such as price, manufacturer, size, color, and so forth; it differs for items depending on what kind of thing they are — cars will have a color but not a gender, clothing might have a gender but not a type of transmission.

Because the attribute information you wish to save for each product varies with the product category — for books you might want to know the author and number of pages, whereas for toys you might want to know recommended age — you cannot simply save this information in the product table, even if it would normally be one-to-one data. To maintain flexibility, each product category (books, toys, and so on) has a number of attributes associated with it, and then each product is associated with a number of attribute values. The relationships look schematically like this:

```
Category
  category_id
  category_name
```

```
Product
  product_id
  product_name
  category_id

Attribute
  attribute_id
  attribute_name
  category_id

Attribute_value
  attrib_val_id
  attrib_val
  attribute_id

Product_attribute
  product_id
  attrib_val_id
```

The product point editor will query the database for all attributes associated with this product and construct a series of pull-down menus with all possible attribute values neatly laid out for the editor to select from. If an attribute value is already associated with this product, that value will be prese-lected in the HTML form field.

Listing 35-1 is a file of functions that are used in both the product point editor and the product batch editor in the next section. This file should be saved under the name `oci8_funcs.php`.

LISTING 35-1

Common Oracle functions (oci8_funcs.php)

```php
<?php

/************************************
 * Functions for Oracle-based tools *
 ************************************/

putenv("ORACLE_HOME=/tools/oracle");


// Use when fetching data from the db
function unescape_quotes($str)
{
    $esc_str = str_replace("''", "'", $str);
    $esc2_str = str_replace("\"\"", "\"", $esc_str);
    return $esc2_str;
}
```

```
// Use when inserting data into the db
function escape_sq($str)
{
    $esc_str = str_replace("'", "''", $str);
    return $esc_str;
}


function escape_html($str)
{
    $gt_str = str_replace("&gt;", ">", $str);
    $lt_str = str_replace("&lt;", "<", $gt_str);
    $dq_str = str_replace("&quot;", "\"", $lt_str);
    $esc_str = str_replace("&amp;", "&", $dq_str);
    return $esc_str;
}


// Use this one for INSERTs, UPDATEs, and DELETEs
function parse_exec_free($conn, $query, &$error_str)
{
    $stmt = OCIParse($conn, $query);
    OCIExecute($stmt, OCI_DEFAULT);
    $err_array = OCIError($stmt);
    if ($err_array) {
        $err_message = $err_array['message'];
        $$error_str = $err_message;
        OCIFreeStatement($stmt);
        $stmt = FALSE;
    } else {
        OCIFreeStatement($stmt);
        $stmt = TRUE;
    }
    return $stmt;
}


// Use this one for SELECTs
function parse_exec_fetch($conn, $query, &$error_str, &$res,
$nulls=0)
{
    $stmt = OCIParse($conn, $query);
    OCIExecute($stmt, OCI_DEFAULT);
    $err_array = OCIError($stmt);
    if ($err_array) {
        $err_message = $err_array['message'];
        $$error_str = $err_message;
        OCIFreeStatement($stmt);
        $stmt = FALSE;
```

```php
   } else {
      if ($nulls == 1) {
         OCIFetchStatement($stmt, $res, OCI_RETURN_NULLS);
      } else {
         OCIFetchStatement($stmt, $res);
      }
   }
   return $stmt;
}


// For batch_upload.php, which writes a separate error log
function choke_and_die($conn, $fp, $error_str)
{
   OCIRollback($conn);
   OCILogoff($conn);
   $error_line = $error_str."<BR>\n";
   echo $error_line;
   fwrite($fp, $error_line);
   fwrite($fp, "</HTML>\n");
   fclose($fp);
   exit;
}


// For all nonlogwriting uses (which is most of them)
function die_silently($conn, $error_str)
{
   OCIRollback($conn);
   OCILogoff($conn);
   // You can uncomment these when debugging
   //$error_line = $error_str."<BR>\n";
   //echo $error_line;
   exit;
}


// Excel sometimes adds random quotes around field contents
function unquote($str)
{
   $pos = strpos($str, "\"");
   if ($pos === 0) {
      $qstr = substr($str, 1, -1);
      return trim($qstr);
   } else {
      return trim($str);
   }
}
```

```
// Excel sometimes doubles double-quotes in an attempt to close
// them
function strip_db($str)
{
    $esc_str = str_replace("\"\"", "\"", $str);
    return $esc_str;
}

?>
```

Listing 35-2 is the actual product point editor itself.

LISTING 35-2

Product point editor (prod_point.php)

```php
<?php

/*****************************************************
 * This is the product point editor.                *
 * The purpose of this tool is to edit all the data  *
 * associated with a single product.  It will mostly *
 * be used for trivial fixes (e.g., spelling errors) *
 *****************************************************/

include("oci8_funcs.php"); //common functions
$thisDB = "dev";
$thisDBuser = "oci_user";
$thisDBpassword = "sesame";

// ------------------
// EDIT PRODUCT DATA
// ------------------
if($_POST['submit'] == "Submit") {
    // Get a timestamp
    $begin_time = time();
    // Open the pipe
    $conn = OCILogon($thisDBuser, $thisDBpassword, $thisDB)
     or die("Can't get a database connection.");

    // UPDATE PRODUCT TABLE
    $product_id = $_POST['product_id'];
    $product_name = escape_sq($_POST['product_name']);
    $sku = escape_sq($_POST['sku']);
    $itemurl = escape_sq($_POST['itemurl']);
    $itemimage = escape_sq($_POST['itemimage']);
```

```
$desc_text = escape_sq($_POST['desc_text']);

$query = "UPDATE product
          SET product_name = '$product_name',
              sku = '$sku',
              itemurl = '$itemurl',
              itemimage = '$itemimage',
              desc_text = '$desc_text',
              modified = SYSDATE
          WHERE product_id = $product_id";
$stmt = parse_exec_free($conn, $query, &$error_str);
if (!$stmt) {
   die_silently($conn, $error_str);
}

// UPDATE PRODUCT_ATTRIB_VAL TABLE
// First blow away all existing rows for this product
$query = "DELETE FROM product_attrib_val
          WHERE product_id = $product_id";
$stmt = parse_exec_free($conn, $query, &$error_str);
if (!$stmt) {
   die_silently($conn, $error_str);
}
if (is_array($_POST['attrib']) &&
   count($_POST['attrib']) > 0) {
   foreach ($_POST['attrib'] as $attrib_id=>$av_id_array) {
     if (is_array($av_id_array) && count($av_id_array) > 0) {
         foreach ($av_id_array as $attrib_val_id) {
             // If attrib value is not Delete All,
             // add new rows
             if ($attrib_val_id != -1) {
                 $query =
                 "INSERT INTO product_attrib_val
                 (attrib_val_id, product_id, modified, created)
                 VALUES($attrib_val_id, $product_id,
                        SYSDATE, SYSDATE)";
                 $stmt =
                   parse_exec_free($conn, $query, &$error_str);
               if (!$stmt) {
                   die_silently($conn, $error_str);
               }
           }
         }
      }
   }
}

OCICommit($conn);
OCILogoff($conn);
```

589

```
    /*
    // Uncomment this block for debugging
    // Get a second timestamp, and do the math
    $end_time = time();
    echo "DONE!  This operation took "
        .($end_time - $begin_time)
        ." seconds to complete.";
    exit;
    */

    // Redisplay the form
    header("Location:  $PHP_SELF?url=$prod_url");
}


// ---------
// SHOW FORM
// ---------
elseif (!isSet($_POST['submit']) ||
        $_POST['submit'] != "Submit") {
    set_time_limit(0);
    // Get a timestamp
    $begin_time = time();
    // Open the pipe
    $conn = OCILogon($thisDBuser, $thisDBpassword, $thisDB)
            or die("Can't get a database connection.");

    // Get the product data based on a unique URL
    //passed in the GET vars
    $url = $_GET['url'];
    if ($url == "") {
        // If a URL isn't passed, spit out a message and quit
        echo "<HTML>\n<BODY>";
        echo '<P>You need to designate a product to edit by
                passing a url like this:
                http://localhost/tools/prod_point.php'.
                '?url=book_PHP5_Bible.</P>';
        echo "</BODY>\n</HTML>";
        exit;
    }
    $query = "SELECT product_id, name, sku, itemurl, itemimage,
                    desc_text, category_id
              FROM product
              WHERE url = '$url'";
    $stmt = parse_exec_fetch($conn, $query, &$error_str, &$res);
    if (!$stmt) {
        die_silently($conn, $error_str);
    } else {
        OCIFreeStatement($stmt);
```

```
    $product_id = $res['PRODUCT_ID'][0];
    $product_name = $res['PRODUCT_NAME'][0];
    $sku = $res['SKU'][0];
    $itemurl = $res['ITEMURL'][0];
    $itemimage = $res['ITEMIMAGE'][0];
    $desc_text = $res['DESC_TEXT'][0];
    $category_id = $res['CATEGORY_ID'][0];
}


// Get attributes for all products in this category
$query = "SELECT attribute_id, attribute_name
           FROM attribute
           WHERE category_id = $category_id";
$stmt = parse_exec_fetch($conn, $query, &$error_str, &$res1);
if (!$stmt) {
    die_silently($conn, $error_str);
    exit;
} else {
    OCIFreeStatement($stmt);
}
if (is_array($res1['ATTRIBUTE_ID']) &&
    count($res1['ATTRIBUTE_ID']) > 0) {
    foreach ($res1['ATTRIBUTE_ID'] as $key=>$attrib_id) {
        $attrib_name = $res1['ATTRIBUTE_NAME'][$key];
        // Get attrib values for this product
        $query = "SELECT product_attrib_val.attrib_val_id
                   FROM product_attrib_val, attrib_val
                   WHERE product_attrib_val.attrib_val_id =
attrib_val.attrib_val_id
                   AND attrib_val.attrib_id = $attrib_id
                   AND product_attrib_val.product_id =
$product_id";
        $stmt = parse_exec_fetch($conn, $query, &$error_str,
&$res2);
        if (!$stmt) {
            die_silently($conn, $error_str);
        } else {
          OCIFreeStatement($stmt);
          // Get all possible attribute values
          //for this attribute
          // and construct nice pulldown lists
          $query = "SELECT attrib_val_id, name
                     FROM attrib_val
                     WHERE attrib_id = $attrib_id
                     ORDER BY name";
          $stmt = parse_exec_fetch($conn, $query,
&$error_str, &$res3);
          if (!$stmt) {
```

591

```
            die_silently($conn, $error_str);
      } else {
         OCIFreeStatement($stmt);
         // This stuff is for Case 2 below
         $is_vals = $res2['ATTRIB_VAL_ID'];
         $num_is_vals = count($is_vals);
         $poss_vals = $res3['ATTRIB_VAL_ID'];
         $num_poss_vals = count($poss_vals);
         $nonmatching = array_diff($poss_vals, $is_vals);

         if ($num_poss_vals > 0) {
           foreach ($poss_vals as
                   $av_key=>$avalue_id) {
             $av_name = $res3['NAME'][$av_key];
             // Existing values are selected in
             // this list.
             // Case 0:  if no existing value
             // then don't highlight any
             if (!is_array($is_vals) ||
                 $num_is_vals == 0) {
               $av_str .= "<OPTION
               VALUE=\"$avalue_id\">$av_name</OPTION>\n";
             }
             // Case 1:  single attrib value
             elseif ($num_is_vals == 1) {
                if ($is_vals[0] == $avalue_id) {
                  $av_str .= "<OPTION VALUE=\"$avalue_id\"
                       SELECTED>$av_name</OPTION>\n";
                } else {
                  $av_str .=
   "<OPTION VALUE=\"$avalue_id\">$av_name</OPTION>\n";
                }
             }
             // Case 2:  multiple attrib values
             // A bit messy because I have to avoid
             // multiple nonmatching options
             elseif ($num_is_vals > 1) {
               foreach ($is_vals as $avid) {
                 if ($avid == $avalue_id) {
                   $av_array[] =
"<OPTION VALUE=\"$avalue_id\" SELECTED>$av_name</OPTION>";
                 }
               }
               if (count($nonmatching) > 0) {
                 foreach ($nonmatching as $avid){
                   if ($avid == $avalue_id) {
                     $av_array[] =
"<OPTION VALUE=\"$avalue_id\">$av_name</OPTION>";
                   }
```

```
                    }
                }
                $av_str = implode("\n", $av_array);
            }
        }
    }
}
$attrib_str .= "$attrib_name ($num_is_vals):  <SELECT
    NAME=\"attrib[$attrib_id][]\"
SIZE=5 MULTIPLE>\n<OPTION VALUE='-1'>Delete All</OPTION>
\n$av_str</SELECT><BR><BR>\n";
    unset($av_array);
    unset($av_str);
    }
    }
    }
OCILogoff($conn);


// ------------
// DISPLAY FORM
// ------------
$php_self = $_SERVER['PHP_SELF'];
// Superglobals don't work with heredoc
$form_str = <<< EOFORMSTR
<HTML>
<HEAD>
<TITLE>Product Point Editor</TITLE>
<STYLE>
<!--
.header    {font-family: verdana, arial, sans-serif;
font-size: 14pt; font-weight: bold; color: #000000;
text-align: left}
.subheader    {font-family: verdana, arial, sans-serif;
font-size: 12pt; font-weight: bold; color: #000000;
background: #ebeef1; text-align: left}
-->
</STYLE>
</HEAD>

<BODY BGCOLOR="#FFFFFF">
<P class="header">Product point editor</P>

<P>The database is <B>$thisDB</B></P>

<P><B>PRODUCT DATA</B></P>
<FORM ACTION="$php_self" METHOD="post">
<INPUT TYPE=HIDDEN NAME="product_id" VALUE="$product_id">
Name:  <INPUT TYPE=TEXT NAME="product_name" SIZE=30
```

```
VALUE="$product_name"><BR><BR>
SKU #:  <INPUT TYPE=TEXT NAME="sku" SIZE=70
VALUE="$sku"><BR><BR>
Item URL:  <INPUT TYPE=TEXT NAME="itemurl" SIZE=70
VALUE="$itemurl"><BR><BR>
Item Image:  <INPUT TYPE=TEXT NAME="itemimage" SIZE=70
VALUE="$itemimage"><BR><BR>
Description:  <TEXTAREA NAME="desc_text" COLS=50
ROWS=5>$desc_text</TEXTAREA><BR><BR>

<P><B>ATTRIBUTES</B></P>
$attrib_str
<INPUT TYPE=SUBMIT NAME="submit" VALUE="Submit">
</FORM>
</BODY>
</HTML>
EOFORMSTR;
    echo $form_str;

    // Get a second timestamp, and do the math
    $end_time = time();
    echo "DONE!  This operation took ".
        ($end_time - $begin_time)
        ." seconds to complete.<BR>\n";
}

?>
```

For single-product editing, a PHP form that makes a direct connection to the Oracle database is not noticeably slower than one that employs a stored procedure. This tool also has the advantage that it can be altered by a PHP developer alone, whereas use of a stored procedure usually also requires time from a PL/SQL programmer.

Project: Batch Editor

To edit data for more than one product at a time, you might want to use stored procedures. This tool, the *product batch editor*, has two main parts. A script called `header_download.php` downloads the attributes for a particular category to a spreadsheet. A corresponding script called `batch_upload_new.php` allows the user to upload data from a spreadsheet to the server, where it is loaded into the database. This is the simplest use of stored procedures.

Listing 35-3 is the first stored procedure called in `header_download.php`. It is called `get_categories.sql`, and it merely downloads a complete list of all the product categories in this schema.

LISTING 35-3

Stored procedure (get_categories.sql)

```
CREATE OR REPLACE PROCEDURE get_categories(
  category_list_out OUT pack.my_targets)
IS

BEGIN
  OPEN category_list_out FOR
    SELECT category_id, category_url
    FROM CATEGORY
    ORDER BY category_url;
END;
/
show errors
```

After you use the form to select a category, `header_download.php` calls the stored procedure called `get_cat_header` shown in Listing 35-4. This takes the category ID as an input and returns two cursors: one consisting of some common product fields, and the other consisting of attribute types associated with this category.

LISTING 35-4

A second stored procedure (header_download.php)

```
create or replace procedure get_cat_header(
    category_id_in INTEGER,
    cat_header_out OUT pack.my_targets,
    cat_attrib_out OUT pack.my_targets,
IS
    v_action VARCHAR2(2) := '01';
BEGIN
    IF category_id_in is NULL THEN
     OPEN cat_header_out for select NULL from DUAL;
     OPEN cat_attrib_out for select NULL from DUAL;
    END IF;

    open cat_header_out for
     SELECT
     column_name,
     column_display_name,
     column_order
     FROM
     event_table_columns
     WHERE
     table_name = 'product'
     ORDER BY column_order;
```

595

```
    open cat_attrib_out for
     SELECT attribute_id, attribute_name
     FROM attribute
     WHERE category_id = category_id_in;
END;
/
show errors
```

Listing 35-5 is `header_download.php`, which shows the form to call both stored procedures.

LISTING 35-5

New product form (header_download.php)

```php
<?php

/********************************************
 * New product download attributes script. *
 * The purpose of this tool is to download  *
 * a spreadsheet with the product data      *
 * header.  Editors will use this to add    *
 * new items to a category.  Use script     *
 * batch_upload_new.php to upload data.     *
 ********************************************/

include("oci8_funcs.php"); //common functions for Oracle tools
$thisDB = "dev";
$thisDBuser = "oci_user";
$thisDBpassword = "sesame";

// Open the pipe
$conn = OCILogon($thisDBuser, $thisDBpassword, $thisDB);

// ----------------------
// GET THE CATEGORY HEADER
// ----------------------
if ($_POST['submit'] == "Add") {
   // Call stored procedure for this category
   $cat_id_in = $_POST['cat_id'];
   $request = "begin DEV.get_cat_header($cat_id_in,
:OUT1, :OUT2); end;";
   $cursor1 = OCINewCursor($conn);
   $cursor2 = OCINewCursor($conn);
   $stmt = OCIParse($conn, $request);
   OCIBindByName($stmt, ":OUT1", &$cursor1, -1, OCI_B_CURSOR);
   OCIBindByName($stmt, ":OUT2", &$cursor2, -1, OCI_B_CURSOR);
   OCIExecute($stmt);
```

```
    OCIExecute($cursor1);
    OCIExecute($cursor2);
    $err_array = OCIError($conn);
    if ($err_array) {
        $err_message = $err_array['message'];
        echo $err_message;
        OCIFreeCursor($cursor1);
        OCIFreeCursor($cursor2);
        OCIFreeStatement($stmt);
        OCILogoff($conn);
        exit;
    }
    while (OCIFetchInto($cursor1,&$data1)) {
        $p_array[] = $data1[1];
    }
    while (OCIFetchInto($cursor2,&$data2)) {
        $a_array[] = $data2[1]."|".$data2[0];
    }
    OCIFreeCursor($cursor1);
    OCIFreeCursor($cursor2);
    OCIFreeStatement($stmt);
    OCILogoff($conn);
    // ASSEMBLE THE DOWNLOAD
    $init_p_str = implode("\t", $p_array);
    $p_str = str_replace("CATEGORY_ID", $cat_id_in, $init_p_str);
    if (count($a_array) > 0) {
        $a_str = implode("\t", $a_array);
    }
    $full_header = implode("\t", array($p_str, $a_str));

    // SEND THE FILE
    $header_file = 'header.xls.Z';
    $zp = gzopen($header_file, "w+");
    gzwrite($zp, $full_header);
    gzclose($fp);
    header("Location:  header.xls.Z");
    // For IE5.x, this is the correct way to trigger a download
    //--by simply directing the browser to download a file type
    // that the browser cannot open
}


// -----------------
// CHOOSE A CATEGORY
// -----------------
elseif (!isSet($_POST['submit'])) {
    // Call stored procedure get_categories
    $request = "begin DEV.get_categories(:OUT1); end;";
    $cursor1 = OCINewCursor($conn);
    $stmt = OCIParse($conn, $request);
```

597

```
    OCIBindByName($stmt, ":OUT1", &$cursor1, -1, OCI_B_CURSOR);
    OCIExecute($stmt);
    OCIExecute($cursor1);
    $err_array = OCIError($conn);
    if ($err_array) {
        $err_message = $err_array['message'];
        echo $err_message;
        OCIFreeCursor($cursor1);
        OCIFreeStatement($stmt);
        OCILogoff($conn);
        exit;
    }
    while (OCIFetchInto($cursor1, &$cat_array)) {
        $opt_str .= "<OPTION VALUE=\""
                    .$cat_array[0]."\">"
                    .$cat_array[1]."</OPTION>\n";
    }
    OCIFreeCursor($cursor1);
    OCIFreeStatement($stmt);
    OCILogoff($conn);


    // CHOOSE CATEGORY FORM
    $php_self = $_SERVER['PHP_SELF'];
    // Superglobals don't work with heredoc
    $form_str = <<< EOFORMSTR
<HTML>
<HEAD>
<TITLE>Batch Editor New: Download</TITLE>
<STYLE>
<!--
.header    {font-family: verdana, arial, sans-serif
font-size: 14pt; font-weight: bold; color: #000000;
text-align: left}
.subheader    {font-family: verdana, arial, sans-serif;
font-size: 12pt; font-weight: bold; color: #000000;
background: #ebeef1; text-align: left}
.ftrnote    {font-family: verdana, arial, sans-serif;
font-size: 8pt; color: #000000; text-align: left}
LI     {line-height:200%}
-->
</STYLE>
</HEAD>

<BODY BGCOLOR="#FFFFFF">
<P class="header">Batch editor new:  download</P>

<P>The database is <B>$thisDB</B></P>
```

```
<FORM ACTION="$php_self" METHOD="POST">
<SELECT NAME="cat_id" SIZE=1>
<OPTION VALUE="-1" SELECTED>Choose one</OPTION>
$opt_str
</SELECT><BR>
<INPUT TYPE=SUBMIT NAME="submit" VALUE="Add">
</FORM>
</BODY>
</HTML>
EOFORMSTR;

    echo $form_str;
}
?>
```

The result of `header_download.php` is a zipped spreadsheet with one row. It resembles the first row of the image in Figure 35-1.

FIGURE 35-1

Category header with some data filled in

We have filled in some data in the preceding figure, using the header row as a guide.

Note that the attributes at the top have names like `"Price|376"`. The number is the attribute ID, which will be used when we upload data with Listing 35-6, `batch_upload_new.php`. This script enables the user to upload a spreadsheet to the server, using the normal HTML file upload capability of his browser. The server will then convert the spreadsheet into rows of data, which are inserted into the database.

LISTING 35-6

Spreadsheet upload script (batch_upload_new.php)

```php
<?php

/*******************************************
 * New product batch upload script. The    *
 * purpose of this tool is to upload a     *
 * spreadsheet with new product data.      *
 * Editors will use this to add new items  *
 * to a category.  Use script              *
 * header_download.php to get attributes.  *
 *******************************************/

include("oci8_funcs.php"); //common functions for Oracle tools
$thisDB = "dev";
$thisDBuser = "oci_user";
$thisDBpassword = "sesame";


// HEADER
$header_str = <<< ENDOFHEADER
<HTML>
<HEAD>
<TITLE>Batch Editor New: Upload</TITLE>
<STYLE>
<!--
.header    {font-family: verdana, arial, sans-serif;
font-size: 14pt; font-weight: bold; color: #000000;
text-align: left}
.subheader    {font-family: verdana, arial, sans-serif;
font-size: 12pt; font-weight: bold; color: #000000;
background: #ebeef1; text-align: left}
-->
</STYLE>
</HEAD>

<BODY BGCOLOR="#FFFFFF">
<P class="header">Batch editor new: upload</P>
```

```
<P>The database is <B>$thisDB</B></P>
ENDOFHEADER;
echo $header_str;


// ADD NEW PRODUCTS
if($_POST['submit'] == "Upload") {
   set_time_limit(0);
   echo "<P>Check the error log
      (<A HREF\"upload_log.html\">upload_log.html</A>)
       for problems.</P>";

   // Copy uploaded file to a specific directory
   $tempfile = $HTTP_POST_FILES[file][tmp_name];
   $localfile = $HTTP_POST_FILES[file][name];
   if(!copy($tempfile, "/tmp/$localfile")) {
      echo "<P>Error writing file to upload directory.
            Quitting.</P>\n";
      exit;
   }

   // Start an error log
   $error_log = 'upload_log.html';
   $fp = fopen($error_log, "w+") or die("Can't open error log.");
   fwrite($fp, "<HTML>\n");

   // Open the pipe
   $conn = OCILogon($thisDBuser, $thisDBpassword, $thisDB)
        or die("Can't get a database connection.");

   // Get a timestamp
   $begin_time = time();

   // Read in the data file as an array
   $uarray = file("/tmp/$localfile");

   // Parse the header for cat_id and attributes
   $header = array_shift($uarray);
   $harray = explode("\t", $header);
   $num_ha = count($harray);
   $cat_id = $harray[0];
   $attrib_array = array();
   for($i = 6; $i <= ($num_ha - 1); $i++) {
      $a_array = explode("|", $harray[$i]);
      $attrib_array[] = $a_array[1];
   }
   $num_attribs = count($attrib_array);
```

601

```
$error_str = "";
$res = array();
// Get all the attrib values and stick them in a
// multidimensional array
foreach($attrib_array as $attrib) {
    $query = "SELECT attrib_value_id, name
              FROM attrib_value
              WHERE attrib_id = $attrib";
    $stmt =
      parse_exec_fetch($conn, $query, &$error_str, &$res);
    if (!$stmt) {
       choke_and_die($conn, $fp, $error_str);
    } else {
      foreach($res['NAME'] as $key => $val) {
         $ava[$attrib][$val] = $res['ATTRIB_VALUE_ID'][$key];
      }
      OCIFreeStatement($stmt);
    }
}
reset($attrib_array);

// Shove the data down the pipe.
foreach ($uarray as $valrow) {
    // Get a fresh product id from Oracle
    $query = "begin :new_id := newid('product'); end;";
    $sth = OCIParse($conn, $query);
    OCIBindByName($sth, ":new_id", &$new_id, 200);
    OCIExecute($sth);
    if (!$sth) {
       choke_and_die($conn, $fp, $error_str);
    } else {
      $rowid = $new_id;
      OCIFreeStatement($sth);
    }

    // Format new product data.
    $val_array = explode("\t", $valrow);
    $prod_name = unquote($val_array[1]);
    $prod_name = strip_db($prod_name);
    $prod_name = escape_sq($prod_name);
    echo "Working on $prod_name<BR>\n";
    $sku = unquote($val_array[2]);
    $itemurl = unquote($val_array[3]);
    $itemimage = unquote($val_array[4]);
    $desc = unquote($val_array[5]);
    $desc = escape_sq($desc);

    // PRODUCT
    $query = "INSERT INTO product (
```

```
                    product_id, name, sku, itemurl, itemimage,
                    desc, created, modified, category_id)
                    VALUES (
                    $rowid, '$prod_name', '$sku', '$itemurl',
                    '$itemimage', '$desc', SYSDATE, SYSDATE,
                    $cat_id)";
        $stmt = parse_exec_free($conn, $query, &$error_str);
        if (!$stmt) {
            choke_and_die($conn, $fp, $error_str);
        }

        // PRODUCT_ATTRIB_VALUE
        for ($i = 6; $i <= (6 + $num_attribs - 1); $i++) {
            $av = unquote($val_array[$i]);
            if($av != "") {
                $temp_q = explode("|", $av);
                foreach($temp_q as $av) {
                    $av = unasterisk($av);
                    $av = escape_sq($av);
                    $akey = $i - 6;
                    $attrib_id = $attrib_array[$akey];
                    if($ava[$attrib_id][$av]) {
                     $pav = $ava[$attrib_id][$av];
                     $query = "INSERT INTO product_attrib_value (
                             attrib_value_id, product_id, created,
                             modified)
                             VALUES(
                             $pav, $rowid, SYSDATE, SYSDATE)";
                     $stmt =
                       parse_exec_free($conn, $query, &$error_str);
                     if (!$stmt) {
                         choke_and_die($conn, $fp, $error_str);
                     }
                  }
               }
            } else {
               //echo "Null attrib value.<BR>\n";
            }
        }
    }

// Get a second timestamp, and do the math
$end_time = time();
echo "DONE!  This operation took "
    .($end_time - $begin_time)
    ." seconds to complete.";

OCICommit($conn);
OCILogoff($conn);
```

```
    fwrite($fp, "</HTML>\n");
    fclose($fp);
}


// SHOW FILE UPLOAD FORM
elseif($_POST['submit'] != "Upload") {
    $upload_str = <<< ENDOFUPLOAD
<P>Upload new product data:</P>
<FORM ACTION="$PHP_SELF" METHOD="post"
ENCTYPE="multipart/form-data">
<INPUT TYPE=HIDDEN NAME="max_file_size" VALUE="1000000">
<INPUT TYPE=FILE NAME="file" SIZE=50><BR><BR>
<INPUT TYPE=SUBMIT NAME="submit" VALUE="Upload">
</FORM>
ENDOFUPLOAD;
    echo $upload_str;
}

?>
</BODY>
</HTML>
```

Notice particularly in this script how we call for a new product ID using Oracle's incrementor (search for the comment `// Get a fresh product id from Oracle`). This is in marked contrast to, for instance, MySQL, where the auto-incrementor is generally built into a column's definition and can be kicked off by merely entering a null value in that column.

Summary

Oracle is one of the biggest names in enterprise databases. The PHP OCI8 extension is quite fast and powerful, offering some functionality not found in most other database extensions. However, the use of Oracle has many implications and should not be undertaken lightly.

There are significant differences between the syntax of Oracle functions and those of lighter-weight databases. These differences include so-called Sybase-style string escaping, separate parsing and execution steps, manual memory management, cursors, transactionality, and the capability to fetch data sets into one big array. We demonstrated the use of all these Oracle features in two examples: a single-product data editor and a mass data entry tool.

Chapter 36

An Introduction to SQLite

W hen coding PHP, you may sometimes find that you want a database that's both easy to use and lightweight, without having the complexity of a larger RDBMS like MySQL or Oracle. Small applications, written for one-off solutions, lend themselves to these types of databases, as do applications that just need a powerful query functionality. Enter SQLite. SQLite is a lightweight database system that's bundled with PHP by default. As you'll see in this chapter, SQLite is a lightweight database system appropriate for many uses in PHP (and other languages, too).

An Introduction to SQLite

SQLite is a database engine that operates without the use of a database server. You may find yourself reading that sentence again, but it's true. SQLite has no server component, but rather directly interacts with its data through files or directly in memory. SQLite works well with small (to medium) data sets or in low-volume transactions, but because the storage medium is usually disk, accessing the data can be slower than a more robust and traditional RDBMS like MySQL.

SQLite doesn't have a privilege system in the way that MySQL does. This means that it isn't possible to limit who has access to certain databases or tables within a given database. The files holding data in SQLite are typically owned and readable by the user under which the web server is running. With these issues in mind, it's not appropriate to use SQLite to store sensitive data.

This chapter gives a brief introduction to the most basic usage of SQLite. More information can be found in the online manual at `www.php.net/sqlite`.

Using SQLite-related Functions

SQLite is simple to use. So much so that working with SQLite may confuse those readers who are accustomed to some level of complexity when working with databases.

Creating Databases

SQLite databases are created simply by calling the `sqlite_open()` function, like this:

```
$sqldb = sqlite_open("mydatabase.sqlite");
```

Alternately, an object-oriented interface is also available to create new databases:

```
$sqldb = new SQLiteDatabase("mydatabase.sqlite");
```

You will concentrate on the object-oriented interface to SQLite in this chapter. More information on the functional interface can be found in the online PHP manual for SQLite at `http://us2.php.net/sqlite`.

Running Queries

In the realm of SQLite, queries, including most interactions with the database such as `SELECT`, `INSERT`, `UPDATE`, and `DELETE` operations, take place through the `query()` method. For example:

```
$sqldb->query("SELECT * FROM foo");
```

This section examines the `query()` method as it is used to create tables and insert, and retrieve data.

SQLite also has a method called `exec()`, which can be used for statements that return no results, such as statements that create tables and the like. Examining SQLite code "in the wild," it's much more common to see the `query()` method used, for better or worse, than it is to see `exec()`; therefore, it's `query()` that you'll concentrate on in this section.

Creating Tables

SQLite stores its data as strings regardless of the type specified in the table creation statement. This means that SQLite can be slow when sorting data.

Automatically incremented columns, like those frequently used for identifier columns, are created by SQLite without need for an additional keyword such as `auto_increment`, simply by specifying that an integer column is the primary key.

Here's an example that creates a table to store the geopolitical states of the union in the United States. In this example, it is assumed that the `$sqldb` handle has been created.

```
$sqldb->query("CREATE TABLE mytable (id integer primary key, stateabb
char(2), state char(50))");
```

Table creation need only take place once. Attempting to create a table that already exists causes a PHP warning, like the one shown in Figure 36-1.

FIGURE 36-1

Attempting to create a table that already exists results in a warning

Therefore, it's up to you either to ignore the warning or comment out code that creates tables if they already exist. One method for ignoring the error (though we definitely do not recommend ignoring errors) would be to precede the call to the query() method with an "at" sign, @.

```
@$sqldb->query("CREATE TABLE mytable (id integer primary key, stateabb
char(2), state char(50))");
```

Another method available in newer versions of SQLite is to use the 'if not exists' syntax when creating a table. Using 'if not exists' tells SQLite to create the table only if the table isn't already there. It is preferred to use this syntax rather than ignoring the error using the @ method shown previously.

Inserting Data

Like other statements, inserting data into an SQLite database is also accomplished using the query() method. Going back to the example table created in the last section, you could insert some states and their respective abbreviation with these statements:

```
$sqldb->query("INSERT INTO mytable (stateabb,state) VALUES
('WI','Wisconsin')");
$sqldb->query("INSERT INTO mytable (stateabb,state) VALUES
('CA','California')");
$sqldb->query("INSERT INTO mytable (stateabb,state) VALUES
('IL','Illinois')");
$sqldb->query("INSERT INTO mytable (stateabb,state) VALUES
('OR','Oregon')");
$sqldb->query("INSERT INTO mytable (stateabb,state) VALUES
('MN','Minnesota')");
$sqldb->query("INSERT INTO mytable (stateabb,state) VALUES
('AL','Alabama')");
$sqldb->query("INSERT INTO mytable (stateabb,state) VALUES
('WA','Washington')");
```

Whenever you're working with input from the web, as you frequently are in PHP, it's critical to make sure that the data has been checked to ensure that it contains valid and acceptable values. For SQLite, this means using the sqlite_escape_string() function on any input from outside of your program, such as $_POST and $_GET, but also from cookies or even server and environment variables in a shared hosting environment.

Assume that you're accepting a POST from a web form containing the elements username and password. This code would take those input values and make them safe to use in SQLite statements:

```
$safe_user = sqlite_escape_string($_POST['username']);
$safe_pass = sqlite_escape_string($_POST['password']);
```

Now instead of using $_POST['username'] within your SQL statement, you would use $safe_user instead.

Fetching Data

While there are several means by which data can be retrieved from SQLite tables, the simplest method for retrieving data from SQLite is accomplished by using the fetch() method. The fetch() method is an alias for sqlite_fetch_array, which means that data retrieved by this method will be placed into an array.

Going back to the state example from earlier, the data could be retrieved simply with the following statement:

```
$result = $sqldb->query("SELECT * from mytable");
```

Then the result is iterated by looping through the result set with the `fetch()` method.

```
while ($row = $result->fetch()) {
    print_r($row);
}
```

The output from this is shown in Figure 36-2, again assuming the inserts from earlier were executed in order to insert some sample data.

FIGURE 36-2

Printing data from an SQLite table using the fetch() method

Another method for retrieval of data from SQLite is the `fetchAll()` method. `Fetchall()` differs from `fetch()` in that with `fetchAll()`, all data is brought back immediately. This is appropriate for smaller data sets where you want to grab everything and then work with it as necessary in your program. The implication of fetching an entire data set, if that data set is large, is that it may consume unnecessary resources versus simply looping through the results with `fetch()`. The `fetchAll()` method retrieves data as an array of arrays.

Other methods also exist for retrieving data from SQLite. These additional methods won't be covered here but you can find them online at `www.php.net/sqlite`.

More on SQLite

You may find that you'll never encounter a need for SQLite in your PHP projects, especially with the ubiquity and ease of use of MySQL. If you do, however, there is much more to SQLite than we've covered here, in a book on PHP and MySQL. If you're going to use SQLite, please visit the SQLite section of the PHP manual for more information, `www.php.net/sqlite`.

Summary

This chapter gave a brief overview of SQLite and its basic usage in PHP. We looked at creation of databases in SQLite using PHP, how to create a table, insert data, and how to select data from that table. SQLite isn't appropriate for all usages of database systems, however. For example, SQLite doesn't contain a privilege system, so anyone that has access to the file containing your data has access to all of the data.

More information on SQLite and PHP can be found in the online manual at `www.php.net/sqlite`.

Part V

Connections

Chapter 37

Sending E-Mail with PHP

This chapter is all about using PHP (and, in some cases, databases) to send e-mail. If you're looking to receive e-mail with PHP, by using a protocol such as POP3 or IMAP, please refer to the PHP manual for those functions.

Sending E-Mail with PHP

Sending mail is where PHP really comes into its own. But before you can send any mail from your server, you need to tweak the configuration file a little.

Windows configuration

In Windows, you need to set two variables in the `php.ini` file:

- `SMTP`: A string containing the DNS name or IP address of an SMTP server that relays for the Windows machine on which PHP is installed. If it is on the PHP server, specify `localhost`.

- `sendmail_from`: A string containing the e-mail address of your default PHP mail sender (for example, `mailbot@example.com`).

> **TIP** IIS4+ has an SMTP server built in, which is lighter than Exchange Server, if you don't need the power of the latter.

Linux configuration

You need to check and possibly change one variable in the `php.ini` file if you're using Unix: `sendmail_path`, a string containing the full path to your sendmail program (usually `/usr/sbin/sendmail` or `/usr/lib/sendmail`), a replacement, or a wrapper (such as `/var/qmail/bin/sendmail`).

The mail function

The `mail()` function is the primary function used to send mail with PHP. This function, which returns a Boolean, attempts to send one message using the data within the parentheses. The simplest use of this function (keeping in mind that this is a dummy address and should not be used for testing purposes) is:

```php
<?php
mail('receiver@example.com', 'A Sample Subject Line',
"Body of e-mail\r\nwith lines separated by the newline
character.");
?>
```

This is the default and minimum format: address of recipient, subject line, and body. In this case, PHP will automatically add a `From: me@sendhost` line to each message header.

You can also, as always, use variables instead of hardcoded values:

```php
<?php
$address = 'santa@example.com';
$subject = 'All I want for Christmas';
$body = "Is my two front teeth.\r\nSincerely, Joey";
$mailsend = mail($address, $subject, $body);
echo $mailsend;
?>
```

Multiple recipients all go into the address field, with commas separating them (this feature is not supported by all MTAs; if you want to be sure, use `cc:` instead):

```php
<?php
$address1 = 'receiver@receipthost';
$address2 = 'jane@example.com';
$address3 = 'john@example.org';
$all_addresses = "$address1, $address2, $address3";
$subject = 'A Sample Subject Line';
$body = "Body of e-mail\r\nwith lines separated by the
newline character.";

$mailsend = mail($addresses, $subject, $body);
echo $mailsend;
?>
```

Remember to ensure that the multiple addresses are one string, as in the preceding code lines. You do *not* want to do this:

```php
<?php
$address1 = 'receiver@receipthost';
$address2 = 'jane@example.com';
$address3 = 'john@example.org';
$subject = 'A Sample Subject Line';
$body = "Body of e-mail\r\nwith lines separated by the newline
character.";

// This is wrong, don't do it!
$mailsend = mail($address1, $address2, $address3, $subject,
$body);
echo $mailsend;
?>
```

Most people would like more control over the addresses, appearance, and format of their e-mails. You can do that by putting an additional header *after* the three default headers.

```php
<?php
$address = 'receiver@receipthost';
$subject = 'A Sample Subject Line';
$body = "Body of e-mail\r\nwith lines separated by the newline
character.";
$extra_header_str = "From: me@sendhost\n\nbcc:
phb@sendhost\r\nContent-type: text/plain\r\nX-mailer: PHP/"
. phpversion();

$mailsend = mail($address, $subject, $body, $extra_header_str);
echo $mailsend;
?>
```

This "additional header" field is somewhat odd because it crams in several types of information that would normally be given their own fields. Ours is not to wonder why; ours is but to explain the kinds of things you might want to put in this field.

- Your name
- The To: address
- Your e-mail address
- A reply-to or bounce-to address
- X-mailer and version number
- MIME version
- Content-type
- Charset (which uses a = to assign a value and not a : like the other headers)

- Content-transfer-encoding
- Carbon-Copy (cc:) and blind carboncopy (bcc:) recipients

CAUTION The `mail()` function returns `1` (`TRUE`) *when PHP believes it has successfully sent mail.* *This has no relationship to any mail actually being sent or received. There are still an* endless number of things that can go wrong: bad e-mail address, SMTP daemon incorrectly designated or configured, local Internet conditions, and so on. Think of `1` as a message meaning no more than "PHP has applied the function to the inputs successfully."

Sending Mail from a Form

Sending mail from a form is quite likely the single most popular application of PHP's `mail()` function. It's a far more functional alternative than HTML's `mailto` link tag, which of course results in e-mail being sent from the client machine's mail program.

Listing 37-1 is a simple example form of the type that often sends e-mail.

LISTING 37-1

e-mail form (titlehelp.html)

```
<html>
<head>
<title>titlehelp.html</title>
</head>

<body>
<center>
<table width="550">
<tr bgcolor= #FF9933><td align="center"><BR>
<H3>The Thriller <BR>
"What was the name of that thriller?"<BR>
Form</H3></td></tr>
<tr><td>
Did you once read an unforgettable thriller, but now you can't
remember the name?  Fill out as many of the fields below as you
can, press the button to submit, and we'll search our sources
and e-mail you back.
</td></tr></table>
</center>

<FORM METHOD=post ACTION="titlehelp.php">
<P>First name: <input type="text" size=30 name="FirstName">
<P>Last name: <input type="text" size=30 name="LastName">
<P>Your Email Address: <input type="text" size=30 name="Email">
<P>In approximately what year did the action of the book occur?
```

```
<input type="text" size=4 name="Year">
<P>Can you remember any settings from the book?
<input type="text" size=30 name="Setting">
<P>The gender of the protagonist(s) was: <br>
<ul>
<input TYPE="radio" NAME="Gender" VALUE=1>Female<br>
<input TYPE="radio" NAME="Gender" VALUE=2>Male<br>
<input TYPE="radio" NAME="Gender" VALUE=3>One of each<br>
<input TYPE="radio" NAME="Gender" VALUE=4>Two males<br>
<input TYPE="radio" NAME="Gender" VALUE=5>Two females<br>
</ul>
<P>When the book first came out, it was: <br>
<ul>
<input TYPE="radio" NAME="Status" VALUE=1>A bestseller<br>
<input TYPE="radio" NAME="Status" VALUE=2>A critic's darling<br>
<input TYPE="radio" NAME="Status" VALUE=3>Neither<br>
<input TYPE="radio" NAME="Status" VALUE=4>I don't know<br>
</ul>
<P>Please tell us anything else you can remember about this
title (plot, characters, settings, cover, movie versions,
etc.):
<br><textarea NAME="Other" ROWS=6 COLS=50></textarea>
<P><input type="submit" name="Submit">
</body>
</html>
```

Listing 37-1 submits to a form handler, shown in Listing 37-2.

LISTING 37-2

E-mail form handler (titlehelp.php)

```
<html><head>
<title>titlehelp.php</title>
</head>

<body>
<?php
// If you wished, you could also save this information to
// a database
$LastName = $_POST['LastName'];
$FirstName = $_POST['FirstName'];
$Year = $_POST['Year'];
$Setting = $_POST['Setting'];
$Gender = $_POST['Gender'];
$Status = $_POST['Status'];
$Other = $_POST['Other'];
```

```
$formsent = mail('help@example.com',
                 'What was the name of that thriller?',
                 "Request from: $LastName $FirstName\r\n
                   Year: $Year\r\n
                   Setting(s): $Setting\r\n
                   Protagonist gender: $Gender\r\n
                   Book status: $Status\r\n
                   Other identifying characteristics: $Other",
                 "From: help@example.com");
if ($formsent) {
  echo "<P>Hi, $FirstName.  We have received your request for
help, and will try to respond within 24 hours.  Thanks for
visiting!";
} else {
  echo "I'm sorry, there's a problem with your form.  Please try
again.";
}
?>
</body>
</html>
```

There are security implications for using any $_POST variables when sending mail. Spammers may use this form for sending spam. A spammer would be able to craft a message into the value of $FirstName and have it sent out. One solution is to validate the input and accept only single words for first names and then only of a certain length.

Summary

E-mail is one of the most useful and attractive functions of the Internet. PHP gives you the ability to both send and receive e-mail from a web page. One of the most common uses of PHP's mail function is to send an e-mail (often to yourself) with values generated from a Web form.

Chapter 38

Integrating PHP and Java

The relationship between PHP and Java has changed significantly with each new release. Unsurprisingly, given the source code, PHP initially had much more in common with C. PHP4 supported integration of PHP and Java using a Java servlet environment or, more experimentally, directly into PHP. Finally, with the overhaul of the object model in PHP5, there was a distinctly Java feel to the PHP approach to object-oriented programming. Java users will find the improved object model in PHP6 very familiar, although there are important differences.

Given these changes, as PHP takes on a more Java-like cast, there are two possibilities for which a discussion of PHP and Java might be pertinent. You might need to work on a project that requires PHP and Java or Java Server Pages (JSP) to work in tandem. Or you may be approaching PHP from a Java background and want to know about the similarities and differences in order to learn PHP faster. We will deal with both needs in this chapter.

If you don't have a need to use Java, or aren't already familiar with the language, this chapter won't do much for you.

PHP for Java programmers

Most projects won't require integration of Java and PHP, unless there is some specific need due to preexisting architecture. The Java programmer approaching PHP for the first time may still want to know more about how PHP compares to Java for the purposes of learning PHP scripting.

Similarities

In this section, we discuss some ways in which PHP and Java are similar.

Syntax

Though PHP syntax is much closer to C, many conventions used in Java apply to PHP as well. Code is whitespace-insensitive, statements are terminated with semicolons, function calls have a similar structure (`my_function(expression1, expression2)`), and curly braces (`{` and `}`) make statements into blocks. PHP supports C and C++-style comments (`/* */` as well as `//`), which are also used in Java.

Operators

The assignment operators (`=`, `+=`, `*=`, and so on), the Boolean operators (`&&`, `||`, `!`), and the basic arithmetic operators (`+`, `-`, `*`, `/`, `%`) all behave as they do in Java. Other operators are similar, with some syntax differences. The string concatenation operator, for example, in PHP is a period (`.`) rather than a plus sign (`+`) as in Java.

Object model

The Java programmers coming to PHP with version 5 can rejoice! You no longer need to unlearn your approach to OOP in order to deal with the often crude OOP support in PHP4 and earlier versions. The PHP5 addition of keywords (such as `private`, `protected`, and `public`) for dealing with member variables should prove familiar to Java coders. New error-handling methods, including the built-in `Exception` class, will also be familiar to Java programmers.

CROSS-REF **See Chapter 20 for much more detailed information about PHP's object model.**

Memory management

Under normal circumstances, PHP's garbage-collected environment ensures that you do not need to explicitly free allocated memory. If you're used to Java's mostly automated garbage-collected heap, you'll be right at home here.

Packages and libraries

Many web-specific libraries are built into PHP and are available by default or with minor changes. This works similarly to the standard Java packages that are available through JAR files and `CLASSPATH` references.

Differences

Although many of the features of PHP have a Java-like feel, there are plenty of notable exceptions to the way Java and PHP operate. As a general rule, never assume that a Java feature or concept will carry over completely into PHP.

Compiled versus scripting

Unlike Java, PHP is a *scripting* language. The development cycle is *edit-execute* rather than *edit-compile-execute*, as in Java. PHP code is automatically compiled at execution time and does not produce native standalone executables. As a result, the developer is not subjected to rigorous compile-time error checking as in Java; many of the errors that you are used to seeing at compile time will not rear their ugly heads until the code is executed in PHP.

Variable declaration and loose typing

Get used to that leading $. Unlike in Java, all variables in PHP must begin with a $. Variables need not be declared before use, nor cast to a different type as in Java. Rather than the Java code:

```
String preamble = new String();
Preamble = "We, the people...";
```

or:

```
String preamble = "We, the people...";
```

the corresponding PHP code would be simply:

```
$preamble = "We, the people...";
```

PHP utilizes dynamic typing; the variable has no intrinsic type and can change with each new statement. While the following code is perfectly legal in PHP:

```
$type = 11;
$type = "11";
```

you would need to use separate variables in Java, or attempt to recast the variable as a `String`, with potential problems and a resultant performance hit.

Variables can be declared and typed as in Java, but this is not required in PHP.

Java Server Pages and PHP

PHP can fulfill many functions similarly to Java Server Pages (JSP). The JSP servlet engine serves as a scripting language for use with Java, and, just like PHP, is often used in frontend applications.

Embedded HTML

PHP is more similar to JSP than Java itself in that you are allowed to write HTML directly rather than using endless `print` statements. Unlike Java, but like JSP, variables can also be referenced from within a block of HTML. A simple HTML page using JSP script might look like this:

```
<%
  String greeting = "Hello, world";
%>
```

621

```
<HTML>
<HEAD>
  <TITLE>Fun with JSP</TITLE>
</HEAD>
<BODY>
<H1><%= greeting %></H1>
</BODY>
</HTML>
```

Similarly, using PHP, you can write:

```
<?php
  $greeting = "Hello, World";
?>
<HTML>
<HEAD>
  <TITLE>Fun with PHP</TITLE>
</HEAD>
<BODY>
<H1><?php echo $greeting ?></H1>
</BODY>
</HTML>
```

Pages can freely alternate between HTML and JSP, just as when you use HTML and PHP.

Choose your scripting language

PHP can actually be used with Java in lieu of JSP, although support is much less robust and subject to change in future releases. Many of the *shortcuts* available to JSP are not available through PHP. For example, many of the standard Java class packages are automatically available as in Java Server Pages but must be implicitly referenced using PHP Java support.

CAUTION **Don't let new syntax and structure similarities lull you into believing that PHP is truly like Java. The forgiving nature of PHP and the loose treatment of variable typing mean that code must be treated very differently. Those errors that Java demanded you to fix before it would compile may not show up in PHP until they are in an end user's browser!**

Integrating PHP and Java

In the course of your development work, you may run across a situation in which you will be required to use PHP and Java together, or this combination may prove advantageous for some reason. You basically have two options to accomplish this tricky undertaking, which are outlined below. Java environments are inherently complicated and vary according to servlet engine and server. Since these issues are well beyond the scope of this book, in all further discussions we will assume that you already have a working web server that supports servlets, an installed Java virtual machine (JVM), and a working familiarity with Java.

The Java SAPI

The most stable solution is to integrate PHP into a Java servlet environment using the Java Service Access Point Identifier (SAPI). This allows the PHP processor to run as a servlet and builds on the PHP Java extension (described following). The servlet will run from within a Java servlet engine, such as Apache Tomcat.

Installation and setup

As with the Java extension, SAPI module support is not built into PHP by default. You will need to rebuild PHP with the necessary options (`-with-servlet -with java`) as well as any options you may require for other uses. (See Chapter 3 for more information on building and installing PHP.) In your environment variables, make sure that `servlet.jar` is included in your `CLASSPATH`, and add the PHP directory containing the `libphp6.so` file to `LD_LIBRARY_PATH`.

Windows users will need to build the `php_java.dll` file and copy it into their `extension_dir` directory and enable the extension in the `php.ini` file. Also be sure that `servlet.jar` is included in your `CLASSPATH`, and add the PHP directory containing the PHP DLL files to `PATH`.

Building the module will also create a JAR file called `phpsrvlt.jar`, which must also be included in your `CLASSPATH`. Additional setup specific to your servlet engine will probably be required. Check the PHP web site and mailing lists for comments from other users who have successfully configured your Java servlet engine for use with this module.

Further information

Once you have the module up and running, you should be able to view PHP files normally. Point to an existing PHP page, or create a test by printing `phpinfo()` to see if you have succeeded.

Usage of this module is still considered experimental, and there isn't a lot of documentation. You probably won't build or use this module unless you already have a specific need for it. As the module is under constant revision, additional useful notes may be found in the `README` file and other sources located in your PHP directory under `/sapi/servlet`. Users also sometimes add comments to the User Contributed Notes in the online manual, so it's probably a good idea to check there every so often. Official notes on Java/PHP integration can be found on the PHP web site at `www.php.net/manual/en/ref.java.php`.

The Java extension

If you're feeling particularly adventurous, you can build Java support directly into PHP, using the experimental Java extension. Once enabled, the extension allows you to create and call Java objects and methods from within PHP. The advantages are obvious to those familiar with Java, but use of this extension is not without some pain upfront and the need for some serious care on your part.

CAUTION **The Java extension for PHP is subject to continuing revision as it is fine-tuned for future versions. Committing to the use of this feature in your application implies added diligence to avoid future code breakage. It's not labeled EXPERIMENTAL in the PHP manual for nothing!**

Installation and setup

Use of the Java extension will require some modifications to your PHP installation and environment. Before rebuilding PHP, it's a good idea to make sure that you have access to pertinent information on your Java Development Kit (JDK) environment. Make sure that you know the following information:

- The base directory of your JDK installation (typically `/usr/java/j2sdk<version>` in Linux)

- The `JAVA_HOME` and `CLASSPATH` environment variables (`JAVA_HOME` should be set to the above directory)

- Location of the Java library (typically in `JAVA_HOME/jre/lib/i386` on Linux)

Java support is not enabled by default, and you will need to rebuild PHP in order to take advantage of its features. During the installation, you must specify the option `-with-java=(base direc-tory)` in addition to any other options you may require. (See Chapter 3 for more information on building and installing PHP.)

Modifications are also required to the `php.ini` configuration file in order to enable the extension. Open your `php.ini` file in your favorite editor and search for the `[Java]` subheading under `Module Settings`. Here's where the information you collected earlier will come in handy. A typical modification on a Linux server might look something like this:

```
[java]
java.home = /usr/java/j2sdk1.4.0
java.library = /usr/java/j2sdk1.4.0/jre/lib/i386/libjava.so
java.library.path = /usr/lib/php/extensions/no-debug-non-zts-
20020429
extension_dir = /usr/lib/php/extensions/no-debug-non-zts-
20020429
extension=libphp_java.so
```

Windows users should also add the following line under `Windows Extensions`:

```
extension=php_java.dll
```

> **NOTE** Just to add some confusion for fun, note that the `java.library` **variable pertains to the Java installation, while** `java.library.path` **refers to the PHP extension directory where the optional PHP support files reside. Most problems with getting the Java extension to work seem to revolve around successfully editing** `php.ini`**.**

Remember that these variables must correspond to the settings on your particular server.

> **CAUTION** **Windows users must be sure to enclose the path names in quotation marks.**

If all goes well, you should be ready to try calling a Java method from within PHP!

Testing

A simple invocation of `java.lang.System` in a JSP environment would look something like this:

```
<%
  String version = System.getProperty("java.version");
  String os = System.getProperty("os.name");
%>
<HTML>
<HEAD>
  <TITLE>Fun with Java and JSP</TITLE>
</HEAD>
<BODY>
<H3>We are running Java version <%= version %> on the
  <%= os %> platform, and it's working!</H3>
</BODY>
</HTML>
```

Similar code using PHP would, by necessity, be a bit more involved. Create a new PHP file called `javatest.php` and insert the following:

```
<?php
  $system = new Java('java.lang.System');
  $version = $system->getProperty('java.version');
  $os = $system->getProperty('os.name');
?>
<HTML>
<HEAD>
  <TITLE>Fun with Java and JSP</TITLE>
</HEAD>
<BODY>
<H3>We are running Java version <?php echo $version ?> on the
<?php echo $os ?> platform, and it's working!</H3>
</BODY>
</HTML>
```

With luck, your browser will output something like the following when you access `javatest.php`:

```
We are running Java version 1.4.0 on the Linux platform, and
it's working!
```

If not, it's time to troubleshoot! Consult the PHP manual or other resources listed in Appendix D for help and suggestions.

The Java object

The `Java` object becomes available with installation of the Java extension and is used to instantiate a Java class within PHP. The format is:

```
new Java(class, parameters)
```

625

where class is the class being invoked and the parameters are arguments to be passed to that object's constructor. Parameters are optional, providing that a default constructor is available.

NOTE   It's important to note that no Java packages are available to PHP by default. Although, as in the previous example, the java.lang.* package is always available to Java and JSP and therefore does not need to be referenced implicitly, the complete package and class name must always be specified from within PHP.

The previous example was a simple one, since we provided no arguments to the System class. System cannot be instantiated in Java and is referenced through static methods just as getProperty() is. Let's take a look at a more involved example.

With the deprecation of several Date() constructors, time reporting and formatting grew in complexity in Java. Here's an example in which we print the current date and time in JSP:

```
<%
Calendar cr = new Calendar();

String date_time = "yyyy-MM-dd HH:mm:ss";

java.text.SimpleDateFormat date =
  new java.text.SimpleDateFormat(date_time);

String current = date.format(cr.getTime()));
%>
<HTML>
<HEAD>
  <TITLE>Got the time?</TITLE>
</HEAD>
<BODY>
<H3>The current date and time is: <%= current %>.</H3>
</BODY>
</HTML>
```

Once again, as written in PHP:

```
<?php

  $cr = new Java('java.lang.Calendar');

  $date_time = "yyyy-MM-dd HH:mm:ss";

  $date = new Java('java.text.SimpleDateFormat',$date_time);

  $current = date->format($cr->getTime()));
?>

<HTML>
<HEAD>
```

```
   <TITLE>Got the time?</TITLE>
</HEAD>
<BODY>
<H3>The current date and time is: <?php echo $current ?>.</H3>
</BODY>
</HTML>
```

Errors and exceptions

Because Java is being accessed through PHP, a Java `Exception` appears as a PHP warning in the browser. While a reference to a class that is misspelled or not in the `classpath` might generate an error such as this within Java (accompanied by a lovely stack trace):

```
/var/tomcat4/work/webroot/_/test/test$jsp.java:57: Class
org.apache.jsp.SomethingAmiss not found.
```

when referenced from PHP, it will simply be displayed in the browser:

```
Warning: java.lang.ClassNotFoundException
```

While this warning will at least notify you that there is a problem, it doesn't provide much useful troubleshooting information. You can suppress the PHP warnings by using an @ prefix with your method calls, although we recommend against it:

```
@$trouble = $output->println();
```

Another method for suppressing errors is to use `display_errors` and set the value to 0 to cause errors not to display. This is accomplished within your program with the following statement:

```
ini_set('display_errors', 0);
```

When an exception is thrown, it's also possible to obtain the `Exception` object from Java for more pertinent information. To accomplish this, the PHP Java extension provides two functions to retrieve the last `Exception` and then to clear it: `java_last_exception_get()` and `java_last_exception_clear()`. Neither function accepts parameters.

With version 5, PHP gained an `Exception` object of its own! You can use both the Java and PHP objects in conjunction to provide more helpful error information, as in Java:

```
// check for a thrown exception in Java
$exception = java_last_exception_get();
if ($exception) {
  $ex_msg = $exception->getMessage();
  // use the Java exception to throw an exception in PHP
  throw new Exception($ex_msg);
  //clear this Java exception
  java_last_exception_clear();
}
```

627

The `getMessage()` **method will provide more information than given in the warning but still might not be enough. Use** `toString()` **if your exceptions are not providing enough useful information.**

By using both the @ prefix and these handy PHP functions, it becomes possible to exert more control over Java errors within PHP.

Potential gotchas

Expect to run into problems while trying to integrate Java and PHP. Judging from comments on PHP mailing lists and on various development boards, even seasoned professionals find they must experiment with both installation and implementation in order to achieve what they want. There are a few problems that seem to crop up most often, and you can learn from the experiences of others.

Installation problems

Assuming a preexisting servlet engine and working environment, most problems in getting the Java extension to work properly seem to begin and end in `php.ini`. Your particular servlet engine and/or platform may require more configuration. Again, check user notes online and in mailing lists and experiment on your own.

It's the classpath, stupid

A common error, especially for those not all that familiar with Java, is to neglect to include relevant packages, libraries, or JAR files within the `classpath` specified as an environment variable. If you receive a dread `ClassNotDefined` error or something similar, check the `CLASSPATH` first. If not, then perhaps you misspelled the class name. Hey, it happens.

Here comes that loose typing again

PHP may not care what type your variable is, but Java certainly does. It's probably a good idea (and good form, when calling Java methods) to typecast your PHP variables before passing them to a Java object:

```
$value = (double) $value;
$sum = (int) $sum;
$name = (String) $name;
```

Get used to typecasting variables for use with Java. As a worst-case scenario, the code may generate errors. At best, it may behave unpredictably.

Speed

Excessive referencing of Java objects can sacrifice some of the performance that PHP aficionados have grown so fond of. Use Java objects and methods only when necessary!

The sky's the limit

Obviously, the complexity only increases when you begin to create more involved scripts. There's a lot of uncharted territory that you can choose to explore. Many creative uses of the Java extension continue to be uncovered as time passes. You can even use the `java.awt.*` packages to create graphical interfaces through PHP, though much of this is limited to CGI mode. If you can dream it up in Java, there's a chance you just might be able to get it to work through PHP as well.

Here we enter the realm of experimentation. While managing to exploit the Java extension to its full potential may prove enjoyable and interesting for the programmer, it often doesn't make for a very stable or efficient application. Use it in a production environment at your own risk, and keep abreast of any changes posted on the PHP site. Meanwhile, load up a development machine and go to town!

Summary

While many similarities in syntax and object models exist between PHP and Java, differences abound in typing, compilation, and methodology.

Java programmers who are working with PHP for the first time will find server-side scripting much more intuitive if they have experience using Java Server Pages (JSP). PHP fulfills a similar function and can be embedded within HTML.

Users seeking to integrate Java and PHP have two options: the Java SAPI module and the Java extension. Both are optional, and PHP must be rebuilt to support these options. Use of Java assumes a web server with installed JVM and a servlet engine such as Apache Jakarta Tomcat. Modifications must be made to environment variables as well as the `php.ini` configuration file, and are specific to your particular platform and Java servlet engine.

Objects and methods are called from PHP, instantiating the `Java` object. Java packages are not directly available to PHP and must be correctly referenced. Parameters are optional if a default constructor is to be used.

Errors generated by Java are reported as PHP warnings but can be suppressed by attaching the @ prefix to your PHP statements, although we strongly recommend against suppressing errors. Java `Exception` objects can be accessed through PHP using built-in functions, and they can be used in conjunction with the PHP `Exception` object.

Java support in PHP is experimental and, as such, is subject to change in future releases. Those wishing to integrate Java and PHP should keep track of new releases and potential changes that could break their code.

Chapter 39

Integrating PHP and JavaScript

I n this chapter, we try to get the best of both client-side and server-side scripting by combining PHP with JavaScript. We briefly touch on the question of when to use which scripting language, and stylistic points that may be helpful when writing the code. Then we move on to pragmatic examples of the type you might see on a real-world PHP site.

> **TIP**   If you've never worked with JavaScript before, you won't learn how just by reading this chapter. We will only touch on aspects of JavaScript that materially impinge upon PHP. If you're wondering what an `onBlur` **event is, we recommend Danny Goodman's superlative** *JavaScript Bible, Sixth Edition* **(Wiley, 2007) or Steve Suehring's** *JavaScript Step by Step* **(Microsoft, 2008).**

Outputting JavaScript with PHP

Because PHP is server-side and JavaScript is client-side, you may expect to have problems using both on the same page. In actuality, it's this difference that makes them such a good match.

Although PHP offers plenty of power for creating dynamically generated web pages, it is strictly a server-side language. There's a common category of web site tasks that perhaps don't require all the processing power of a server and would best be done quickly — for instance, changing the look of a button on mouseover. JavaScript, a purely client-side language (there's a server-side version, but we're assuming that you've already chosen PHP on that end), can be easily integrated into PHP to fill in many of these gaps.

On the other hand, client-side JavaScript (aka Javascript, JScript, ECMAScript) itself has many limitations. For example, because it can't

communicate directly with a database, JavaScript cannot update itself with fresh data, depending on the page. Even worse, it's impossible to depend on client-side technologies, because they may not be present in a visitor's browser or may be disabled. Conscientious client-side web developers must either decide to code probabilistically (and accept complaints from minority-browser users) or maintain several versions of a site at the same time. (Nonconscientious developers simply adapt themselves to the market-leading browser's full capabilities and damn the torpedoes . . . but that's another story.) PHP can help to mitigate the effects of client-side indeterminacy.

Dueling objects

Perhaps the biggest divergence between JavaScript and PHP is in the area of object models. The two are quite divergent conceptually, and they use different styles of notation. Some people consider this a good thing, because at least there's no chance of mixing up objects that look similar (as there is with, say, ASP and JavaScript). Probably just as many consider it a pain, an incompatibility, or a design flaw. In any case, there is no chance that you can access the same object with both PHP and JavaScript — so forget it.

JavaScript is consistently object-oriented from top to bottom. Every statement requires an object and a method or function to be specified and may also have event handlers. JavaScript uses the so-called *dot* object notation (`object.method`), which is similar to that of other common programming languages such as Java, Python, and Microsoft's VBScript.

The downside is that JavaScript's document object model has been shakily standardized: Although, in theory, ECMA and the W3C shepherd the international standard, in practice the various browser manufacturers violate/add to this core at their whim. Proficient JavaScripters spend a good deal of energy keeping track of incompatibilities and workarounds for various browsers and platforms.

PHP doesn't care what it outputs

The main thing to keep in mind is that PHP doesn't know or care what it returns. You can (and people do) use PHP to write out plain text, HTML, XHTML, DHTML, JavaScript, XML, MathML, various graphical formats, CSS, XSL, or even (for the ironic ironists among us) ASP. No real technical barrier exists to having PHP output C code, although it's probably not a usage whose popularity is going to sweep the nation. Remember, PHP does not always output PHP — its ultimate end product is usually code that will be run by another application, usually a browser.

There are a couple of ways to write out the JavaScript with PHP. The simplest is to escape from PHP whenever you get the urge to go client-side. This is accomplished in precisely the same way you would escape from HTML.

```php
<?php
echo("Imagine tons of complex PHP code in this block.");
?>
<script language="JavaScript">
<!-- Hide from JavaScript disabled browsers
document.write("Strict separation of client-side and server-side
```

```
code is a good thing.")
// end hiding -->
</script>
<?php
echo("More PHP in this block.");
?>
```

Even this example doesn't show the fullest extent of PHP/JavaScript separation. A lot of JavaScript is actually defined within the `<HEAD>` element of an HTML page and simply called in the `<BODY>`, whereas PHP is generally used in the latter.

As with HTML, there are occasions when you don't want to escape from PHP — or this style may just be your personal preference. In that case, you can use PHP's `echo` or `print` statements to output JavaScript.

```
<?php
echo("This is some complex PHP code.");
echo("<script language=\"JavaScript\">\n");
echo("<!-- Hide from JavaScript disabled browsers\n");
echo("document.write(\"Strict separation of client-side and
server-side code is a good thing.\\n\")");
echo("// end hiding -->\n");
echo("</script>\n");
echo("More PHP in this block.");
?>
```

CAUTION You may run into trouble if you use script tags (for instance, `<script language="PHP">`) to delineate PHP chunks — the PHP parser may have a hard time figuring out which `</script>` tag goes with what `<script>` tag. Whenever possible, use the canonical `<?php ?>` tag.

This style is not at all incorrect, but it can be considerably harder to keep everything straight. Unless you're an experienced programmer, you might want to limit this style to occasions in which you simply call predefined JavaScript functions, such as `onSubmit` events.

TIP Remember to escape double quotation marks in JavaScript sections if using `echo`/`print` to output code. See line 3 of the preceding snippet.

Where to use JavaScript

Client-side JavaScript doesn't do heavy lifting, but it is faster at certain tasks and also allows for some effects that you can't easily duplicate with PHP. Some places you should definitely consider replacing or enhancing PHP with JavaScript include:

■ Simple arithmetic in forms and calculators (such as a shopping-cart running total or a mortgage calculator)

■ Simple form validation (such as making sure that e-mail addresses have @ symbols)

■ Site navigation (such as pull-down navigation menus)

- ■   Pop-ups (alerts, search boxes)
- ■   Browser events (mouseover, onClick)

PHP as a Backup for JavaScript

The flip side of our *where to use JavaScript* advice is that PHP can help caulk the cracks in JavaScript. Sometimes you can seamlessly implement both client-side and server-side methods of doing a task. If a visitor's browser is JavaScript-enabled, fine — visitors will be able to take advantage of the zippier method, generally without even noticing that they've had a choice. If not, at least you won't suffer the ignominy of totally locking them out of your site's functionality.

A perfect example is the double-barreled pull-down menu for site navigation. JavaScript gives you an instant redirect, whereas PHP provides the same result after a longer wait for those without JavaScript-enabled browsers. This trick takes advantage of the fact that JavaScript has event handlers (for example, `onChange`) that work off the structure of HTML forms without requiring an actual button-clicking submission. Therefore, the Submit button can be reserved for PHP's use. Listing 39-1 shows an HTML page that uses a JavaScript `onChange` redirect and, if that doesn't work, a PHP form handler.

LISTING 39-1

A JavaScript and PHP navigation form (navigation.html)

```
<html>
<head>
<title>Navigation pulldown</title>
<script language="JavaScript">
<!--
function Browse(form, i){
  var site = form.elements[i].selectedIndex;
  if(site > 0){
    top.location = form.elements[i].options[site].value
  }
}
// -->
</script>
</head>

<body>
<form method="post" action="redirect.php">
<select name="category" onChange="Browse(this.form,0)">
<option selected value=0>Choose a Category</option>
<option value="desktop.php">Desktops</option>
<option value="laptop.php">Laptops</option>
```

```
<option value="monitor.php">Monitors</option>
<option value="input.php">Input devices</option>
<option value="storage.php">Storage devices</option>
</select>
<input type="submit">
</form>
</body>
</html>
```

The PHP form handler file, called `redirect.php`, need only have two lines:

```
<?php $category = $_POST['category'];
header("Location: $category"); ?>
```

You could use a similar division of labor with form validation. If JavaScript is enabled, you can use it to make sure that zip codes have nine digits, phone numbers have ten digits, and e-mail addresses have both an @ and a .. If JavaScript is not enabled, you can write a little PHP script that will do the same things when the form is submitted and return the form with warnings if the values are bad.

CAUTION JavaScript form validation should be relied on only for quick convenience reminders, never for data sanitization.

Another kind of form is basically arithmetic — a shopping cart with running totals or a mortgage payment calculator. Again, you can combine both JavaScript and PHP in an arithmetic form to cover all the bases.

Finally, there is one use where PHP is so much faster that you might want to replace JavaScript altogether: browser sniffing. This is done to send different versions of a file (for instance, a style sheet) to a visitor depending on which browser she's using. Server-side browser sniffing is vastly more efficient than client-side because no text is sent until the sniff has occurred. A JavaScript browser sniff can amount to hundreds of lines of JavaScript, which must be sent on every download whether the correct browser version has been detected or not. Listing 39-2 shows a very simple server-side browser sniff, though these can be faked by the client's browser quite easily.

LISTING 39-2

A server-side browser sniff (browsersniff.php)

```
<?php
if (strpos($HTTP_USER_AGENT, 'MSIE') > 0) {
  header("Location: index_ie.html");
} elseif (strpos($HTTP_USER_AGENT, 'Gecko') > 0) {
  header("Location: index_moz.html");
}
?>
```

Static versus Dynamic JavaScript

Although the static JavaScript-PHP form in Listing 39-1 is handy for many applications, there's one big problem with it: You have to maintain it by hand. Every time you decide to add a new page to your site, you'll have to remember to manually add it to the drop-down list. Big deal, you're thinking — but these are the little things that become time-sucking nightmares when you're running a huge and high-traffic site.

With PHP and a database, you can update some of your JavaScript automatically when new data is stored in the database — or, as we might say, *dynamically.* You want to take this option whenever possible, as it will help you save time in the long run. Listing 39-3 is how you'd rewrite the form in Listing 39-1 for even better client/server integration.

LISTING 39-3

Dynamic JavaScript and PHP form (dyn_navigation.html)

```
<html>
<head>
<title>Navigation pulldown</title>
<script language="JavaScript">
<!--
function Browse(form, i){
  var site = form.elements[i].selectedIndex;
  if(site > 0){
    top.location = form.elements[i].options[site].value
  }
}
// -->
</script>
</head>

<body>
<form method="post" action="redirect.php">
<select name="category" onChange="Browse(this.form,0)">
<option selected value=0>Choose a Category</option>
<?php
mysql_connect("localhost", "user", "password");
mysql_select_db("site_db");
$query = "SELECT filename, my_text
          FROM categories
          WHERE display = 1";
$result = mysql_query($query);
while (list($filename, $my_text) = mysql_fetch_array($result)) {
  print("<option value=\"$filename\"> $my_text</option>\n");
}
?>
</select>
```

```
<input type="submit">
</form>
</body>
</html>
```

You will doubtless have realized by now that a similar technique would be valuable even if you were making a straight JavaScript form (such as by using the onSubmit event handler rather than onChange). It would enable you to make a basic JavaScript function more flexible by allowing PHP to change variable values within the function before it was sent to the browser. So feel free to use PHP to output straight JavaScript using variables from a data source, if you like.

Dynamically generated forms

You can usefully extend this train of programming thought even further by setting up a series of dynamic drop-downs that change according to previous form inputs. PHP will fetch all the data from a database and load it into the HTML page, whereas JavaScript will decide which data set should be visible under various conditions.

In this example, we want to help users find information on specific cars. The list of the model of every car made by every manufacturer is dauntingly large, too long for even a well-designed drop-down list. Furthermore, car names tend to be eerily similar to each other, like the first names of a large family of sisters in a Swedish farming village — Integra, Sentra, Jetta, Elantra, Sephia, and so forth. So one way to narrow things down logically is to have the user pick a manufacturer from a pull-down menu, which would narrow the list to only models made by that company.

The database table we need looks like this (actually it probably doesn't if you're using a relational database — but here we want to focus on the JavaScript part, not the database part):

```
 -----------------------
| make     | model      |
 -----------------------
| Audi     | A4         |
| Audi     | A6         |
| Audi     | A8         |
| Audi     | Quattro    |
| Chrysler | Cirrus     |
| Chrysler | Concorde   |
| Chrysler | PT Cruiser|
| Toyota   | Camry      |
| Toyota   | Corolla    |
| Toyota   | Rav4       |
 -----------------------
```

Using this database table and server-side PHP scripts, you would be limited to two suboptimal choices. You could opt for one extremely large list (either drop-down or full page) of manufacturers and models, or you could make the visitor go through two sequential forms. But after we add

JavaScript to the mix, we can start a page with two drop-downs and have the contents of the second list change, depending on what is selected in the first.

Our double drop-down design is based on Andrew King's very clever JavaScript code, available at www.webreference.com under the GNU General Public License. We used PHP simply to connect to the database and fetch data to populate the two-dimensional arrays from which the JavaScript works. All of the interesting functionality here is provided by the JavaScript portion (see Listing 39-4).

LISTING 39-4

A two-dimensional dynamic drop-down (double_drop.html)

```html
<HTML>
<HEAD>
<META NAME="save" CONTENT="history">
<STYLE>
  .saveHistory {behavior:url(#default#savehistory);}
</STYLE>

<SCRIPT LANGUAGE="JavaScript">
<!--
var v=false;
//-->
</SCRIPT>

<SCRIPT LANGUAGE="JavaScript1.1">
<!--
if (typeof(Option)+"" != "undefined") v=true;
//-->
</SCRIPT>

<SCRIPT LANGUAGE="JavaScript">
<!--
// Universal Related Select Menus - cascading popdown menus
// by Andrew King. v1.34 19990720
// Copyright (c) 1999 internet.com LLC. All Rights Reserved.
// Modified by Joyce Park 20000703
//
// This program is free software; you can redistribute it
// and/or modify it under the terms of the GNU General Public
// License as published by the Free Software Foundation; either
// version 2 of the License, or (at your option) any later
// version.
//
// This program is distributed in the hope that it will be
// useful, but WITHOUT ANY WARRANTY; without even the implied
// warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR
```

```
// PURPOSE.  See the GNU General Public License for more
// details.
//
// You should have received a copy of the GNU General Public
// License along with this program; if not, write to the Free
// Software Foundation, Inc., 59 Temple Place, Suite 330,
// Boston, MA  02111-1307  USA
//
// Originally published and documented at www.webreference.com
// see www.webreference.com/dev/menus/intro2.html for changelog

if(v){a=new Array(22);}

function getFormNum (formName) {
  var formNum =-1;
  for (i=0;i<document.forms.length;i++){
    var tempForm = document.forms[i];
    if (formName == tempForm) {
      formNum = i;
      break;
    }
  }
  return formNum;
}

function jmp(form, elt) {
// The first parameter is a reference to the form.
  if (form != null) {
    with (form.elements[elt]) {
      if (0 <= selectedIndex)
        var location = options[selectedIndex].value;
    }
  }
}

var catsIndex = -1;
var itemsIndex;

if (v) { // ns 2 fix
function newCat(){
  catsIndex++;
  a[catsIndex] = new Array();
  itemsIndex = 0;
}

// Andrew chose to name this function "O", presumably standing
// for "Option".  It's not a zero, here or in the array below!
function O(txt,url) {
  a[catsIndex][itemsIndex]=new myOptions(txt,url);
```

```
    itemsIndex++;
  }

  function myOptions(text,value){
    this.text = text;
    this.value = value;
  }

  // fill array
  <?php
  mysql_connect("localhost", "db_user");
  mysql_select_db("auto_db");
  // Get the makes
  $i = 0;
  $make_query = "SELECT DISTINCT make FROM cars";
  $make_result = mysql_query($make_query);
  while ($make_row = mysql_fetch_array($make_result)) {

    $make[$i] = $make_row[0];
    // Now fill the array with models for each make
    echo "newCat();\n";
    $model_query = "SELECT model
                    FROM cars
                    WHERE make = '$make[$i]'
                    ORDER BY model";
    $model_result = mysql_query($model_query);
    while(list($model) = mysql_fetch_array($model_result)) {
      echo "O(\"$model\", \"/$model.php\")\n";
    }
    echo "\n";
    $i++;
  }
  ?>
  } // close if (v)

  function relate(formName,elementNum,j) {
  if(v){
  var formNum = getFormNum(formName);
    if (formNum>=0) {
      formNum++; // reference next form, assume it follows in HTML
      with (document.forms[formNum].elements[elementNum]) {
        for(i=options.length-1;i>0;i--) options[i] = null;
        // null out in reverse order (bug workarnd)
        for(i=0;i<a[j].length;i++){
          options[i] = new Option(a[j][i].text,a[j][i].value);
        }
        options[0].selected = true;
      }
    }
  } else {
```

640

```
  jmp(formName,elementNum);
}
}

// BACK BUTTON FIX for ie4+- or
// MEMORY-CACHE-STORING-ONLY-INDEX-AND-NOT-CONTENT
// see www.webreference.com for full comments
function IEsetup(){
  if(!document.all) return;
  IE5 = navigator.appVersion.indexOf("5.")!=-1;
  if(!IE5) {
    for (i=0;i<document.forms.length;i++) {
      document.forms[i].reset();
    }
  }
}

window.onload = IEsetup;

//-->
</SCRIPT>
</HEAD>
<BODY BGCOLOR="#ffffff">

<CENTER>
<TABLE BGCOLOR="#DDCCFF" BORDER="0" CELLPADDING="8"
CELLSPACING="0">
<TR VALIGN="TOP">
<TD>Choose a make:<BR>
<FORM NAME="f1" METHOD="POST" ACTION="redirect.php"
onSubmit="return false;">
<SELECT NAME="m1" ID="m1" CLASS="saveHistory"
onChange="relate(this.form,0,this.selectedIndex)">
<?php
while (list($key, $val) = each($make)) {
  echo "<OPTION VALUE=\"/{$val}.php\">{$val}</OPTION>\n";
}
?>
</SELECT>
<INPUT TYPE=SUBMIT VALUE="Go" onClick="jmp(this.form,0);">
</FORM>
</TD>

<TD BGCOLOR="#FFFFFF" VALIGN=MIDDLE><B>---&gt;</B></TD>

<TD>Choose a model:<BR>
<FORM NAME="f2" METHOD="POST" ACTION="redirect.php"
onSubmit="return false;">
<SELECT NAME="m2" ID="m2" CLASS="saveHistory"
onChange="jmp(this.form,0)">
```

```
// These are placeholder values for the first time the page is
// loaded.  They will not change when the form values change.
// If you delete them, the forms will still work, but the
// second select menu would come up empty until changed.
// These values could be generated dynamically, but we wanted
// to show them in place.
<OPTION VALUE="/A4.php">A4</OPTION>
<OPTION VALUE="/A6.php">A6</OPTION>
<OPTION VALUE="/A8.php">A8</OPTION>
<OPTION VALUE="/Quattro">Quattro</OPTION>
</SELECT>
<INPUT TYPE=SUBMIT VALUE="Go" onClick="jmp(this.form,0);">
<INPUT TYPE="hidden" NAME="baseurl" VALUE="http://localhost">
</FORM>
</TD>
</TR>
</TABLE></CENTER>

</BODY>
</HTML>
```

If you were to add to or change any of the data in the database, the JavaScript would change automatically. Dynamic integration of new data makes this a very powerful tool and keeps page maintenance to a minimum.

Passing data back to PHP from JavaScript

Finally, we close the data loop by passing form values back to PHP with JavaScript. Listings 39-5, 39-6, and 39-7 use JavaScript to force at least one check box to be checked at all times. In addition, it passes an array to a PHP script. We've chosen to use frames here to maximize the speed of the changes, and we wrote all values out by hand for clarity, rather than assembling them dynamically from a data source.

LISTING 39-5

Frameset (sandwich_frames.html)

```
<HTML>
<HEAD>
<FRAMESET ROWS="50%, 50%" FRAMEBORDER="no" BORDER=0>
<FRAME SRC="sandwichorder.html" NAME="main" SCROLLING="auto">
<FRAME SRC="results.php" NAME="results" SCROLLING="auto">
</FRAMESET>
</HEAD>
<BODY></BODY>
</HTML>
```

Form page (sandwichorder.html)

```
<HTML>
<HEAD>
<SCRIPT LANGUAGE="JavaScript">
<!--

function deselectAllOthers(boxVals) {
  for (var x = 1; x < boxVals.length; x++) {
    boxVals[x].checked=false;
  }
}

function confirmOne(boxVals) {
  var count = 0;
  for (var x = 1; x < boxVals.length; x++) {
    if (boxVals[x].checked == false) {
      count++;
    }
  }
  if (count == (boxVals.length--1)) {
    boxVals[0].checked = true;
  } else {
    boxVals[0].checked = false;
  }
}

function toArray(boxVals) {
  for (var x = 0; x < boxVals.length; x++) {
    var valArray = boxVals[x].name+"[]";
    boxVals[x].name = valArray;
  }
}

// -->
</script>
</HEAD>

<BODY BGCOLOR=#FCFCF0 onLoad="document.selector.submit();">

<TABLE CELLPADDING=20>
<TR>
<TD VALIGN="top">
<B>Order a sandwich with...</B>
<BR><BR>
<FORM NAME="selector" TARGET="results" METHOD="post"
ACTION="results.php">
<B>Fillings (check one or more)</B><BR><BR>
<INPUT TYPE="checkbox" name="filling" value="everything"
```

```
checked onClick="deselectAllOthers(document.selector.filling);
confirmOne(document.selector.filling);
toArray(document.selector.filling); submit();"> Everything
<BR>
<INPUT TYPE="checkbox" name="filling" value="turkey"
onClick="confirmOne(document.selector.filling);
toArray(document.selector.filling); submit();"> Turkey
<BR>
<INPUT TYPE="checkbox" name="filling" value="roastbeef"
onClick="confirmOne(document.selector.filling);
toArray(document.selector.filling); submit();"> Roast beef
<BR>
<INPUT TYPE="checkbox" name="filling" value="pastrami"
onClick="confirmOne(document.selector.filling);
toArray(document.selector.filling); submit();"> Pastrami
<BR>
<INPUT TYPE="checkbox" name="filling" value="eggplant"
onClick="confirmOne(document.selector.filling);
toArray(document.selector.filling); submit();"> Eggplant
<BR>
</TD>
<TD VALIGN="top"><BR><BR>
<B>Cheese</B><BR><BR>
<SELECT NAME="cheese" onChange="submit();">
<OPTION VALUE="none">None</OPTION>
<OPTION VALUE="cheddar">Cheddar</OPTION>
<OPTION VALUE="swiss">Swiss</OPTION>
<OPTION VALUE="camembert">Camembert</OPTION>
<OPTION VALUE="bleu">Blue</OPTION>
<OPTION VALUE="cottage">Cottage</OPTION>
</SELECT>
<BR><BR>
</TD>
<TD VALIGN="top"><BR><BR>
<B>Bread</B><BR><BR>
<SELECT NAME="bread" onChange="submit();">
<OPTION VALUE="white">White</OPTION>
<OPTION VALUE="wheat">Wheat</OPTION>
<OPTION VALUE="rye">Rye</OPTION>
<OPTION VALUE="kaiser">Kaiser roll</OPTION>
<OPTION VALUE="onion">Onion roll</OPTION>
<OPTION VALUE="dutch">Dutch crunch</OPTION>
</SELECT>
</FORM>
<BR><BR>
</TD>
</TR></TABLE>
</BODY>
</HTML>
```

LISTING 39-7

Results listing (results.php)

```
<HTML>
<HEAD></HEAD>

<BODY BGCOLOR=#666680 TEXT=#ffffff>
<TABLE CELLPADDING=30><TR><TD>
<B>Da Results</B><BR><BR>
<?php
$filling = $_POST['filling'];
$cheese = $_POST['cheese'];
$bread = $_POST['bread'];

if ($filling) {
  if (is_array($filling)) {
    reset($filling);
    while (list($key, $value) = each($filling)) {
      echo("$value<BR>\n");
    }
  } else {
    echo($filling);
  }
}
?>
<BR><BR></TD>
<TD VALIGN=top><BR><BR>
<B>Cheese</B><BR><BR>
<?php echo($cheese); ?>
<BR><BR></TD>
<TD VALIGN=top><BR><BR>
<B>Bread</B><BR><BR>
<?php echo($bread); ?>
</TD></TR>
</TABLE>
</BODY>
</HTML>
```

This form admittedly doesn't actually do very much yet — but the point is that it would obviate one or two trips from client to server and back. It also demonstrates another of the interesting effects PHP developers can get by experimenting with JavaScript.

Summary

JavaScript is a client-side scripting language that is highly efficient at many tasks that do not require server-side processing. Not everyone will want to use JavaScript, which has longstanding usability and security issues, but for those who do, the combination of client-side and server-side programming languages can provide an attractive variety of functionalities.

PHP and JavaScript have different object notations. JavaScript uses the so-called *dot* notation, whereas PHP uses the *arrow* or C++ style notation. JavaScript is thoroughly object-oriented, whereas PHP treats objects as an optional feature. The good news is that you'll never confuse a JavaScript object for a PHP object, or vice versa. The bad news is that you cannot access the same object from both languages.

It's often possible to implement a feature in both a client-side and a server-side way. Users with JavaScript-enabled browsers can enjoy greater speed and convenience, whereas those without can still get the functionality. This makes it possible to consider using JavaScript without its greatest drawback, which is unpredictability leading to alienation of segments of the user base.

Perhaps the greatest service PHP can perform for JavaScript is to enable database connectivity — resulting in what we might call *Dynamic JavaScript*. JavaScript, being purely a client-side technology, cannot query a server-side database for variable data with which to dynamically generate content. Without a server-side helper like PHP, JavaScripts must be updated by hand whenever variable data is changed. PHP's capability to pass in up-to-date variables from a data store can make it considerably less labor-intensive to maintain JavaScript-enabled forms and functions.

Chapter 40

Integrating PHP and XML

XML is one of the hottest buzzwords in the software business today; but what does it mean for Joe or Jane Average PHP Developer? Well, it could very well be the necessary precondition for a better Internet — one that is faster to develop, more interactive, less junky, and more accessible to a larger audience. With PHP, you're already in an excellent position to smoothly integrate XML into your web development arsenal as the technology matures.

What Is XML?

XML stands for *Extensible Markup Language.* XML is a form of *SGML*, the *S*tandard *G*eneralized *M*arkup *L*anguage, but you don't need to know anything about SGML to use XML. It defines syntax for structured documents that both humans and machines can read.

> **NOTE** Our explanation of XML will necessarily be extremely brief (because this is a book about PHP rather than XML). For those who want to learn more, we highly recommend Elliotte Rusty Harold's *XML 1.1 Bible, Third Edition* (Wiley, 2004). Although this book is neither short nor a specific guide to programming XML-based applications, it will give you a firm conceptual grasp of XML that should set you up nicely for any particular XML-based task.

Perhaps the easiest way to understand XML is to think about all the things HTML can't do. HTML is also a markup language, but HTML documents are anything but structured. HTML tags (technically known as elements) and attributes are just simple identification markers for the browser. For instance, a pair of matched <H1> and </H1> tags designates a top-level heading. Browsers interpret this to mean you want heading text to be

displayed in a really big, bold, possibly italicized font. HTML does not, however, indicate whether the text between those tags is the title of the page, the name of the author, an invitation to enter the site, a pertinent quotation, a promise of special sale prices, or what. It's just some text that happens to be big.

One implication of HTML's lack of structure is that search engines have little built-in guidance about what's important on each page of your site or what each chunk of text means in relation to the others. They use various methods they can use to guess, none of which is foolproof. `<META>` tags are notoriously prone to abuse — porn sites often load popular but irrelevant search terms into their headers to fool unwary web surfers — and spiders can end up giving too much weight to portions of the page that designers might think are unimportant. If XML becomes ubiquitous, it could eliminate many of these problems and lead the way to much more meaningful web searching.

Let's say you work for a content web site that has just signed a major distribution deal with a top-five portal. After you wake up from the champagne hangover, you're faced with the hard question of how you plan to deliver the content. HTML isn't going to do the job: Obviously the portal's page design and web-serving technology are different from your site's, and they won't be able to just plug your HTML into theirs. Just to make things really interesting, let's assume that you and Big Portal Company use different programming languages, different data stores, different HTML editors, different style sheets — in short, different everything. The necessary bridge is a data-exchange format that is easy for you to output with your technical setup, clearly understood by both parties with their existing software, and equally easy for the Big Portal Company to convert to its own purposes and designs. XML is that data exchange format.

You could, of course, write a script to dump data from your data store into a tab-delimited file. Then you could write out the details of your custom data format and send it with the tab-delimited file to Big Portal Company. There one of its engineers would try to figure out your schema and write code to transform your data into its format. However, anyone who has actually done this knows how much fiddly work it requires, how many tests need to be performed, and how much time even the tiniest error can suck up. On the other hand, you could just output your data in XML, and the Big Portal Company engineer could write a very short script — perhaps just three functions long — to transform your XML tags to its corresponding XML tags. Then Big Portal Company could treat the data just like its own data. XML is an attempt to move toward a common language and set of methods for performing tasks like these, instead of having data exchange involve a series of custom jobs each time.

We hope these examples begin to answer the "Why XML?" question. If you forget the hype and focus on what problems XML might begin to solve, you'll be in a much better position to assess whether it can help you today or sometime in the future. In the simplest terms, XML is a flexible data exchange format that is not dependent on any particular software or domain, can be parsed easily by both machines and humans, and allows content providers to include information about the structure of the data along with the data itself.

The next question about XML is typically, "What does XML look like anyway?" Actually, XML looks a lot like HTML. A simple XML file, such as the one shown in Listing 40-1, is easy for HTML users to understand.

A simple XML file

```
<?xml version="1.0"?>
<book>
 <publisher>IDG Books</publisher>
 <title>PHP5 Bible</title>
 <chapter title="PHP and XML">
  <section title="What is XML?">
    <paragraph>
If you know HTML, you're most of the way to understanding XML.
    </paragraph>
    <paragraph>
They are both markup languages, but XML is more structured
than HTML.
    </paragraph>
  </section>
 </chapter>
</book>
```

As you can see, XML has tags and attributes and the hierarchical structure that you're used to see-ing in HTML. In XML, each pair of tags (`<paragraph></paragraph>`) is known as an *element*. Actually, this is true in HTML, too, but most people strongly prefer the term *tag* (the construction that marks an element) over *element* (the conceptual thing that is being marked by a tag) — we're not picky. Use whatever term you want as long as you know what you mean. The biggest difference is that XML tags are self-defined; they carry absolutely no display directive to the web browser or other viewing application.

XML makes the following minimal demands:

■ There must be a single root element that encloses all the other elements, similar to `<HTML></HTML>` in HTML documents. This is also sometimes called the *document element*.

■ Elements must be hierarchical. That is, `<X> <Y> </Y> </X>` is allowed, but `<X> <Y> </X> </Y>` is not. In the first example, `<X>` clearly contains all of `<Y>`. In the second example, `<X>` and `<Y>` overlap. XML does not allow overlapping tags.

■ All elements must be deliberately closed (in contrast to HTML, which allows some unclosed elements such as `<OPTION>` or `<LI>`). This can be accomplished with a closing tag (`<title></title>`) as in HTML or by using an XML feature with no HTML equiva-lent called a *self-closing element* (`<logo href="graphic.jpg"/>`). A self-closing element is also known as an *empty element*.

■ Elements can contain elements, text, and other data. If an element encloses something that looks like it might be XML — such as `<hello>` — but isn't, or if you don't want something parsed, it must be escaped.

> **CAUTION**   The &, <, >, ', and " characters are all restricted in XML. You can use them in your data by *escaping* them — using codes such as `&` and `<` — or by putting them in CDATA sections, which we discuss in the section "Documents and DTDs," later in this chapter.

In addition to these mandatory requirements for what is called *well-formedness*, the XML standard also suggests that XML documents should start with an identifying XML declaration. This is a processing instruction giving the MIME type and version number, such as `<?xml version="1.0"?>`. This is not required, but some parsers complain if it isn't present. Also, XML is case-sensitive; some variants, such as XHTML, require lowercase tags and attributes. Lowercase tags are not absolutely required by the XML standard itself, but unless you have a good reason to do otherwise you should use lowercase tags and attributes.

> **NOTE**   It's the XML declaration, and other processing instructions with the same format, that prevents you from using PHP's short tags with XML. Because the two tag styles are identical (`<? ?>`), it would be unclear whether this character sequence set off a PHP block or an XML processing instruction.

XML documents are usually text. They can contain binary data, but they aren't really meant to. If you want to put binary data in your XML documents, you have to encode it first and decode it later. Note that including binary data may break some of the platform-independence of pure XML.

Working with XML

By now you may or may not think XML is the greatest thing since cinnamon toast, but in either case you're probably asking yourself, "OK, but what can I actually *do* with it?" This is actually not such an easy question to answer. In theory, you can do three main things with XML: manipulate and store data, pass data around between software applications or between organizations, and display XML pages in a browser or other application using style sheets to apply display directives.

In practice, almost no one actually uses XML as a primary data store when SQL is so ubiquitous. It's possible, although still difficult, to manipulate data using XML — for instance, to edit documents by creating and manipulating XML nodes rather than straight text — but again many users don't see a tremendous amount of extra value to this practice. A great deal of progress has been made in displaying XML in the browser, generally in the form of XHTML, in the last couple of years, but there are still significant issues with this practice. For more information about displaying XML, see the sidebar "The Promises and Pitfalls of Displaying XML."

This leaves one main job for XML right now: exchanging data between applications and organizations. This happens to be the area in which PHP can have the most immediate impact. For instance, a C program might perform some operations on data from a data store and then output the results in XML, which PHP could transform into HTML for display in a browser or other application.

This data flow actually makes sense if substantial amounts of computation need to happen behind the scenes, because you do not want to have a big program both performing complex operations and outputting HTML if you can possibly help it.

PHP can also read in data from a data store and write XML documents itself. This can be helpful when transferring content from one web site to another, as in syndicating news stories. You can also use this functionality to help nontechnical users produce well-formed XML documents with a Web form frontend. At the moment, writing XML might well be the most common category of XML-related PHP task.

Finally, data is beginning to be manipulated and exchanged across human and nonhuman endpoints via the Internet itself. This technology is called *Web services*, and it is the subject of Chapter 41.

Documents and DTDs

As we explained earlier, the requirements for a well-formed XML document are fairly minimal. However, XML documents have another possible level of "goodness," which is called *validity*. A valid XML document is one that conforms to certain stated rules that together are known as a *document type definition* (*DTD*).

To get in the mood to understand the value of DTDs, imagine that you are the head of an open source project that exists to make books and other documents freely available in electronic form on the Internet. You're very excited about XML from the moment you learn about it because it seems to meet your need for a data exchange format that can adapt easily to new display technologies as they evolve. Your group members vote to encode all the project's books and documents in XML, and soon the XMLized documents start to pour in.

But when you look at the first couple of submissions, you get a rude shock. One of them is in the same format as Listing 40-1, earlier in this chapter, but one of them looks like what you see in Listing 40-2.

LISTING 40-2

A book in XML format

```
<?xml version="1.0"?>
<book title="PHP5 Bible">
 <publisher name="Wiley Publishing"/>
 <chapter number="40">
  <chapter_title>PHP and XML</chapter_title>
  <p>
   <sentence>If you know HTML, you're most of the way to
understanding XML.</sentence>
   <sentence>They are both markup languages, but XML is more
structured than HTML.</sentence>
  </p>
 </chapter>
</book>
```

The two XML files express similar, but not identical, hierarchical structures using similar but not identical tags. This is the potential downside of the self-defined markup tags that XML enables: random variation that makes it difficult to match up similar kinds of information across files. You quickly realize that you will need to implement some rules about what kinds of information should be in a book file and what the relationships between these elements will be. You've just realized you need a DTD.

A DTD describes the structure of a class of XML documents. A DTD is a kind of formal constraint, guaranteeing that all documents of its type will conform to stated structural rules and naming conventions. A DTD enables you to specify exactly what elements are *allowed,* how elements are *related,* what *type* each element is, and a *name* for each element. DTDs also specify what attributes are required or optional, and their default values. You could of course just write down these rules in a text file:

```
The top-level object of this document is a BOOK
A BOOK has one and only one TABLE OF CONTENTS
A BOOK has one and only one TITLE
A BOOK is composed of multiple CHAPTERS
CHAPTERS have one and only one CHAPTERTITLE
All CHAPTERTITLEs are listed in the TABLE OF CONTENTS
etc.
```

You could give a copy of the list to anyone who might need it. A DTD is just a more concise, well-defined, generally agreed upon grammar in which to do the same thing. It's a useful discipline to apply to XML documents, which can be chaotic because of their entirely self-defined nature. Furthermore, if you can get a group of people to agree on a DTD, you are well on the way to having a standard format for all information of a certain type. Many professions and industries, from mathematicians to sheet music publishers to human resources departments, are eager to develop such domain-specific information formats.

In the previous example, which uses XML to store books electronically, your group members may have to argue for months before hashing out the details of a DTD that perfectly describes the relationships between the table of contents, chapters, titles, and headings, indexes, appendices, sections, paragraphs, forwards, epilogues, and so on. You can, of course, iterate on DTDs as frequently as necessary.

But after your DTD is finalized, you can enjoy another value-add of XML. You can now run any XML document through a so-called "validating parser," which will tell you whether it's meeting all the requirements of its DTD. So instead of a human editor having to read each electronic book submission to see whether it has the required elements and attributes in the correct relationship, you can just throw them all into a parser and let it do the formal checking. This won't tell you anything about the quality of the content in the XML document, but it will tell you whether the form meets your requirements.

In order to work with XML in PHP, you need to learn about the basic structure of DTDs and the XML documents they describe whether you choose to validate or not.

The structure of a DTD

A *document type definition* is a set of rules that defines the structure of a particular group of XML documents. A DTD can be either a part of the XML document itself (in which case it is an *internal DTD*), or it can be located externally, in another file on the same server or at a publicly available URL anywhere on the Internet (in which case it is an *external DTD*).

NOTE **Although a DTD can be internal (part of the XML document itself), making it external (a separate file) is usually better. DTDs are meant to define a class of documents, so separating them from the XML saves you from editing every XML document of that class if you need to change the DTD later on. Because demonstrating on an internal DTD is easier for readers to follow in a book format, however, we use both as examples in this chapter.**

You can start by looking at a simple XML document with an internal DTD in Listing 40-3.

LISTING 40-3

An XML document with internal DTD (recipe.xml)

```
<?xml version="1.0"?>

<!DOCTYPE recipe [
<!ELEMENT recipe (ingredients, directions, servings)>
<!ATTLIST recipe name CDATA #REQUIRED>
<!ELEMENT ingredients (#PCDATA)>
<!ELEMENT directions (#PCDATA)>
<!ELEMENT servings (#PCDATA)>
]>

<recipe name ="Beef Burgundy">
 <ingredients>Beef</ingredients>
 <ingredients>Burgundy</ingredients>
 <directions>
 Add beef to burgundy. Serve.
 </directions>
 <servings>12</servings>
</recipe>
```

We've divided the XML document into three subsections for easier reading. The first section is the standard one-line XML declaration that should begin every XML document. The second section is the internal DTD, marked by lines beginning with the `<!` sequence. The third section is the XML itself, strictly speaking. For the moment, we are focusing on the second section, the DTD. In our example, the stuff outside the square brackets is a *document type declaration* (not to be confused with *document type definition*): `<!DOCTYPE recipe [...]>`. The document type declaration gives information about the DTD this document is using. Because this is an internal DTD, we simply give the name of the root element (`recipe`) and then include the rest of the definition within square brackets. If you are using an external DTD, however, you use the document type declaration

to state the type and location of the DTD. Two example document type declarations referring to external DTDs are:

```
<!DOCTYPE recipe SYSTEM "recipe.dtd">
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01 Transitional//EN"
 "http://www.w3.org/TR/html4/loose.dtd">
```

External document type declarations give a root element name, the type (`SYSTEM`, meaning on the server, or `PUBLIC`, meaning a standardized DTD) and the location where it can be found. You are doubtless familiar with document type declarations because, without exception, you always include one, like the preceding example, in every single HTML or XHTML document you write — right?

The DTD proper consists of the lines inside the square brackets. These lay out the elements, element types, and attributes contained in the XML document.

- **Element**: A start and end tag pair — for example, `<b> something </b>` — or an empty element (`<br/>`). Elements have types and sometimes content and attributes.
- **Element Type**: A constraint on the content and attributes of an element. A type can be used to specify what kind of data it can contain and to specify what attributes it can have.
- **Attribute**: A name and value pair associated with an element, in the form `<element attributename="attributevalue">`.

In the example DTD in Listing 40-3, we've declared that the root element, `recipe`, contains three child elements — `ingredients`, `directions`, and `servings` — and has one required attribute, `name`. Each child element is of the parsed character data type, and the attribute is of the character data type.

If you wanted to split up Listing 40-3 into an XML document and an external DTD, it would look much the same, except that, instead of providing the definition in square brackets, you would give a reference to the external DTD file. The result would look like Listings 40-4 and 40-5.

LISTING 40-4

An XML document with external DTD (recipe_ext.xml)

```
<?xml version="1.0"?>
<!DOCTYPE recipe SYSTEM "recipe.dtd">

<recipe name ="Beef Burgundy">
 <ingredients>Beef</ingredients>
 <ingredients>Burgundy</ingredients>
 <directions>
 Add beef to burgundy. Serve.
 </directions>
 <servings>12</servings>
</recipe>
```

LISTING 40-5

An external DTD (recipe.dtd)

```
<!ELEMENT recipe (ingredients, directions, servings)>
<!ATTLIST recipe name CDATA #REQUIRED>
<!ELEMENT ingredients (#PCDATA)>
<!ELEMENT directions (#PCDATA)>
<!ELEMENT servings (#PCDATA)>
```

Because the XML used in both examples conforms to the internal and external DTDs, both documents should be declared valid by a validating parser.

You could learn a lot more about the specifics of DTDs and XML documents, but these basics should enable you to understand most of PHP's XML functions.

Validating and nonvalidating parsers

XML parsers come in two flavors: *validating* and *nonvalidating*. Nonvalidating parsers care only that an XML document is well formed — that it obeys all the rules for closing tags, quotation marks, and so on. Validating parsers require well-formed documents as well, but they also check the XML document against a DTD. If the XML document doesn't conform to its DTD, the validating parser outputs specific error messages explaining what has gone wrong.

PHP's SAX parser, `libxml2`, is nonvalidating (as was the `expat` parser used in PHP4). That doesn't mean that you should ignore DTDs. Going through the process of creating a DTD for each of your document types is a good design practice. It forces you to think out the document structure very carefully. And if your documents ever need to go through a validating parser, you're covered. In fact, many experts recommend that you put all XML documents through a validating parser even if you never plan to use one again.

Most validating parsers are written in Java and are a pain to set up and use. The easiest way to validate your XML is to use an online validator. A well-known one is the STG validator at `www.stg.brown.edu/service/xmlvalid`.

Actually, using Gnome `libxml` to validate an XML document is possible — but it takes some work. Examples of validation using C are on the `libxml` web site (at `www.xmlsoft.org`).

SAX versus DOM

There are three common APIs for handling XML and XML documents: SimpleXML, the *Document Object Model* (DOM), and the *Simple API for XML* (SAX). All three modules are now included in all PHP distributions.

655

You can use the DOM, SAX, or SimpleXML API to parse and change an XML document. To create or extend an XML document entirely through the PHP interface (in other words, without writing any of it by hand), you must use the DOM. Each API has advantages and disadvantages:

- **SAX**: SAX is much more lightweight and easier to learn, but it basically treats XML as flowthrough string data. So if, for instance, you want to parse a recipe, you could whip up a SAX parser in PHP, which might enable you to add boldface to the ingredient list. Adding a completely new element or attribute would be very difficult, however, and even changing the value of one particular ingredient would be laborious.

- SAX is very good for repetitive tasks that can be applied to all elements of a certain type — for instance, replacing a particular element tag with HTML tags as a step toward transforming XML into HTML for display. The SAX parser passes through a document once from top to bottom — so it cannot "go back" and do things based on inputs later in the document.

- **DOM**: PHP's DOM extension reads in an XML file and creates a *walkable* object tree in memory. Starting with a document or an element of a document (called *nodes* in the DOM) you can get or set the children, parents, and text content of each part of the tree. You can save DOM objects to containers as well as write them out as text. DOM XML works best if you have a complete XML document available. If your XML is streaming in very slowly or you want to treat many different XML snippets as sections of the same document, you want to use SAX. Because the DOM extension builds a tree in memory, it can be quite the resource hog with large documents.

- **SimpleXML**: The SimpleXML API makes it easy to quickly open an XML file, convert some of the elements found there into native PHP types (variables, objects, and so on) and then operate on those native types as you would normally. The SimpleXML API saves you the hassle of making a lot of the extra calls that the SAX and DOM APIs require, uses far less memory than DOM XML, and often is the simplest way of accessing XML data quickly. There are limitations, though, including some quirky behavior related to attributes and deeply nested elements.

DOM

The Document Object Model is a complete API for creating, editing, and parsing XML documents. The DOM is a recommendation of the World Wide Web Consortium. You can read all about it in the W3's inimitable prose at `www.w3.org/DOM`.

Basically the idea is that every XML document can be viewed as a hierarchy of nodes resembling leaves on a tree. Starting with the root element, of which all other elements can be expressed as children, any program should be able to build a representation of the structure of a document. Attributes and character data can also be attached to elements. This *tree* can be read into memory from an XML file, manipulated by PHP, and written out to another XML file or stored in a container.

The parser behind the scenes in PHP's DOM extension is *gnome-libxml2* (aka Gnome *libxml2*), which is supposedly less memory-intensive than others. This is available at www.xmlsoft.org.

DOM XML is the only entirely object-oriented API in PHP, so some familiarity with object-oriented programming helps when using it. However, there are a limited number of objects and methods, so you do not need any particularly deep knowledge of object-oriented programming to use DOM XML.

Using DOM XML

How you use the DOM will depend on your goals, but these steps are common:

1. Open a new DOM XML document, or read one into memory.
2. Manipulate the document by nodes.
3. Write out the resulting XML into a string or file. This also frees the memory used by the parser.

The simple example in Listing 40-6 shows some basic DOM XML functions in use. Make sure that your server has its file permissions set in such a way that the web server can write a file.

LISTING 40-6

A simple DOM XML example (dom_example.php)

```php
<?php
$doc = new DomDocument("1.0");
$root = $doc->createElement("HTML");
$root = $doc->appendChild($root);
$body = $doc->createElement("BODY");
$body = $root->appendChild($body);
$body->setAttribute("bgcolor", "#87CEEB");
$graff = $doc->createElement("P");
$graff = $body->appendChild($graff);
$text = $doc->createTextNode("This is some text.");
$text = $graff->appendChild($text);
$doc->save("test_dom.xml");
?>
```

DOM functions

Table 40-1 lists the most common DOM functions. You must call one of these functions before you can use any of the other DOM XML functions!

DOM XML Top-Level Function Summary

Function	Behavior
`domxml_open_mem(string)`	Takes a string containing an XML document as an argument. This function parses the document and creates a Document object.
`domxml_open_file(filename)`	Takes a string containing an XML file as an argument. This function parses the file and creates a Document object.
`domxml_xmltree(string)`	Takes a string containing an XML document as an argument. Creates a tree of PHP objects and returns a DOM object. *Note:* The object tree returned by this function is read-only.
`domxml_new_doc(version)`	Creates a new, empty XML document in memory. Returns a Document object.

Table 40-2 lists the most important classes of the DOM API.

XML DOM Class Summary

Class	Behavior
`DomDocument`	This class encapsulates an XML document. It contains the root element and a DTD if any.
`DomNode`	Encapsulates a node, aka an element. A node can be the root element or any element within it. Nodes can contain other nodes, character data, and attributes.
`DomAttr`	This class encapsulates a node attribute. An *attribute* is a user-defined quality of the node.

Table 40-3 lists the most important methods of the `DomDocument` class.

DomDocument Class Summary

Method	Behavior
`createElement(name)`	Creates a new element whose tag is the passed string. You must append this element to another element using `DomNode->appendChild()`.

Method	Behavior
`createTextNode` `(character_data)`	Creates a new text node (`DomText` object). You must append this node to another node using `DomNode->appendChild()`.
`save(filename)`	Dumps XML from memory to a designated file.
`saveXML([node])`	Dumps XML from memory to a string. Optional parameter is a DomNode object.

Table 40-4 lists the most important methods of the `DomNode` class.

TABLE 40-4

DomNode Class Summary

Method	Behavior
`appendChild(newnode)`	Attaches a node to another node.
`removeChild(child)`	Removes the child node.

Table 40-5 lists the most important methods of the `DomAttr` class.

TABLE 40-5

DomAttr Class Summary

Method	Behavior
`name()`	Returns an attribute name.
`value()`	Returns the value of an attribute.

SAX

The Simple API for XML is widely used to parse XML documents. It is an event-based API, which means that the parser calls designated functions after it recognizes a certain trigger in the event stream.

SAX has an interesting history, especially in contrast to the DOM. The SAX API is not shepherded by an official standardizing body. Instead, it was hammered out by a group of programmers on the XML-DEV mailing list, many of whom had already implemented their own XML parsers (in Java first!) without a standard API. You can learn more at the web sites of SAX team members, such as `www.saxproject.org`.

SAX works from a number of *event hooks* supplied by you via PHP. As the parser goes through an XML document, it recognizes pieces of XML such as elements, character data, and external entities. Each of these is an event. If you have supplied the parser with a function to call for the particular kind of event, it pauses to call your function after it reaches that event. The parsed data associated with an event is made available to the called function. After the event-handling function finishes, the SAX parser continues through the document, calling functions on events, until it reaches the end. This process is unidirectional from beginning to end of the document — the parser cannot back up or loop.

A very simple example is an event hook that directs PHP to recognize the XML element `<paragraph></paragraph>` and substitute the HTML tags `<p></p>` around the character data. If you wrote this event hook, you could not specify a particular paragraph — instead, the function is called for every instance of this event.

The parser behind the scenes in the PHP SAX extension is `libxml2`, which you can read about on its project site at `www.xmlsoft.org`.

Prior to version 5, PHP used James Clark's `expat`, a widely used XML parser toolkit. More information about `expat` can be found on Clark's web site at `www.jclark.com/xml`. If you compile with `libxml2`, you should be able to use all your PHP4 SAX code in PHP6 without problems.

CAUTION	Unfortunately, the term *parser* can refer either to a software library such as `libxml2`, or to a block of XML-handling functions in PHP. Verbs such as *create* and *call* indicate the latter, more specific meaning. Any PHP XML function that uses the term *parser* also refers to the latter meaning.

Using SAX

How you use the SAX will depend on your goals, but these steps are common:

1. Determine what kinds of events you want to handle.

2. Write handler functions for each event. You almost certainly want to write a character data handler, plus start element and end element handlers.

3. Create a parser by using `xml_parser_create()` and then call it by using `xml_parse()`.

4. Free the memory used up by the parser by using `xml_parser_free()`.

The simple example in Listing 40-7 shows all the basic XML functions in use (see recipe.xml from earlier in the chapter for the XML file used in this example).

LISTING 40-7

A simple XML parser (simpleparser.php)

```php
<?php
$file = "recipe.xml";

// Call this at the beginning of every element
```

```
function startElement($parser, $name, $attrs) {
    print "<B>$name =></B>  ";
}

// Call this at the end of every element
function endElement($parser, $name) {
    print "\n";
}

// Call this whenever there is character data
function characterData($parser, $value) {
    print "$value<BR>";
}

// Define the parser
$simpleparser = xml_parser_create();
xml_set_element_handler($simpleparser, "startElement",
"endElement");
xml_set_character_data_handler($simpleparser, "characterData");

// Open the XML file for reading
if (!($fp = fopen($file, "r"))) {
  die("could not open XML input");
}

// Parse it
while ($data = fread($fp, filesize($file))) {
if (!xml_parse($simpleparser, $data, feof($fp))) {
  die(xml_error_string(xml_get_error_code($simpleparser)));
  }
}

// Free memory
xml_parser_free($simpleparser);
?>
```

SAX options

The XML parser in the SAX API has two configurable options: one for case folding and the other for target encoding.

Case folding is the residue of a series of past decisions and may not be relevant now that XML has been definitely declared case-sensitive. Early versions of SGML and HTML were not case-sensitive and, therefore, employed case folding (making all characters uppercase or lowercase during parsing) as a means of getting a uniform result to compare. This is how your browser knew to match up a <P> tag with a </p> tag. Case folding fell out of favor due to problems with internationalization, so after much debate XML was declared case-sensitive. When case folding is enabled, node names

passed to event handlers are turned into all uppercase characters. A node named `mynode` would be received as MYNODE. When case folding is disabled, a `<paragraph>` tag will not match a `</PARAGRAPH>` closing tag.

> **NOTE**   **Case folding is enabled by default, which violates the XML 1.0 specification. Unless you disable it by using** `xml_parser_set_option()` **as explained in a moment, your event handlers receive tags in uppercase letters.**

Event handlers receive text data from the XML parser in one of three encodings: *ISO-8859-1*, *US-ASCII*, or *UTF-8*. The default is ISO-8859-1. The encoding of text passed to event handlers is known as the *target encoding.* This is, by default, the same encoding as in the source document, which is known as the *source encoding.* You can change the target encoding if you need to process the text in an encoding other than the encoding it was stored in.

PHP and Internationalization

Computer programs store letters as integers, which they convert back to letters according to *encodings*. Early programs used English, which conveniently needs only 1 byte (actually only 7 bits) to represent all the common letters and symbols. This encoding standard was promulgated in 1968 as ASCII (*American Standard Code for Information Interchange*).

However, programmers soon found that English has an unusually small number of characters, and thus the only languages that can be expressed with any completeness in ASCII are Hawaiian, Kiswahili, Latin, and American English. Ever since then, programmers concerned with internationalization have tried to promote encoding standards that promise to assign a unique integer to every one of the letters of every one of the world's alphabetical languages. The result of this effort is referred to as *Unicode*.

The three encodings supported by PHP's XML extension are *ISO-8859-1*, *US-ASCII*, and *UTF-8*. US-ASCII is the simplest of these, a slight renaming of the original 7-bit ASCII set. *ISO-8859-1* is also known as the Latin1, Western, or Western European encoding. It can represent almost all western European languages adequately. *UTF-8* allows the use of up to 4 bytes to represent as many of the world's languages as possible. If your XML document is written in Han-gul or Zulu, you have no choice but to use UTF-8.

Encoding options are retrieved and set with the functions `xml_parser_get_option()` and `xml_parser_set_option()`. Case folding is controlled by using the constant `XML_OPTION_CASE_FOLDING`, and target encoding by using the constant `XML_OPTION_TARGET_ENCODING`.

In the following example, we create an XML parser that reads in data as ASCII, turns off case folding, and spits out the output as UTF-8.

```
$new_parser = xml_parser_create('US-ASCII');
$case_folding = xml_parser_get_option(XML_OPTION_CASE_FOLDING);
echo $case_folding;
```

```
$change_folding = xml_parser_set_option($new_parser, XML_OPTION_CASE_
FOLDING,0);

$target_encoding = xml_parser_get_option(XML_TARGET_ENCODING);
echo $target_encoding;
$change_encoding = xml_parser_set_option($new_parser, XML_OPTION_
TARGET_ENCODING, 'UTF-8');
```

SAX functions

Table 40-6 lists the most important SAX functions, with descriptions of what they do.

TABLE 40-6

XML SAX Function Summary

Function	Behavior
xml_parser_create([encoding])	This function creates a new XML parser instance. You may have several distinct parsers at any time. The return value is an XML parser or false on failure. Takes one optional argument, a character-encoding identifier (such as UTF-8). If no encoding is supplied, ISO-8859-1 is assumed.
xml_parser_free(parser)	Frees the memory associated with a parser created by xml_parser_create().
xml_parse(parser, data[, *final*])	This function starts the XML parser. Its arguments are a parser created by using xml_parser_create(), a string containing XML, and an optional finality flag. The finality flag indicates that this is the last piece of data handled by this parser.
xml_get_error_code(parser)	If the parser has encountered a problem, its parse fails. Call this function to find out the error code.
xml_error_string(errorcode)	Given an error code returned by xml_get_error_code(), it returns a string containing a description of the error suitable for logging.
xml_set_element_handler(parser, start_element_handler, *end_element_handler*)	This function actually sets two handlers, which are simply functions. The first is a start-of-element handler, which has access to the name of the element and an associative array of its elements. The second is an end-of-element handler, at which time the element is fully parsed.
xml_set_character_data_handler(parser, cd_handler)	Sets the handler function to call whenever character data is encountered. The handler function takes a string containing the character data as an argument.
xml_set_default_handler(parser, handler)	Sets the default handler. If no handler is specified for an event, the default handler is called if it is specified. Takes as arguments the parser and a string containing unhandled data, such as a notation declaration or an external entity reference.

SimpleXML API

The SimpleXML API was new in PHP5. Characterized as an object-mapping API, SimpleXML dispenses with web standards and absolute flexibility in favor of simplicity and modest memory usage. If you just need to read some data from an XML document and write some other data back in, the SimpleXML likely will require the fewest lines of code of all possible approaches to the problem.

Here's the idea behind SimpleXML: As in the DOM approach, SimpleXML parses an XML document and holds the whole thing in memory. However, rather than hold the document as a DOM object (which you must further manipulate before you can use its contents), its elements are stored as native PHP variables and so are immediately usable. Because many DOM tasks do not actually require you to traverse all the children and parents of a document, but rather perform repetitive tasks on well-defined nodes, SimpleXML ultimately constitutes a PHP-specific compromise between the SAX and DOM approaches.

Using SimpleXML

When using SimpleXML, you read a passage of XML text — either a string or a file — into a variable with the function `simplexml_load_string()` or `simplexml_load_file()`. You then have a local object you can refer to directly. Listing 40-8 shows how the SimpleXML API can be used to get variable values out of an XML file with just a few lines of code.

Listing 40-8 demonstrates a typical use of SimpleXML.

LISTING 40-8

SimpleXML sample (simplexml.php)

```php
<?php

$recipe = simplexml_load_file("recipe.xml");

$ingredients = $recipe->ingredients;
$directions  = $recipe->directions;
$servings    = $recipe->servings;

foreach ($ingredients as $ingredient)
{
print "<P>Ingredient: $ingredient";
}

print "<P>Directions: $directions";
print "<P>Serves $servings";

?>
```

SimpleXML functions

Table 40-7 lists the most important SimpleXML functions, with descriptions of what they do.

TABLE 40-7

SimpleXML Function Summary

`simplexml_load_file(file)`	Import and parse a file.
`simplexml_load_string(string)`	Import and parse a string.
`simplexml_import_dom(DomDocument)`	This function allows you to convert a DomDocument object into a SimpleXML object, and then treat it just like an imported XML file or string.

A Sample XML Application

This series of scripts will write out XML to a file by using data from an HTML form and then will allow you to edit the values in that file.

Listing 40-9 is an HTML form that can be used by nontechnical users to define forms. (They don't care that this data will be formatted and stored in XML.) Listing 40-10 is a script to write out the XML file.

LISTING 40-9

A form to collect values for an XML file (pollform.php)

```
<HTML>
<HEAD>
<TITLE>Make-a-poll</TITLE>
</HEAD>

<BODY>
<CENTER><H3>Make-a-poll</H3></CENTER>

<P>Use this form to define a poll:</P>
<FORM METHOD="post" ACTION="writepoll.php">

<P>Give this poll a <B>short</B> name, like <FONT COLOR="red">Color Poll</
FONT>.<BR>
<INPUT TYPE=TEXT NAME="PollName" SIZE=30>
```

```
</P>

<P>This poll should <B>begin</B> on this date (MM/DD/YYYY):
<INPUT TYPE=TEXT Name="Poll_Startdate" SIZE=10>
</P>

<P>This poll should <B>end</B> on this date (MM/DD/YYYY):
<INPUT TYPE=TEXT NAME="Poll_Enddate" SIZE=10>
</P>

<P>This is the poll question (<FONT COLOR="blue">e.g. Why did the chicken
cross the road?</FONT>):
<INPUT TYPE=TEXT NAME="Poll_Question", size=100>
</P>

<P>These are the potential answer choices you want to offer (<FONT
COLOR="darkgreen">e.g. Yes, No, Say what?</FONT>).  Fill in only as many as
you need.  Keep in mind that brevity is the soul of good poll-making.<BR>
<INPUT TYPE=TEXT NAME="Raw_Poll_Option[]" SIZE=25><BR>
<INPUT TYPE=TEXT NAME="Raw_Poll_Option[]" SIZE=25><BR>
<INPUT TYPE=TEXT NAME="Raw_Poll_Option[]" SIZE=25><BR>
<INPUT TYPE=TEXT NAME="Raw_Poll_Option[]" SIZE=25><BR>
<INPUT TYPE=TEXT NAME="Raw_Poll_Option[]" SIZE=25><BR>
<INPUT TYPE=TEXT NAME="Raw_Poll_Option[]" SIZE=25><BR>
</P>

<INPUT TYPE="submit" NAME="Submit" VALUE="Add a poll">
</FORM>

</BODY>
</HTML>
```

LISTING 40-10

A script to write out an XML file (writepoll.php)

```
<html>
<head>
<title>Write an XML file</title>
</head>

<body>
<?php

$pollfile = "poll.xml";
```

```php
// Reading in the xml file as a string
$fd = fopen($pollfile, "r") or die("Can't open file.");
$fstr = fread($fd, filesize($pollfile)) or die("Can't read file, check
permissions.");
fclose($fd);

// Format response sets.
$PollName = str_replace("\'", "", $_POST["PollName"]);
$PollName = str_replace(" ", "_", $_POST["PollName"]);

$RespSet = "";

for ($r=0; $r<=5; $r++) {
    $currentRawPollOption = $_POST["Raw_Poll_Option"][$r];
  if (!empty($_POST["Raw_Poll_Option"][$r])) {
      $Poll_Option[$r] = "$_POST[PollName]-".str_replace("'", "",
$currentRawPollOption);
    $Poll_Option[$r] = "$_POST[PollName]-".str_replace(" ", "_",
$currentRawPollOption);
    $currentPollOption = $Poll_Option[$r];

    $RespSet .= "\t<response id=\"$currentPollOption\">$currentRawPollOption</
response>\n";
  }


}

//Add new poll data
$separator = "</Poll>";
$divide = explode($separator, $fstr);
$glue =
"\t<Poll name=\"$_POST[PollName]\"/>
</PollList>

<Poll id=\"$_POST[PollName]\">
\t<StartDate>$_POST[Poll_Startdate]</StartDate>
\t<EndDate>$_POST[Poll_Enddate]</EndDate>
\t<name>$_POST[PollName]</name>
\t<text>$_POST[Poll_Question]</text>
\t<display type=\"Bar_Graph\"/>
\t<responseSet resource=\"$PollName-responseSet\"/>
</Poll>

<responseSet id=\"$PollName-responseSet\">
$RespSet</responseSet>
";
```

```
$newxml = implode($glue, $divide);

//Write to file
$fd = fopen($pollfile, "w") or die("Can't open file for writing; check file
permissions");
$writestr = fwrite($fd, $newxml);

//Message
echo "Wrote $writestr chars to $pollfile.";
?>

</body>
</html>
```

Listing 40-11 shows the XML file where our polls are stored, with one poll already defined for you.
If you add a new poll, it will be appended near the top of this file, and its name will be added to the
PollList.

LISTING 40-11

An XML file (poll.xml)

```
<?xml version="1.0"?>
<PollDefs>
<Poll id="Best_Text_Editor">
  <StartDate>01/01/2003</StartDate>
  <EndDate>01/31/2004</EndDate>
  <question>Which is the best programmer's editor?</question>
  <display type="Bar)Graph"/>
  <responseSet>
    <response id="Best_Text_Editor-emacs">emacs</response>
    <response id="Best_Text_Editor-vim">vim</response>
    <response id="Best_Text_Editor-notepad">notepad</response>
    <response id="Best_Text_Editor-kate">kate</response>
    <response resource="Best_Text_Editor-BBEdit">BBEdit</response>
  </responseSet>
</Poll>

<Poll id="Best_Pointer_Device">
  <StartDate>02/01/2004</StartDate>
  <EndDate>02/29/2004</EndDate>
  <question>Which is the best pointer device?</question>
  <display type="Bar_Graph"/>
  <responseSet>
    <response id="Best_Pointer-mouse">Mouse</response>
    <response id="Best_Pointer-trackball">Trackball</response>
    <response id="Best_Pointer-touchpad">Touchpad</response>
```

```
    <response id="Best_Pointer-trackpoint">TrackPoint</response>
    <response id="Best_Pointer-pen">Pen</response>
    <response id="Best_Pointer-stylus">Stylus</response>
  </responseSet>
</Poll>

</PollDefs>
```

Listing 40-12 shows a script that will allow you to edit the XML file in Listing 40-10 using DOM XML.

LISTING 40-12

XML editor (dom_polledit.php)

```php
<html>
<head>
<title>Poll XML editor</title>
</head>

<body>
<?php

$doc = new DomDocument();
$pollfile = "poll.xml";


// Handle form submission
if ($_POST['stage'] == 1) {
  // Reading in the XML file as a DOM object
  if (!$doc->load($pollfile)) {
    echo "Cannot read XML file.";
    exit;
  }

  // Once a poll is created, the user will only be able to
  // change the StartDate, EndDate, Question, and response values.

  // Format the data
  $pollname = $_POST['poll_name'];
  $startdate = $_POST['Poll_Startdate'];
  $enddate = $_POST['Poll_Enddate'];
  $question = $_POST['Poll_Question'];

  // Replace the values as text nodes
  $poll_list = $doc->getElementsByTagname("Poll");
  foreach ($poll_list as $poll_obj) {
```

669

```
// Figure out which poll we're editing, then work on its children
$pollname_value = $poll_obj->getAttribute("id");
if ($pollname_value == $pollname) {
  $children = $poll_obj->childNodes;
  foreach ($children as $child_obj) {
    $node_name = $child_obj->nodeName;
    $value = $child_obj->nodeValue;
    if ($node_name == "StartDate") {
      if ($value == $startdate) {
        // Do nothing
      } else {
        $sd_textnode = $child_obj->firstChild;
        $new_startdate = $doc->createTextNode($startdate);
        $child_obj->replaceChild($new_startdate, $sd_textnode);
      }
    }
    if ($node_name == "EndDate") {
      if ($value == $enddate) {
        // Do nothing
      } else {
        $ed_textnode = $child_obj->firstChild;
        $new_enddate = $doc->createTextNode($enddate);
        $child_obj->replaceChild($new_enddate, $ed_textnode);
      }
    }
    if ($node_name == "question") {
      if ($value == $enddate) {
        // Do nothing
      } else {
        $q_textnode = $child_obj->firstChild;
        $new_question = $doc->createTextNode($question);
        $child_obj->replaceChild($new_question, $q_textnode);
      }
    }
    if ($node_name == "responseSet") {
      $old_responses = $child_obj->childNodes;
      $i=0;
      foreach ($old_responses as $delete_responses) {
        if ($delete_responses->nodeName == 'response') {
          $r_textnode = $delete_responses->firstChild;
          $new_response = $doc->createTextNode($_POST['response'][$i]);
          $delete_responses->replaceChild($new_response, $r_textnode);
          $i++;
        }
      }
    }
  }
}
```

```
  }

  // Write out the file
  $doc->save($pollfile);

}

// This stuff happens every time, whether a submission
// has occurred or not.

// Reading in the XML file as a DOM object
// Must read fresh every time
if (!$doc->load($pollfile)) {
  echo "Cannot read XML file.";
  exit;
}

// Get a list of the polls in this XML document
// and then pull out the start date, end date,
// poll question, and possible responses.
$poll_list = $doc->getElementsByTagname('Poll');
foreach ($poll_list as $poll_obj) {
  $id = $poll_obj->getAttribute("id");
  $children = $poll_obj->childNodes;
  foreach ($children as $key=>$child_obj) {
    $node_name = $child_obj->nodeName;
    if ($node_name != "#text" && $node_name != 'responseSet') {
      $content_str = $child_obj->nodeValue;
      $poll_array["$node_name"] = $content_str;
    } elseif ($node_name == 'responseSet') {
      // Get the responses
      $responselist = $child_obj->childNodes;
      foreach ($responselist as $responses) {
        $response_name = $responses->nodeName;
        if ($response_name != "#text") {
          $response_array[] = $responses->nodeValue;
        }
      }
    }
  }

  // Arrange all the data nicely
  $poll_startdate = $poll_array['StartDate'];
  $poll_enddate = $poll_array['EndDate'];
  $poll_name = $poll_array['name'];
  $poll_question = $poll_array['question'];
  $poll_question = stripslashes($poll_question);
```

671

```
  foreach ($response_array as $key=>$val) {
    $resp_str .= "Option:  <INPUT TYPE=\"text\" SIZE=25 NAME=\"response[$key]\"
VALUE=\"$val\"><BR>\n";
  }

  // Display form with old values
  $php_self = $_SERVER['PHP_SELF'];
$form = <<< EOFORM
<FORM METHOD="post" ACTION="$php_self">
Start Date:  <INPUT TYPE="text" SIZE=10 NAME="Poll_Startdate" VALUE="$poll_
startdate"><BR>
End date:  <INPUT TYPE="text" SIZE=10 NAME="Poll_Enddate" VALUE="$poll_
enddate"><BR>
Poll question:  <INPUT TYPE="text" SIZE=100 NAME="Poll_Question" VALUE="$poll_
question"><BR>
$resp_str
<INPUT TYPE="hidden" NAME="poll_name" VALUE="$id">
<INPUT TYPE="hidden" NAME="stage" VALUE=1><BR>
<INPUT TYPE="submit" VALUE="Presto-chango">
</FORM>


EOFORM;
  echo $form;
  unset($resp_str);
  unset($response_array);
}
?>

</body>
</html>
```

Gotchas and Troubleshooting

The DOM and SAX parsers will only parse a well-formed XML document. If the parser rejects your XML, make sure that it is well formed. If it looks good to your eye, run it through a different validating parser or an online XML checker, such as the one at www.xml.com/xml/pub/tools/ruwf/ check.html.

If you cannot read and write XML documents to disk, check that the web server process has permission to do so.

If the DOM API returns a fatal *function not found* error, the DOM XML module may not be installed. Use the phpinfo() function to check for a domxml entry. If it isn't there, you will have to recompile PHP with the DOM XML module (on Unix) or uncomment the php_domxml.dll line in php.ini (on Windows).

Summary

XML is an application-independent data exchange format that promises to make web development faster and easier in the future. XML and HTML are both descended from SGML, accounting for their close resemblance at first glance. Both have tags (more correctly called elements) and attributes, although XMLs are self-defined and structured whereas HTMLs are defined by the HTML standard and contain no information about document structure.

XML has only a few minimal requirements for well-formedness. These include closed elements, no overlapping elements, escaped special characters, and the presence of a single root element for each document. XML can also be valid, however, in the sense of conforming to a formal declaration of its structure in a document type definition or DTD. DTDs can be internal or external to the XML document and even located on another server. They contain declarations of the types, attributes, and names of the various elements within the XML file.

For the present, few prefabricated tools are available to help you write, edit, and display XML. You can use one of the three PHP XML APIs — SAX, DOM, and SimpleXML — to write your own tools. The APIs have different tradeoffs and uses. SAX is an event-based parser, whereas DOM XML creates an object tree in memory. SimpleXML is easy to use and requires little code, but it is relatively limited in its capability. It's mainly useful for quick reads of simple XML files.

At the moment, PHP with the SAX extension can be used to write out well-formed XML from values entered into a Web form and to edit XML documents. DOM XML can be used to create complete XML documents programmatically. The SAX parser is also commonly employed to transform XML into HTML for less problematic display in current web browsers. Another possible task for PHP's XML extensions is to pull data from a data store and write it out as XML for exchange with another organization.

Chapter 41

Creating and Consuming Web Services with PHP

Web services are an emerging field of programming that seeks to apply the benefits of the web to bigger problems than merely displaying data in a browser. PHP, which has already proven itself as a core *glue* component of the web, has the opportunity to grab even more market share in the Web services arena. As is true of other hot technologies such as XML, however, a world of hype surrounds Web services. Here we try to cut through the buzzwords and analyst predictions to look at what Web services mean to the average PHP developer.

The End of Programming as We Know It

The title of this section is a bit of a joke — one of us works in the Web services field and often hears presentations that assert things such as: "In 10 years, we will have no more need for programmers, because Web services will eliminate duplication of effort." Many people have thought that programming was about to die out, and all of them have been wrong so far — but hope springs eternal in the pundit's breast. Notching down the hyperbole to manageable levels, we can say that Web services could make some common but hard tasks in commercial computer programming a lot easier.

The ugly truth about data movement

Joking aside, Web services do solve some problems — at the moment largely in the realm of moving data around. Later in this chapter, for example, we offer code for a client to the Amazon REST service. This code enables you to

grab the latest data about a given product or group of products — photos, current prices, availability, and so on — up to once per second via an automatic process.

If you're a first-time author who has a small informational site with one link to Amazon, this isn't really going to help you much. But there are Amazon Associates who link to thousands or even millions of products. They did so until recently by horrible hacks involving downloading all those web pages and using some kind of string or XPATH parsing to pick out the three or four pieces of data they wanted from each page. Furthermore, each client organization did all this work for itself — because, among other things, this is a totally unauthorized use of Amazon's copyrighted material, so they can hardly expect Amazon to help them. Harvesting data from full HTML pages is tremendously wasteful and expensive for both Amazon and the Associate — so much so that it's a good way to get banned from Amazon altogether. Slamming the door on requests from a particular IP block is almost the only way to control access to a public web server.

Even if an organization wants to give you large amounts of information in a data feed, the mechanics right now are not very elegant. We are aware of many large and well-respected data-related businesses that move data around in text files (or spreadsheets) that are downloaded via some mechanism such as FTP (or e-mail) and parsed on both ends by custom Perl code (or by hand). Often, there is no way to send only data that has changed — the feeds are dumped out and processed in a dumb way every so often, rather than updating only if and when changes occur. Obviously, these are all batch processes, which have no possibility of working in real time. XML-based Web services promise to offer a common language, a common transport mechanism, a common authentication and authorization method, and potentially common code for organizations to access each other's data.

If Web services were just about moving data around, the idea would be extremely useful but not at all sexy. What excites everyone about Web services is the promise that they can help solve the hardest problems of distributed computing once and for all.

Brutal simplicity

Think back, if you can, to the bad old days before the web. If you can go back far enough, think back to the days when the Internet itself was a rarity and networking something limited to high-end universities (it may help to remember that for a long time one of Apple's selling points in college computer labs was AppleTalk).

Back in those dark days, my children, things such as operating systems and programming languages were major barriers to integration — they were little islands in the sea of incompatibility. If you were going to write an application, you were specifically writing it for a particular platform and language — sometimes even for a particular version of a compiler. It was very, very hard to make one program talk to another program. If you wrote a COBOL program on a VAX, that was where it was going to stay. With a great deal of effort you could get one program to send something simple, such as ASCII data, to another — but any little thing could mess up your interapp communication. If you changed anything on one side, it might mean that you had to change a bunch of stuff on the other side, too. These programs were said to be *tightly coupled*.

This meant a lot of duplication of effort. Porting was technically difficult, and the market was fragmented. So a team that wrote an application — say, an accounting program — for Minicomputer X was not necessarily going to have the resources to do the same for Microcomputer Y. Lots of teams wrote lots of accounting programs, and all the formats were proprietary. None of them could exchange data with each other, much less share tasks easily.

Slowly, mankind groped toward a way to make programs talk to each other. The blanket term for this activity was *distributed computing*. It took until the mid-1990s for these methods to reach the common programmer, in the form of standards such as DCOM, CORBA, and Java RMI. These standards enabled all programs that shared a common architecture to call each other's methods and send data back and forth. When you are able to embed a spreadsheet inside a word processor document, it's via the magic of DCOM.

These common object models, however, had three major problems. They were still more or less tied to particular platforms or programming languages; they were considered difficult to learn, and they reached general usability at the same moment that the web arrived to tantalize us with the possibility of Internet-scale loosely coupled, distributed computing based on open standards.

The web is the biggest, most open, most loosely coupled — and most successful — distributed architecture of all time. With few exceptions, no web server cares which web browser is asking it for a page, or what operating system that browser is running on, or what chip is running the hardware on which that operating system lives. The application asking for the page doesn't even need to be a browser — it may be a spider from a search engine, it may be an `fopen()` call from a PHP command-line script, or it may be a cellular phone. The HTML it sends may not render nicely on every device, but that has nothing to do with whether Apache or IIS is serving up the page.

Applying the lessons of the web to applications, you come up with something quite a lot like Web services. The beauty of Web services is its brutal simplicity, which squashes everything down to a lowest common denominator. A Web services architecture doesn't care about the benefits of any particular platform, and it doesn't care about the pitfalls. Those are your problems. All that matters to the outside world is that a program can send and receive text messages across HTTP or SMTP and that these text messages can trigger computational actions.

An archetypal Web service would be something like a Japanese-to-English translation service. It lives somewhere on the Internet, on some unknown platform, and is written in some unknown programming language. You don't need to know or care about that stuff. All you care about is that your browser or your mail client knows that you only read English — so every time you get a web page or an e-mail in Japanese, these applications automatically send their contents to this translation service and then display the translated results to you. You don't ever see or care that part of the processing is happening at some remote location — to you, the end user, it just looks as though your application is handling it seamlessly. Instead of your web browser using Babelfish to translate Japanese to English and your mail client having a little built-in dictionary and your local department store's inventory management system using a third-party program that runs only on Solaris — all of them can just call this translation Web service.

Web services should also enable much easier integration. Say that you work for a university alumni office that is still using an alumni database written in COBOL on a VAX. (You may laugh, but Y2K wouldn't have been such a big deal if there weren't so many legacy systems lying around.) It works perfectly well, and you don't have budget to replace it — but those VAX terminals are getting old. It would sure be great if you could query your alumni database via an ordinary web browser — but there's no way you're ever going to be able to squeeze a full web server onto that old VAX, even if someone wrote one. With Web services, if you can get the VAX to understand just a little bit of XML and spit out its data as XML, you're all set. You can exchange instructions and data via XML by using some other machine that does have a web server, and that other machine can communicate with the rest of the world. Someday when you are ready to replace that VAX with a newer machine and a different programming language, no one need ever know. As long as the service is reachable at the same address by using the same method invocations, it doesn't matter whether it's a VAX or a PC, whether the application was written in COBOL or whether it's just a thin shell of PHP on top of a database.

Integration between businesses particularly benefits the smaller parties involved. Web services are easy to implement because corporate firewalls already have holes punched through them for HTTP and SMTP and because they can be implemented by using inexpensive software such as PHP. Say that you run a small business that makes widgets. You want your widgets distributed by a large retailer, Humongous Widget Depot. Until recently, for you (and the gazillions of other manufacturers who supply goods to Humongous Widget Depot) to provide real-time inventory information to the retailer entailed tremendous expense as you bought a large software package such as SAP and integrated it on a private network. Now, in theory, each small manufacturer can merely expose its inventory information via a Web service, and Humongous Widget Depot's humongous IT department merely points a Web services client at them.

That, in a nutshell, is the dream and the promise of Web services. We are quite a way from the actuality, but the outlines of a solution are firming up.

REST, XML-RPC, SOAP, .NET

For Web services to work, every application and many servers need to speak a common language. Everyone agrees that the common language is XML, but there are some philosophical differences about the implementation details. The three main Web services standards are *REST*, *XML-RPC*, and *SOAP*. One of the biggest backers of SOAP is Microsoft, which uses that standard heavily in its .NET services architecture.

REST

REST is an acronym for *RE*presentational *S*tate *T*ransfer. The concept is based on a dissertation by Roy Fielding, and its main point is that we already have everything we need to implement Web services — in HTTP itself. For all intents and purposes, a REST service is just an XML page on the web, although usually not one that is intended to be read by a human being using a browser.

REST is particularly valuable for content-focused services. You can build an XML document on the fly, and your users can access it reliably as a URI. In theory, REST should also be easier for lightly technical users to deal with. On the other hand, REST doesn't have built-in support for complex types — because there's no shared *vocabulary*, there's no particular way to designate an array versus a string.

You can learn more about REST at the RESTWiki:

```
http://rest.blueoxen.net/cgi-bin/wiki.pl
```

XML-RPC refers to a spec for making remote procedure calls over HTTP by using XML encoding. An XML-RPC server takes an input that consists of a simple XML encoding of a method call sent as an HTTP POST. An example is:

```
POST /xmlrpc-epi/xmlrpc-php-epi/sample/server.php HTTP/1.0
User-Agent: xmlrpc-epi-php/0.2 (PHP)
Host: localhost:80
Content-Type: text/xml
Content-Length: 191

<?xml version='1.0' encoding='iso-8859-1' ?>
<methodCall>
<methodName>greeting</methodName>
<params>
  <param>
    <value>
      <string>World</string>
    </value>
  </param>
</params>
</methodCall>
```

Assume that greeting() is a function that takes a string input and returns a string output consisting of the string "Hello, " prepended to the input string. It returns a response that is formatted in a similar way. An example is:

```
<?xml version='1.0' encoding='iso-8859-1' ?>
<methodResponse>
<params>
  <param>
    <value>
      <array>
        <data>
          <value>
            <string>Hello, World</string>
          </value>
        </data>
      </array>
    </value>
```

```
      </param>
    </params>
  </methodCall>
```

Notice that unlike using REST, you are not simply asking for data back — you are calling a specific function on another machine using specified types. This particular function happens to simply return data, but that is entirely arbitrary — any method the server owner is willing to expose as a Web service is fair game. Also unlike REST, XML-RPC supports all PHP native types, except objects and resources, and also a few that PHP doesn't have (structs, date-time, base-64 binary).

XML-RPC can be seen as a compromise between the complexity of SOAP and the simplicity of REST. It is so similar to SOAP, however, that it may simply be absorbed wholesale into the more vendor-friendly concept. As you will see, the PHP XML-RPC server can also deliver SOAP responses.

Learn more about XML-RPC at `www.xmlrpc.org`.

SOAP

SOAP may or may not stand for Simple Object Access Protocol — some members of the committee dispute this — but so many people have said it now that it's become true through usage. SOAP is a proposal of the W3C. SOAP, like XML-RPC, sends messages in XML wrappers with a fairly strict vocabulary that makes extensive use of namespaces. A very simple SOAP request may look like this:

```
POST /xmlrpc-epi/xmlrpc-php-epi/sample/server.php HTTP/1.0
User-Agent: xmlrpc-epi-php/0.2 (PHP)
Host: localhost:80
Content-Type: text/xml
Content-Length: 530

<?xml version='1.0' encoding=''iso-8859-1' ?>
<SOAP-ENV:Envelope
  xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:xsi="http://www.w3.org/1999/XMLSchema-instance"
  xmlns:xsd="http://www.w3.org/1999/XMLSchema">
<SOAP-ENV:Header>
...
</SOAP-ENV:Header>
<SOAP-ENV:Body>
  <greeting>
    <xsd:string>World</xsd:string>
  </greeting>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

The response may be something like the following:

```
<?xml version='1.0' encoding='iso-8859-1' ?>
<SOAP-ENV:Envelope
```

```
     xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
     xmlns:xsi="http://www.w3.org/1999/XMLSchema-instance"
     xmlns:xsd="http://www.w3.org/1999/XMLSchema">
<SOAP-ENV:Header>
...
</SOAP-ENV:Header>
<SOAP-ENV:Body>
   <greetingResponse>
     <SOAP-ENC:Array SOAP-ENC:arrayType="xsd:string[1]">
       <xsd:string>Hello, World</xsd:string>
     </SOAP-ENC:Array>
   </greetingResponse>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

SOAP offers you even more data types than XML-RPC. You must, however, specify a lot more, too. Other than that, the two protocols are very similar. Obviously, SOAP also enjoys greater acceptance from Big Software.

Learn more about SOAP at `www.soapware.org/bdg`.

Current Issues with Web Services

By now, you're probably thinking, "Okay, if Web services are so great, why aren't we using them everywhere?" Well, there are still many issues to be worked out. Web services are in their infancy, and it is likely to be years before we live in a totally Web-serviced world.

Large Footprint

Web services can be rather data-intensive and heavy in terms of wrappings and protocol, even for simple calls. This drives binary programmers especially crazy, accustomed as they are to apps talking to each other in compact binary formats. To a large extent, people just need to get over this, but there are still many situations where data storage, memory, and bandwidth are issues — in cellular phones, for instance.

Potentially heavy load

So far, there is no standard way to cache the results of RPC calls. Even if 80 percent of your clients are asking for the exact same response and, therefore, you can't save resources — every request must handled *de novo*.

REST enables caching via all the methods by which HTML can be cached.

681

Standards

Before Web services can really take off, applications need to handle their results transparently. Because the Web services standards are still somewhat in flux, and there are multiple candidates with competing strengths, this has not yet happened. For smallish web applications, it's not that big a deal if some service changes one of its API methods — but for a big app like Lotus Notes, it's a major investment of resources to transparently deal with SOAP.

The companies that are leading the way in public Web services — Amazon, Salesforce.com, eBay, and Google — have so far used a mixture of Web service APIs. Many of them maintain multiple interfaces for developer convenience; Amazon, for instance, offers all its Web services via both REST and SOAP. While this is extremely developer-friendly of them, in the long run most organizations long for a single, stable standard to conform to.

Hide and seek

The ultimate goal of Web services is to have the application transparently find all the resources it needs. Say that you get an e-mail in Japanese — your mail server or client should be smart enough to find the translation service it needs, get your document translated, and show the final result to you.

To accomplish this, we need some kind of directory system and a standard way for servers to describe themselves. WSDL and UDDI are the technologies that can make this possible. *WSDL* (*Web Services Description Language*) describes a Web service interface, while *UDDI* (*Universal Description, Discovery, and Integration*) is a registry for Web services. Learn more about UDDI and WSDL, respectively, at `http://uddi.xml.org` and `www.w3.org/TR/wsdl`.

We are a long way, however, from automatic discovery and communication by applications. In fact, many businesses that deploy Web services deliberately do so under a veil — for example, FedEx, which needed to take down its public SOAP server after it was used for fraud by crackers. Web services are growing most quickly in the realm of semiprivate transactions — companies set up Web services that are only meant to be accessed by authenticated and authorized business partners.

Who pays and how?

Ultimately, the biggest question about public Web services is: Who pays for them, and how? So far most of the Web services that you can access are things such as weblog entries and simple currency calculators — services for which you normally would not expect to pay. Big Software's answer is to create huge private networks of Web services, similar to those of Hailstorm or the Liberty Alliance. This has serious implications for privacy and open architecture.

Unless and until Web services find a way to pay for themselves, they are likely to continue to be deployed mostly inside corporate firewalls and in nonprofit situations.

Project: A REST Client

Listing 41-1 is a basic client script for Amazon's elegantly simple REST service, which has been available (with some changes) to Amazon Associates and other developers since spring of 2002. You feed the script a search string at the top, and it outputs a CSS-formatted box at the end containing information about the current edition of the book in question.

This service clearly demonstrates the biggest advantage of REST: You can work with it by using the HTTP concepts — and the PHP functions — you're already familiar with. For all intents and purposes, you are simply asking for a web page by using http fopen(). It happens to be well-formed XML instead of HTML — but that is incidental to the transport mechanism.

We chose to parse the XML by using PHP's DOM XML extension, which so far has found relatively few real-world uses (see Chapter 40 for discussion of the DOM). Many other PHP-literate Amazon developers have produced scripts that use other types of parsing, such as string parsing and regex, to extract the desired information — but we want to show you the power of using XML itself.

We should warn you, however, that this type of solution does not scale — DOM XML is a notorious memory hog. (We've heard credible reports that a 1,000-line XML document read into the DOM results in 1MB of memory being appropriated.) However, the Amazon Web services interface will only return a few items at a time, so DOM XML is an appropriate technology for this purpose.

CAUTION   The DOM extension changed significantly in PHP5. This script will not work at all in versions of PHP before 5.0.0b2. Obviously, the script will also not work unless you have previously compiled PHP with the `--with-domxml` flag and libxml2.

LISTING 41-1

Client for Amazon RESTservice (rest_amazon_client.php)

```php
<?php

// Get the xml
$writefile = 'mysqlbible.xml';
$file = "http://ecs.amazonaws.com/onca/xml?Service=AWSECommerceService&AWSAcce
ssKeyId=XXXXXXXXXXXXXXXX&Operation=ItemLookup&IdType=ASIN&ItemId=0764549324&Ve
rsion=2008-04-07&ResponseGroup=Medium,OfferFull";

$fp = fopen($file, "r");
$xml_array = file($file);
fclose($fp);
$xml_str = implode("", $xml_array);
$fp2 = fopen($writefile, "w");
$fw_return = fwrite($fp2, $xml_str);
fclose($fp2);

$xml = file_get_contents($writefile);
```

```
$loaded_xml = simplexml_load_string($xml);

// Load up the xml file into memory
$dom = new DomDocument;
if (!$dom->load($writefile)) {
  echo "Cannot load XML file";
  exit;
}
// Get an immediately available edition
$asin = $loaded_xml->Items->Item->ASIN;
$url = $loaded_xml->Items->Item->DetailPageURL;
$title = $loaded_xml->Items->Item->ItemAttributes->Title;
$imageurl = $loaded_xml->Items->Item->MediumImage->URL;

// Format a nice box
$box_str = <<< EONICEBOX
<HTML>
<HEAD>
<STYLE>
#content {
  float: left;
  padding: 10px;
  margin: 10px;
  background: #FFFFFF;
  border: 4px solid #008000;
  width: 200px; /* ie5win fudge begins */
  voice-family: "\"}\"";
  voice-family:inherit;
  width: 200px;
}
html>body #content {
  width: 170px; /* ie5win fudge ends */
  }
p {
  font-family: Verdana, Arial, sans-serif;
  font-size: 12px;
  line-height: 22px;
  margin-top: 3px;
  margin-bottom: 2px;
}

</STYLE>
</HEAD>
<BODY>
<div id="content">
<p><img src="$imageurl">
<BR><A HREF="$url">$title</A>
</p>
```

```
</BODY>
</HTML>
EONICEBOX;
echo $box_str;
?>
```

Those of you who do not have Amazon Associates accounts can use Listing 41-2 for testing, which you should save as mysqlbible.xml somewhere under your web tree.

LISTING 41-2

XML sample (mysqlbible.xml)

```
<?xml version="1.0" ?>
<ItemLookupResponse xmlns="http://webservices.amazon.com/
AWSECommerceService/2008-04-07">
<OperationRequest>
<RequestId>15409f90-7eca-4743-90e2-c776913b4e6c</RequestId>
<Arguments>
<Argument Name="Operation" Value="ItemLookup"></Argument>
<Argument Name="Service" Value="AWSECommerceService"></Argument>
<Argument Name="Version" Value="2008-04-07"></Argument>
<Argument Name="ItemId" Value="0764549324"></Argument>
<Argument Name="IdType" Value="ASIN"></Argument>
<Argument Name="AWSAccessKeyId" Value="XXXX"></Argument>
<Argument Name="ResponseGroup" Value="Medium,OfferFull"></Argument>
</Arguments>
<RequestProcessingTime>0.0245840000000000</RequestProcessingTime>
</OperationRequest>
<Items>
<Request>
<IsValid>True</IsValid>
<ItemLookupRequest>
<Condition>New</Condition>
<DeliveryMethod>Ship</DeliveryMethod>
<IdType>ASIN</IdType>
<MerchantId>Amazon</MerchantId>
<OfferPage>1</OfferPage>
<ItemId>0764549324</ItemId>
<ResponseGroup>Medium</ResponseGroup>
<ResponseGroup>OfferFull</ResponseGroup>
<ReviewPage>1</ReviewPage>
<ReviewSort>-SubmissionDate</ReviewSort>
<VariationPage>All</VariationPage>
</ItemLookupRequest>
```

```
</Request>
<Item>
<ASIN>0764549324</ASIN>
<DetailPageURL>http://www.amazon.com/MySQL-Bible-CDROM-Steve-Suehring/dp/07645
49324%3FSubscriptionId%3D03S5FBC8SMHRV9HBDWG2%26tag%3Dws%26linkCode%3Dxm2%26ca
mp%3D2025%26creative%3D165953%26creativeASIN%3D0764549324</DetailPageURL>
<SalesRank>780753</SalesRank>
<SmallImage><URL>http://ecx.images-amazon.com/images/I/51WJRNR4ENL._SL75_.
jpg</URL><Height Units="pixels">75</Height><Width Units="pixels">60</Width></
SmallImage>
<MediumImage><URL>http://ecx.images-amazon.com/images/I/51WJRNR4ENL._SL160_.
jpg</URL><Height Units="pixels">160</Height><Width Units="pixels">128</
Width></MediumImage>
<LargeImage><URL>http://ecx.images-amazon.com/images/I/51WJRNR4ENL.jpg</
URL><Height Units="pixels">475</Height><Width Units="pixels">380</Width></
LargeImage>
<ImageSets>
<ImageSet Category="primary"><SwatchImage><URL>http://ecx.images-amazon.com/
images/I/51WJRNR4ENL._SL30_.jpg</URL><Height Units="pixels">30</Height><Width
Units="pixels">24</Width></SwatchImage><SmallImage><URL>http://ecx.images-
amazon.com/images/I/51WJRNR4ENL._SL75_.jpg</URL><Height Units="pixels">75</
Height><Width Units="pixels">60</Width></SmallImage><ThumbnailImage><URL>
http://ecx.images-amazon.com/images/I/51WJRNR4ENL._SL75_.jpg</URL><Height
Units="pixels">75</Height><Width Units="pixels">60</Width></ThumbnailImage><T
inyImage><URL>http://ecx.images-amazon.com/images/I/51WJRNR4ENL._SL110_.jpg</
URL><Height Units="pixels">110</Height><Width Units="pixels">88</Width></Tin
yImage><MediumImage><URL>http://ecx.images-amazon.com/images/I/51WJRNR4ENL._
SL160_.jpg</URL><Height Units="pixels">160</Height><Width Units="pixels">128</
Width></MediumImage><LargeImage><URL>http://ecx.images-amazon.com/
images/I/51WJRNR4ENL.jpg</URL><Height Units="pixels">475</Height><Width
Units="pixels">380</Width></LargeImage>
</ImageSet>
</ImageSets>
<ItemAttributes>
<Author>Steve Suehring</Author>
<Binding>Paperback</Binding>
<DeweyDecimalNumber>005.7565</DeweyDecimalNumber>
<EAN>9780764549328</EAN>
<Edition>Pap/Cdr</Edition>
<ISBN>0764549324</ISBN>
<Label>Wiley</Label>
<Languages><Language><Name>English</Name><Type>Original Language</Type></Langu
age><Language><Name>English</Name><Type>Unknown</Type></Language><Language><Na
me>English</Name><Type>Published</Type></Language></Languages>
<ListPrice><Amount>4999</Amount><CurrencyCode>USD</CurrencyCode><FormattedPric
e>$49.99</FormattedPrice></ListPrice>
<Manufacturer>Wiley</Manufacturer>
<NumberOfItems>1</NumberOfItems>
<NumberOfPages>775</NumberOfPages>
```

```
<PackageDimensions><Height Units="hundredths-inches">170</Height><Length
Units="hundredths-inches">920</Length><Weight Units="hundredths-pounds">240</
Weight><Width Units="hundredths-inches">770</Width></PackageDimensions>
<ProductGroup>Book</ProductGroup>
<ProductTypeName>ABIS_BOOK</ProductTypeName>
<PublicationDate>2002-06-15</PublicationDate>
<Publisher>Wiley</Publisher>
<Studio>Wiley</Studio>
<Title>MySQL Bible with CDROM</Title>
</ItemAttributes></Item></Items></ItemLookupResponse>
```

The result of the REST client script is shown in Figure 41-1.

REST client gets XML and outputs HTML.
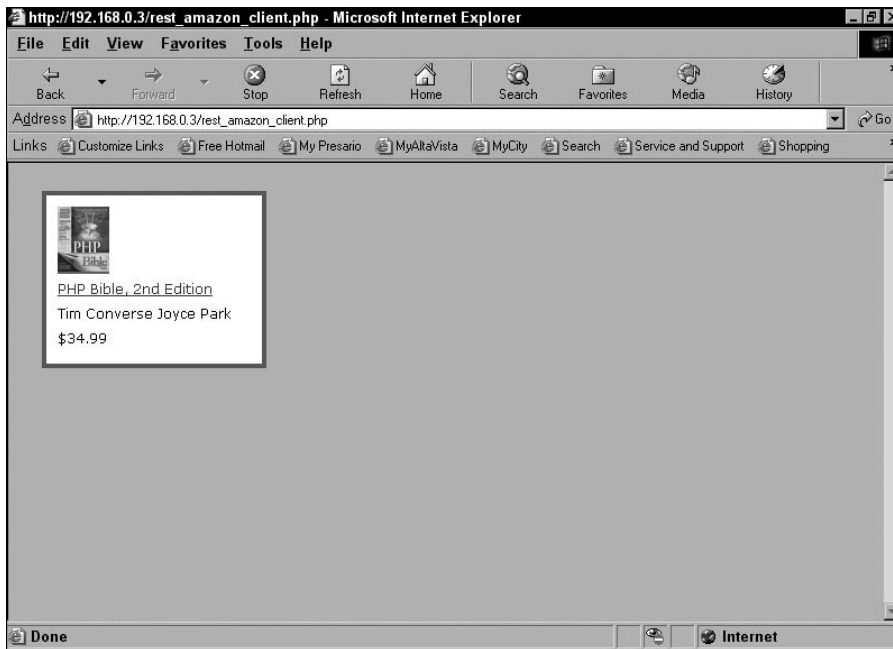
As you can see from the XML sample, there is a lot of potential data from the Web service that we didn't use — suggested price, release date, and so on. This demonstrates how easy it is to pick out just the data you want from the feed using Simple XML. According to the current Amazon Web services rules, you could poll its Web service for fresh XML every second, so you could keep this part of your information very fresh.

Summary

Web services is an emerging field of programming. Web services offer immediate payoffs in data transfer, especially for web applications, and may finally unlock the promise of distributed computing. Although there is a lot of hype and misinformation out there about the technologies, and although much of the action is happening inside intranets and in semiprivate transactions, you can start familiarizing yourself with PHP to both create and consume Web services.

The three main Web services standards in discussion now are REST, XML-RPC, and SOAP. REST is the most lightweight and easy to use but offers the least functionality. SOAP is the most complicated and, as a consequence, has the most interoperability problems, but large vendors such as Microsoft and IBM have thrown their support behind it. XML-RPC offers a nice blend of power and simplicity but lacks big-vendor support.

PHP can be used to create servers and clients in REST, XML-RPC, and SOAP. However, particularly with SOAP, the syntax is so complex that a third-party package may be highly useful for help with the serialization, type-shuffling, and request-creation steps required by Web services.

Chapter 42

Creating Graphics with PHP

In this chapter, we delve into how to use PHP to create graphics of your own and display them to the user. Although we spend a little bit of time on pure HTML "graphics," our primary focus is on creating images on the fly by using the gd library. This library helps you create images such as PNGs and JPEGs, which you can then link to from dynamically generated HTML pages or send to the user as standalone web pages.

Your Options

Just to see where image creation fits into the web-scripting world, look at the following spectrum of choices, in order of increasing dynamics:

- You can have no graphics at all and display purely textual information.

- You can embed static images in your HTML, whether created by yourself or by other people.

- You can write programmatically generated HTML pseudographics.

- You can embed static image graphics (or even image animations, if you insist) in your HTML pages, but display different ones conditionally.

- You can use gd to pre-generate static graphics for all the cases that may possibly arise from your code, store them in files, and display them conditionally.

- You can create graphic images on demand in response to user input.

We start off with the third option (HTML graphics) and then devote most of the rest of the chapter to the last one, which is the most interesting case.

HTML Graphics

You know those horizontal, colored bar graphs you see all over the web, especially in connection with poll results? It looks as though some graphics are being used in creating these graphs, but in truth there's just a couple of canned color images and the magic of image scaling in HTML. This graphing technique is actually very useful, and we include it here because it's very easy to create graphs like this dynamically from PHP.

Before we get into this data visualization technique, we need some data. Listing 42-1 shows a small sample data set, which we imagine has been produced from a survey of programmers asked about their favorite languages and operating systems. The data is stored in a MySQL database, in a single table, with the following definition:

```
CREATE TABLE programmers (
  id int(11) NOT NULL auto_increment,
  sex char(1) default NULL,
  age int(11) default NULL,
  language varchar(30) default NULL,
  os varchar(30) default NULL,
  country varchar(30) default NULL,
  continent varchar(30) default NULL,
  PRIMARY KEY  (id)
);
```

LISTING 42-1

Sample data set

```
+--+---+---+----------+---------+------------+---------------+
|id|sex|age| language | os      | country    | continent     |
+--+---+---+----------+---------+------------+---------------+
|1 |F  |33 | PHP      | Linux   | USA        | North America |
|2 |M  |41 | Java     | Solaris | USA        | North America |
|4 |M  |31 | C++      | Solaris | USA        | North America |
|5 |M  |45 | Lisp     | MacOS   | USA        | North America |
|6 |M  |25 | C        | Solaris | Antarctica | Antarctica    |
|7 |F  |17 | PHP      | Linux   | Denmark    | Europe        |
|8 |M  |21 | Perl     | Linux   | UK         | Europe        |
|9 |M  |14 | PHP      | Linux   | UK         | Europe        |
|10|F  |21 | Perl     | Linux   | Germany    | Europe        |
```

```
|11|F  |38 |  PHP        |  Linux   |  Germany     |  Europe        |
|12|M  |26 |  C++        |  Windows |  USA         |  North America |
|13|M  |22 |  PHP        |  Windows |  France      |  Europe        |
|14|M  |17 |  PHP        |  Linux   |  Japan       |  Asia          |
|15|F  |38 |  C          |  Solaris |  South Korea |  Asia          |
|16|F  |19 |  PHP        |  Linux   |  Canada      |  North America |
|17|F  |32 |  Perl       |  Linux   |  France      |  Europe        |
|18|M  |32 |  Java       |  Solaris |  Mexico      |  North America |
|19|F  |23 |  PHP        |  Solaris |  Brazil      |  South America |
|20|F  |19 |  PHP        |  Linux   |  Finland     |  Europe        |
|21|M  |21 |  PHP        |  Linux   |  Brazil      |  South America |
|22|M  |51 |  Java       |  Linux   |  UK          |  Europe        |
|23|M  |29 |  Java       |  Linux   |  Japan       |  Asia          |
|24|M  |29 |  Java       |  Solaris |  China       |  Asia          |
|25|M  |21 |  C++        |  MacOS   |  Germany     |  Europe        |
|26|M  |21 |  Perl       |  Solaris |  France      |  Europe        |
|27|M  |27 |  PHP        |  Linux   |  India       |  Asia          |
|28|M  |31 |  Perl       |  Linux   |  India       |  Asia          |
|29|M  |17 |  C          |  Linux   |  Pakistan    |  Asia          |
|30|M  |45 |  PHP        |  Windows |  USA         |  North America |
|31|F  |22 |  Java       |  Windows |  Italy       |  Europe        |
|32|F  |33 |  C          |  Linux   |  Spain       |  Europe        |
+--+---+---+-----------+---------+------------+--------------+
```

So say that our goal is to visualize counts of the distribution of values for different columns — we want to know not only how many of our respondents list this or that programming language, but also to see comparisons graphically.

Although our data is in a MySQL database, the display portion of this code need not be tied to that. We may want to use it for a different purpose entirely. So we break out a separate function that produces a bar graph from an array in a particular format, and only later hook that up to code that produces the requisite array via SQL queries. The code to translate an array into a bar graph is shown in Listing 42-2.

LISTING 42-2

bar_graph.php

```php
<?php

function array_to_bar_graph ($array, $max_width) {
  // expects as input an array where the keys
  //    are string labels and the values are
  //    numbers. Values must be non-negative
```

```
// returns an HTML bar graph as a string
// assumes bar[1-5].gif, located in images/

foreach ($array as $value) {
  if ((IsSet($max_value) &&
      ($value > $max_value)) ||
      (!IsSet($max_value))) {
    $max_value = $value;
  }
}
$pixels_per_value = ((double) $max_width)
                    / $max_value;

$string_to_return = "<TABLE CELLPADDING=5>";
$counter = 0;
foreach ($array as $name => $value) {
  $bar_width = $value * $pixels_per_value;
  $image_no = ($counter % 5) + 1;
  $string_to_return .=
   "<TR><TD>$name ($value)</TD>
    <TD><IMG SRC=\"images/bar$image_no.gif\"
             WIDTH=$bar_width
             HEIGHT=10>
    </TD></TR>";
  $counter++;
}
$string_to_return .= "</TABLE>";
return($string_to_return);
}
?>
```

The bar graph code is extremely simple — it iterates through an array, which is assumed to have names for keys and quantities for values. It normalizes the maximum value to a fixed-width bar and calculates the width of all the other bars proportionally. Finally, it displays bars by using the scaling parameters in the `<IMG>` tag to give each variable a fixed height and an appropriate width. It cycles through a list of five images, which are premade one-color GIFs (which could as well be PNGs) and could be created by using your favorite graphics program. As long as these images are monocolor, their size and shape are irrelevant. (If you don't have any such images handy, you can find the ones we used at the code download site: `www.troutworks.com/phpbook`.)

Now that we can display names and associated values in a bar graph, we can hook that up to the database via a Web form and an SQL query. Code for this is shown in Listing 42-3.

LISTING 42-3

bar_graph_form.php

```php
<?php
include_once("dbconnect.inc");
include_once("bar_graph.php");
mysql_connect($hostname, $user, $password);
mysql_select_db("c43");

if (IsSet($_POST['COLUMN_NAME'])) {
$column_name = mysql_real_escape_string($_POST['COLUMN_NAME']);
  $query = "select $column_name, count(*)
            from programmers
            group by $column_name";
  $result = mysql_query($query)
            or die("Error in database interaction<BR>".
                     mysql_error());
  $array_collection = array();
  while ($row = mysql_fetch_row($result)) {
    $name = $row[0];
    $count = $row[1];
    $array_collection[$name] = $count;
  }
  $bar_graph =
    array_to_bar_graph($array_collection,
                       300);
}
else {
  $bar_graph = "";
}

$self = "bar_graph_form.php";
$form = <<<EOT
<FORM METHOD=POST ACTION="$self">
<H3>Choose a table column for graphing</H3>
<SELECT NAME=COLUMN_NAME>
<OPTION VALUE=os>os
<OPTION VALUE=language>language
<OPTION VALUE=continent>continent
<OPTION VALUE=sex>sex
</SELECT>
<INPUT TYPE=SUBMIT NAME=SUBMIT>
</FORM>
EOT;

$page = <<<EOT
<HTML><HEAD><TITLE>Survey data</TITLE></HEAD>
<BODY>
```

```
$form
<BR>
$bar_graph
</BODY></HTML>
EOT;

echo $page;
?>
```

The form is self-submitting and loads a file called `dbconnect.inc`, which contains the appropriate login, password, and database name, like this:

```
<?php
$hostname = "YOURMYSQLHOST";
$user = "YOURMYSQLUSER";
$password = "YOURMYSQLPASSWORD";
?>
```

All that is supplied by the form submission is the name of the column. Starting with that, the code submits an SQL statement to count all the distinct values that occur for that column and then creates an array by using names and the corresponding counts. What remains is to feed the resulting array to the bar graph code from Listing 42-1 and to do some layout. The results for two different columns are shown in Figures 42-1 and 42-2. (Since this is a grayscale book, you won't see interestingly different colors in the diagram, but you should at least see bars of different sizes.)

FIGURE 42-1

HTML bar graph 1
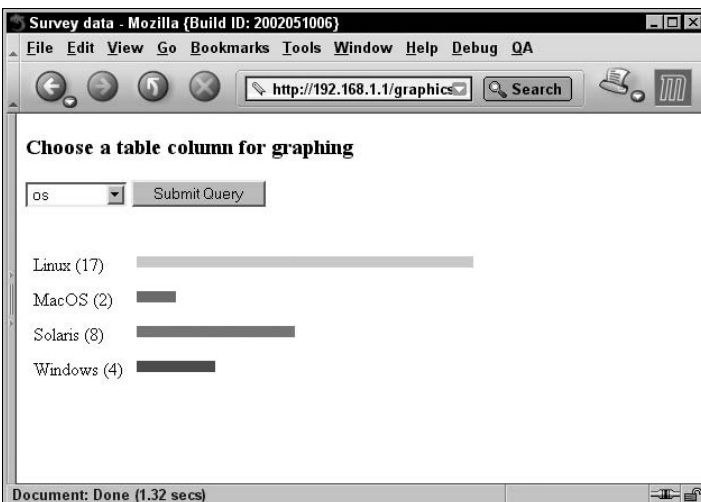
HTML bar graph 2
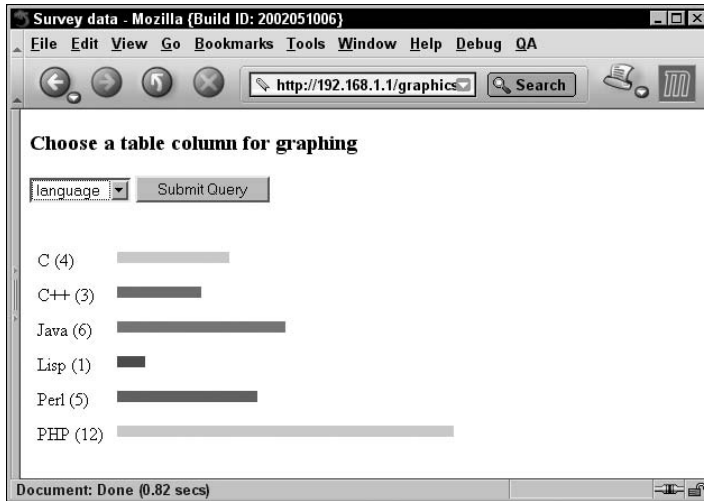
Creating images using gd

Having mostly exhausted the graphic possibilities afforded by vanilla HTML, let's turn our attention to creating real standalone graphics by using the gd library.

What is gd?

What is gd, anyway? The gd toolkit is a C code library for creating and manipulating images, which was originally created by the kind and clever people at Boutell.com (`www.boutell.com`). gd is not a graphics or paint program in and of itself, as it has no standalone application or GUI. Instead, it provides functions that programs can call to do these manipulations, and any C program that wants to can link against that library to use the routines. The PHP developers have done this and, in fact, have written a set of interface functions that make it easy to call gd routines from PHP. But nothing in gd is specific to PHP, and there are interfaces to it from several other languages and environments, including Perl, Tcl, Pascal, Haskell, and REXX.

gd lets you call functions to create images (initially blank, like a clean sheet of paper), draw and paint on those images in various ways, and ultimately convert the image from gd's internal image format to a standard image format, and send it off to its ultimate fate (display in a browser or storage in a file or database). And because all this is under programmatic control rather than human control, these created images can be arbitrarily complex, and they can depend on anything in your program that you would like to have them depend on.

Image formats and browsers

The `gd` library can, in principle, import and output images in a wide variety of formats. The three image formats we talk about at all seriously are GIF, JPEG, and PNG, although for examples we focus mostly on the last of these.

The GIF and PNG formats essentially exist to describe a grid of colored cells corresponding to pixels, with a few complications. The first complication is that the cells may contain actual color values or they may contain indexes in a table of color values. (The former is more expressive because any number of different colors may be used, and the latter is more compact.)

Another complication is that, although the conceptual representation of GIFs and PNGs is fairly simple, in practice they are always read, written, and transferred in compressed form. Compression is necessary because a grid of cells is a costly thing to specify. A simple 500 × 400 pixel image is 200,000 pixels — if each pixel needs 3 bytes to specify, then we're over half a megabyte already. Compression is a large topic, but most compression algorithms take advantage of redundancy in the image to make it smaller. (There are more concise ways to say that every pixel is green than specifying every pixel's green color value individually.) Unfortunately, there is a lot more to compression algorithms than that — enough so that the compression algorithm used for writing GIFs is patented.

Early browsers were written using GIF as the graphics format of choice, and it wasn't until that practice had been under way for a while that it became clear that the patent holder was going to insist on going after people who used the compression algorithm. This left web graphics in a bit of a bind — GIF was the *lingua franca*, but you couldn't legally create such graphics, at least without paying a license fee. The PNG format has come to the rescue in a sense — recent versions of major browsers support it, and with that support it plays much the same role as GIF.

Compression is different in the case of JPEGs, as well, although not for legal reasons. Compression for GIFs and PNGs is lossless, meaning that if you compress and then uncompress an image, you should have your exact original image back. The JPEG compression, on the other hand, is lossy. Essentially, if redundancy helps compression, JPEG compression tries to introduce a little bit of extra redundancy into the image before compression, mostly in ways that the human eye won't notice. This is particularly effective with photographic images, but it does mean that sometimes the compression/uncompression cycle doesn't leave you with exactly what you started with.

Because JPEG is better for photographic images than the kinds of images we're making, and because deciding on the export format is a final step anyway, we've decided to focus on PNG graphics exclusively. If you would rather produce JPEGs, it is a simple matter to change the export functions appropriately.

Installation

Installing `gd` and getting it to work with PHP is, frankly, a pain. This is not because of any weakness in either the PHP codebase or the `gd` codebase but is all about configuration issues: sorting out the likely and actual locations of the libraries `gd` depends on and making sure that everything can build and link appropriately. So the happiest situation possible is to find out that `gd` is already installed,

and PHP already has gd support enabled (whether that's due to the diligence of your webhost or because the PHP you installed by yourself had it included).

So the first step in installing gd is: Check to see if it has already been installed. Whether you are running via a webhost or are in command of your own installation, start off as always by putting the following into a file and viewing the result in a browser:

```php
<?php
    phpinfo();
?>
```

After you have the displayed page, just do a text search for gd in the browser window — you may find a subsection that describes to what extent gd is enabled in your PHP installation. If you only want to produce certain kinds of images (PNGs, for example) and phpinfo() tells you that support for that image type is enabled, then you may be good to go. If the gd version includes the word *bundled*, you are using the gd that is bundled with PHP.

If this fails, and if you are in control of your PHP installation, you will have to install and configure gd. (If, instead, your PHP installation is run by a hosting company, your options may be reduced to asking them to provide gd support, or to switching webhosts.)

Using the PHP-bundled version of gd removes some, but not all, of the hassle of a gd install; if you use the bundled version itself, you have the gd library but not necessarily the libraries that gd needs. The gd library itself depends on several other libraries: libpng (for manipulating PNG images), zlib (used in compression), and jpeg-6b or later (if you want to manipulate JPEG images). (Only gd, libpng, and zlib are necessary for the examples in this chapter.) These will be present already in many Linux installations, and if so it may be sufficient to include a with flag (such as --with-zlib) without specifying the installation directory. If you are configuring PHP yourself, adding the --with-gd flag will cause the bundled version of gd to be included. Use --with-gd=path instead if you want to point to an alternate version.

If you find that you lack one or more of the necessary libraries, you will have to build them. The documentation at www.libgd.org is a good place to start to find the current versions.

gd Concepts

While an image is being constructed or manipulated in the gd toolkit, it is stored in a gd-specific format that doesn't correspond to any conventional image type. Images can in theory be exported in this gd format, but it's unusual to do so because the resulting image is not compressed and cannot be displayed in a browser or conventional graphics program.

An image in the gd toolkit has a width, a height, and color information for all the *width x height* many pixels. (See the "Colors" section for more detail on how colors are stored.) Usually a program starts off its interaction with gd by either creating a new blank image (which is drawn and painted on) or by importing an image from a file. The next steps are typically (1) allocate colors in the image,

(2) draw, paint, or otherwise transform the image, (3) translate the image to a conventional format (for example, PNG, JPEG), and send it to output.

Colors

There are two ways of representing colors in gd images: *palette-based*, which is limited to 256 colors, and *truecolor*, which can store an unlimited number of distinct RBG color values. In gd 1.x, palette-based colors were the only alternative; gd 2.x and the PHP-bundled version of it offers both palette-based images and truecolor images. Note that a given gd image is either palette-based or truecolor; there is no notion of adding true colors to a palette-based image.

To get an initial blank palette-based image, you call the function ImageCreate(); to get a truecolor image, call ImageCreateTrueColor().

Palette-based images

Colors are specified in a red-green-blue (RGB) format, with three numbers between 0 and 255. The color specified by (255, 0, 0), for example, is bright red; (0, 255, 0) is green; (0, 0, 255) is blue; (0, 0, 0) is black; (255, 255, 255) is white; and (127, 127, 127) is gray. You can tweak these values to your heart's content to design new colors.

Any drawing into an image must be done in a particular color, and colors must be allocated in an image before they are used. Also, the first color allocated into an image automatically becomes the background color. So, colors are not optional in any sense, and usually color allocation is the first thing you do after creating a new blank image.

Colors in palette-based images are created by using imagecolorallocate(), which takes as arguments an (already created) image, and three integers specifying the proportion of red, green, and blue. The return value is an integer, which specifies the index of the new color in the image's internal palette. You must hang on to this return value in a variable, because you need the index value for any future drawing using that color. Palette-based images can have a maximum of 256 colors. (It may or may not be obvious what's going on under the hood here, but every pixel in a palette-based image is actually a single byte, which stores an index into the 256-color palette.)

Note that the index returned by allocating a color in an image makes sense only for that image. If you assign an allocated color to the PHP variable $black, it won't work to use that variable as the color input for a drawing command called on a different image.

Truecolor

In gd 2.0 and later, you can also create images that are not palette-based, where every pixel stores an arbitrary RGB color value. In this truecolor format, the number of colors is essentially unlimited. This can be useful not only for the free range of your artistic expression, but for faithfully representing truecolor PNGs and JPEG images that have been loaded into gd.

Aside from the initial function to create an image, and the lack of limitation on distinct colors, working with truecolor images is similar to working with palette-based images. In particular, you still call ImageColorAllocate() to create new colors, and hang on to the return value for later commands

to use; it just so happens that the returned value will be an RGB color rather than an index into a palette. Also, in truecolor images there is no notion of a background color created as a side effect of `ImageColorAllocate()`; all pixels are initialized to black.

Transparency

`gd` 2.x supports transparency in the form of an *alpha value* (in addition to the red, green, blue values) that specifies how transparent the given color is. This allows you, for example, to overlay a shape onto another one without simply occluding the first shape.

Many of the image functions in PHP have an analog with "alpha" in its name, which indicates that it deals with a four-value (R,G,B,A) color. For example, while `ImageColorAllocate()` expects three arguments, `ImageColorAllocateAlpha` expects a fourth argument between 0 and 127. A value of zero indicates that the color is completely opaque; a value of 127 means that the color is completely transparent.

Drawing coordinates and commands

After you create an image within `gd`, you have an implicit coordinate system for drawing on it, determined by the width and height you specified.

In this coordinate system, the origin (0, 0) is at the top-left corner of the image, and the positive direction for *x* values is to the right, whereas the positive direction for *y* values is down. (This is often true of computer graphics coordinate systems, but you may be more accustomed to a lower-left origin if you learned analytic geometry in school.)

There are many drawing commands, including but not limited to drawing line segments, rectangles, arcs, and setting particular pixel values. Note that the end effect of all these painting and drawing commands is to set the value of pixels. There is no memory retained of the commands that changed the pixels and, therefore, no way to undo drawing commands or separate out the effects of distinct commands.

Nothing stops you from drawing outside the bounds of the image you have specified, but such drawing has no visible effect. A rectangle with coordinate values that are all negative, for example, is not visible.

Format translation

All this drawing and image manipulation is done on the image in its `gd`-internal format. After your script is done, it can use one of the translation-and-output commands (`imagetopng`, `imagetojpeg`, and so on) to translate the image to the desired graphics format and echo it out to the user's browser (or to a file).

Freeing resources

After you have sent a translation of your completed `gd` image off to the user, you are done with the internal version and should dispose of it. The right way to do this is to call `imagedestroy()` with the image as an argument.

Functions

We are not planning to individually list and describe all the functions in PHP's gd interface in this chapter; for that, we refer you to the "Image Processing and Generation" section of the manual at www.php.net. Here we summarize the most important functions. Most of the gd functions are in one of the categories shown in Table 42-1. Note that the function names in this table have internal capital letters at word breaks for clarity, but we may not always observe this when writing code because PHP function names are not case-sensitive.

TABLE 42-1

Breakdown of gd Functions

Type	Examples	Notes
Image-creation functions	ImageCreate(), ImageCreateTruecolor(), ImageCreateFromGd(), ImageCreateFromJpeg()	These functions return a new gd image. ImageCreate() takes a width and height as arguments; others take a filepath, URL, or string containing a preexisting image to load in and convert to gd.
Color allocation	ImageColorAllocate(), ImageColorAllocateAlpha(), ImageColorDeallocate()	ImageColorAllocate() takes an image and the desired red, green, and blue color values, and returns the color value to be used for subsequent drawing. ImageColorAllocateAlpha takes an additional transparency value (0–127).
Color matching	ImageColorClosest(), ImageColorClosestAlpha(), ImageColorExact(), ImageColorExactAlpha()	Return the index of a matching color in a palette image. The "Closest" functions return the best-matching color by RGB distance; the "Exact" functions return a color only if it is identical, -1 otherwise. "Alpha" functions operate on 4-value (transparent) colors.
Line-drawing functions	ImageLine(), ImageDashedLine(), ImageRectangle(), ImagePolygon(), ImageEllipse(), ImageArc()	These functions draw lines or curves in the specified shapes. Usually the first argument is an image, the last argument is a color, and the intermediate arguments are x- and y- coordinates.
Pen-setting functions for line drawing	ImageSetStyle(), ImageSetThickness()	These functions alter settings that affect the lines created by later line-drawing commands. (Some of these are available only with gd 2.0.1 or later.)

Type	Examples	Notes

Painting and filling functions	`ImageFilledRectangle(),`
	`ImageFilledEllipse(),`
	`ImageFilledRectangle(),`
	`ImageFilledPolygon(),`
	`ImageFilledArc(),`
	`ImageFill()`
		Usually analogous to corresponding line-drawing functions but with areas filled rather than outlined. The special `ImageFill()` function "flood fills" outward from a specified x-y-coordinate with a given fill color. (Some of these functions require gd 2.0.1 or later.)

Text functions	`ImageString(),` `ImageLoadFont()`	`ImageString` takes as arguments an image, a font number, x- and y-coordinates, a text string, and a color. If the font number is between 1 and 5, one of the five built-in fonts is used to draw the string in the given color. A number greater than 5 indicates a result of loading a custom font with `ImageLoadFont().`
Exporting functions	`ImagePng(),` `ImageJpeg()`	These functions convert the internal gd image to the relevant image format and then send to output. If only one argument (an image) is given, the image is echoed to the user; if an additional path name argument is given, the destination is a file.
Image-destruction function	`ImageDestroy()`	Takes an image argument and frees all resources associated with the image.

Images and HTTP

Before the user's browser can display an image appropriately, it has to know that an image is coming, and what the image format is. So it is, unfortunately, not sufficient to simply embed a call to `ImageToPng()` in your generated HTML and have an image show up. You essentially have three choices in regard to intermixing images with PHP-generated HTML.

Full-page images

You can make the entire generated page an image. In this case, you need to send an HTTP header before the image data, announcing that an image of a certain type is on the way. You may, for example, have lines such as the following near the end of your script:

```
// ... code to create image in $image
header("Content-type: image/png");  // announcement to browser
imagepng($image);  // sending actual PNG-converted image data
imagedestroy($image);  // freeing resources
```

This approach has the benefit that you can use any kind of information, including POST arguments, to decide what the image should contain. The downside is that the resulting page can't contain any conventional HTML. In fact, you need to be careful that no textual output is sent from your scripts before the header and image because this causes content to be sent prematurely. In this case, you get a Headers already sent . . . error.

Embedded images from files

Of course, HTML has had the <IMG> tag for a long time. This enables you to embed an image by specifying its file path or URL, like this:

```
<IMG SRC="my_png.png">
```

This works with static image files, but there is no reason why the image can't have been recently created. So you can have a script that (1) creates an image, (2) writes the image data to a local file, and then (3) produces HTML with an appropriate <IMG> tag referring to the file that you just made.

The only drawbacks to this approach are that you're introducing file writes, which may be time-consuming, into the page-generation process, and that you need to figure out what to do with the files after you are done with them. There is one situation this approach is perfect for, however, which is creating and caching images that represent a finite set of possibilities. In this case, you have some way to map from a situation to an image filename. Whenever a display situation arises, you check to see if you already have the appropriate file — if you do, you simply refer to it by using an <IMG> tag, and if not, you create the image, write it out to a file, and then refer to it. Eventually, you should need to do no more creation.

Embedded images from scripts

Finally, there is no reason why you cannot have a standalone generated image, as in the section "Full-page images," but, in turn, embed that URL in a different dynamic page via an <IMG> tag. The only difficulty lies in how to communicate necessary data to the dependent page. You may, for example, have an embedded image tag like this:

```
<IMG SRC="ballpage.php?ball_color=green&ball_position=5>
```

where ballpage.php happened to return PNG images of colored balls in various positions in the image.

There is a gotcha lurking here because both web servers and browsers sometimes pay attention to the suffix of the served file, and in different ways. You may need the suffix of ballpage to be .php to let Apache (for example) know that the server-side code should be interpreted as PHP (although this behavior can be controlled with configuration files). Some broken browsers, however, may insist that a file that ends in .php cannot be an image despite the headers we are sending. This technique requires some cross-browser testing to make sure that your intended users are seeing the same thing you are.

Now it's high time to move on to an example of using gd to create images.

Example: fractal images

There's a fine tradition of livening up the potentially unexciting topic of line drawing by using fractals as examples, and your authors are not about to mess with tradition. In addition to showing how you can produce a complex image programmatically, this kind of example is also a good fit for PHP because its arrays and loose data types make it very easy to build complex data structures corresponding to fractal images, without a lot of declarations.

What's a fractal? It's a shape that is self-similar, in that the parts of a fractal have a shape similar to the shape of the whole, and the parts of those parts have a similar shape, and so on.

In theory, you can keep zooming into ever-smaller pieces of an ideal fractal, and keep finding the same patterns repeated. In practice, computer-generated fractals bottom out after some limited number of generations into nonfractal shapes like simple curves and line segments.

An example of the kind of image we're going to create is shown in Figure 42-3. Although it may not look like it, this image is simply a lot of small line segments with endpoints connected into a path.

FIGURE 42-3

Fractal 1
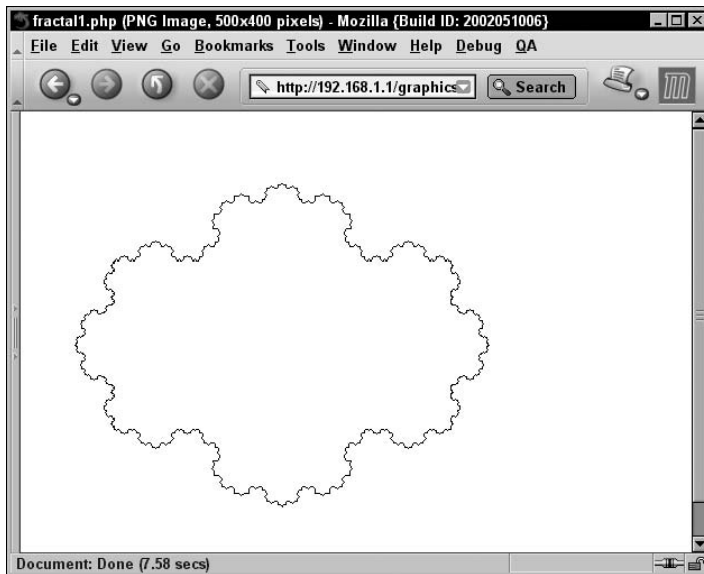
Our job is to calculate the endpoints of all those line segments and then display them appropriately as a PNG image. We're going to be slightly more ambitious than simply creating a one-off piece of fractal display code and construct a little framework that makes it easy to vary the fractal parameters and to generate new kinds of displays.

To start with, we build some data structures to represent the complex shapes that we are displaying. We use these data structures both in our intermediate calculations and for drawing the end result. Let's say somewhat arbitrarily that:

- A coordinate point is a pair of numbers.
- A path is a list of points.

We end up drawing paths by drawing line segments between all the points in a path. If we want to draw a simple line segment, we draw a path that has two points in it; if we want to draw a rectangle, then we draw a path that has five points in it (with the starting point repeated to close off the rectangle). (We could have made a line segment a primitive entity here, but paths seemed more concise for our fractal purposes.)

Now, how shall we represent points and paths? The easiest way to make lists of things in PHP is to use arrays. So we declare that a point is an array that happens to contain two numbers, and a path is an array that happens to contain a sequence of points. The resulting structures are multidimensional PHP arrays, but if we define well-named constructor and accessor functions, we can forget about that and just write code that acts as though these things are genuine data types.

Listing 42-4 shows the code used to create Figure 42-3. The code defines the data types in terms of functions to create them (starting with make_), functions to access their parts, and functions to draw them into an image (starting with display_). Points cannot be drawn and have no display function; paths are drawn by drawing lines between successive pairs of points.

The function transform_path takes an input path as first argument, and as second argument it takes the name of a function that, in turn, is expected to take a path as argument and return a path as a result. The third argument to transform_path() is a number of times that the path-to-path function should be successively applied to create a new path. The reason that this kind of second-order function is useful is that, otherwise, we may find ourselves writing a new looping function every time we want to build a new fractal. With this approach, we can bundle the varying part of the fractal code into a function that we pass into transform_path and avoid duplicating work.

The spike function takes a path as argument and returns a path where every two-point line segment has been replaced by a five-point line segment with a spike in the middle. The top-hat function does something similar, except that six points are involved, and the spike is rectangular. We also include a couple of functions to create rectangular paths of standard sizes, to use as starting points.

After loading all the functions from the included files, the remaining code creates a gd image of specific height and width and allocates colors into that image. (The background is white, and the lines are black.)

The fractal creation code starts off by creating a standard rectangular path (containing five points and, therefore, four [implicit] line segments). It then passes this off to the transform_path function, asking it to return the path that results from applying the spike() function to the rectangle four times. The rectangle path starts with four line segments, and every segment is itself replaced by

four segments. So the four successive iterations have 16 segments, 64 segments, 256 segments, and 1024 segments, respectively.

Then all that remains is to display the complicated path that we've generated. We call our own function `display_path()` to draw all the lines into the image, send off an HTTP header announcing a PNG, call `imagepng()` for the conversion and output, and then dispense with the internal `gd` image.

LISTING 42-4

fractal1.php

```php
<?php

// --- points ----

// A point is just a pair of numerical coordinates

function make_point ($x, $y) {
  return(array($x, $y));
}

function point_x ($point) {
  return($point[0]);
}

function point_y ($point) {
  return($point[1]);
}

// --- paths ---

// A path is a list of points

function make_path () {
  return array();
}

function add_point_to_path ($path, $point) {
  $path[] = $point;
  return($path);
}

function display_path ($image, $path, $color) {
  static $line_count = 0;
  $prev_point = NULL;
  foreach ($path as $point) {
    if ($point && $prev_point) {
      $line_count++;
```

```
        imageline($image,
                  point_x($prev_point),
                  point_y($prev_point),
                  point_x($point),
                  point_y($point),
                  $color);
    }
    $prev_point = $point;
  }
}

function transform_path ($path_input,
                         $function_name,
                         $iterations) {
  // Expects a path, a path-to-path function
  //  and a number of times to apply the
  //  function.
  // Returns a path
  $path_to_return = $path_input;
  for ($i = 0; $i < $iterations; $i++) {
    $path_to_return = $function_name($path_to_return);
  }
  return($path_to_return);
}

function spike ($path) {
  // Takes a path and returns a path
  $path_to_return = make_path();
  $prev_point = NULL;
  foreach ($path as $point) {
    if ($point && $prev_point) {
      $path_to_return =
          add_point_to_path($path_to_return,
                            $prev_point);
      $path_to_return =
          add_point_to_path($path_to_return,
            point_along_segment($prev_point,
                                $point,
                                0.25));
        $path_to_return =
          add_point_to_path($path_to_return,
            point_off_segment($prev_point,
                              $point,
                              0.5, 0.23));
      $path_to_return =
          add_point_to_path($path_to_return,
            point_along_segment($prev_point,
                                $point,
                                0.75));
      $path_to_return =
```

```
                add_point_to_path($path_to_return,
                                  $point);
    }
    $prev_point = $point;
  }
  return($path_to_return);
}

function top_hat ($path) {
  // Takes a path and returns a path
  $path_to_return = make_path();
  $prev_point = NULL;
  foreach ($path as $point) {
    if ($point && $prev_point) {
      $path_to_return =
          add_point_to_path($path_to_return,
                            $prev_point);
      $path_to_return =
        add_point_to_path($path_to_return,
          point_along_segment($prev_point,
                              $point,
                              0.35));
        $path_to_return =
          add_point_to_path($path_to_return,
            point_off_segment($prev_point,
                              $point,
                              0.35, 0.24));
        $path_to_return =
          add_point_to_path($path_to_return,
            point_off_segment($prev_point,
                              $point,
                              0.65, 0.24));
        $path_to_return =
           add_point_to_path($path_to_return,
             point_along_segment($prev_point,
                              $point,
                              0.65));
        $path_to_return =
           add_point_to_path($path_to_return,
                              $point);
    }
    $prev_point = $point;
  }
  return($path_to_return);
}


function point_along_segment ($first_point,
                              $second_point,
                              $proportion)
```

707

```
{
  $delta_x = (point_x($second_point) -
              point_x($first_point));
  $delta_y = (point_y($second_point) -
              point_y($first_point));
  return(make_point(point_x($first_point) +
                    $proportion * $delta_x,
                    point_y($first_point) +
                    $proportion * $delta_y));
}

function point_off_segment ($first_point,
                           $second_point,
                           $proportion,
                           $proportional_distance)
{
  $delta_x = (point_x($second_point) -
              point_x($first_point));
  $delta_y = (point_y($second_point) -
              point_y($first_point));
  return(make_point(point_x($first_point) +
                    $proportion * $delta_x -
                    $proportional_distance *
                    $delta_y,
                    point_y($first_point) +
                    $proportion * $delta_y +
                    $proportional_distance *
                    $delta_x));
}

function make_small_rectangle () {
  $path = make_path();
  $path = add_point_to_path ($path, make_point(75, 275));
  $path = add_point_to_path ($path, make_point(375, 275));
  $path = add_point_to_path ($path, make_point(375, 125));
  $path = add_point_to_path ($path, make_point(75, 125));
  $path = add_point_to_path ($path, make_point(75, 275));
  return($path);
}

function make_large_rectangle () {
  $path = make_path();
  $path = add_point_to_path ($path, make_point(5, 5));
  $path = add_point_to_path ($path, make_point(495, 5));
  $path = add_point_to_path ($path, make_point(495, 395));
  $path = add_point_to_path ($path, make_point(5, 395));
  $path = add_point_to_path ($path, make_point(5, 5));
  return($path);
}
```

```
$IMAGE_WIDTH = 500;
$IMAGE_HEIGHT = 400;

$image = imagecreate($IMAGE_WIDTH, $IMAGE_HEIGHT)
   or die("Could not create image");
$background_color = ImageColorAllocate($image, 255, 255, 255);
$drawing_color = ImageColorAllocate($image, 0, 0, 0);

$path = make_small_rectangle();
$path = transform_path($path, 'spike', 4);
display_path($image, $path, $drawing_color);

header("Content-type: image/png");
imagepng($image);
imagedestroy($image);


?>
```

Although we won't show it as a separate listing, we took a copy of Listing 42-4, changed the function name argument from `spike` to `top-hat`, and renamed the file `fractal2.php`. The resulting image is shown in Figure 42-4.

FIGURE 42-4

Fractal 2
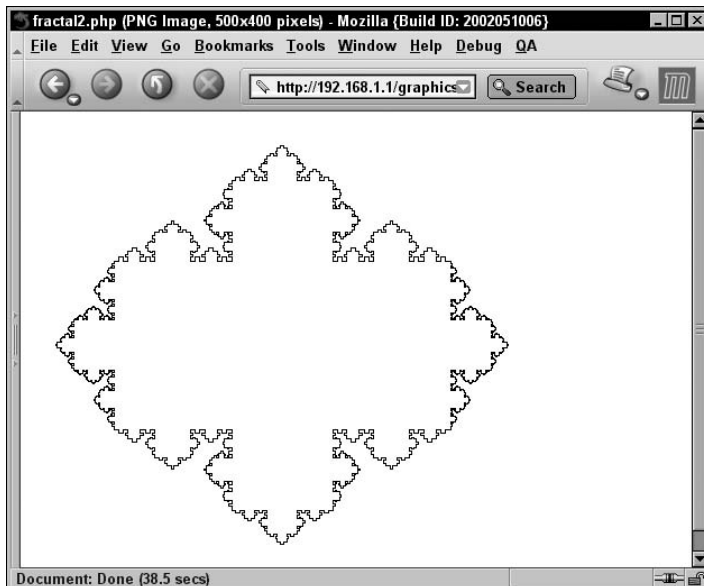
CAUTION   **Creating and displaying these images can be time-consuming and the more so the more line segments are created. Your web server may time out while the creation is happening. Your options then are to decrease the number of generations in the fractal code or to raise the timeouts in your web server or PHP configuration files.**

Tweaking fractal code is definitely an art, and your humble authors are not particularly good artists. We wish you luck in improving on our images.

CROSS-REF   **For a much more extended example of producing graphics with** gd**, see Chapter 45.**

Gotchas and Troubleshooting

Code to produce images can be especially difficult to debug, because some of the simplest tricks (for example, diagnostic print statements) can't be used as easily. What follows is a list of symptoms you may encounter in running gd-enabled PHP code and some things you can try to correct them.

Symptom: completely blank image

Sometimes your code runs without incident or apparent error, but the image that results is a blank slate, although you expected it to be full of graphic wonders. Some things to check (some obvious, some not):

- Are you drawing outside the bounds of the image? (If your image is 100   100, a small circle drawn at (200, 200) cannot be seen.)

- Are you drawing infinitesimally small graphics? (A circle with a center in range of the image with a radius of zero or near-zero may be completely undetectable.)

- Are you drawing by using the background color? (White-on-white is the same as white.)

Symptom: headers already sent

This problem is almost always due to printing text to output before the header call that announces a graphic image. Just as with other HTTP headers (such as those setting cookies), you must ruthlessly root out any printing of text before that call, even if that text is composed of blank lines or spaces.

One common pattern is to see something such as the following in your browser as you test:

```
Warning:  Division by zero in
          /usr/local/apache/htdocs/graphics/fractal1.php on
line 18
Warning:  Cannot add header information -
  headers already sent by (output started at
/usr/local/apache/htdocs/graphics/fractal1.php:18) in
/usr/local/apache/htdocs/graphics/fractal1.php on line 19)
PNG
IHDR [trailing off into binary gibberish]
```

The binary gibberish is, of course, your image data, which is being printed as text in your browser, resulting in nonsense characters. The reason you're seeing it as text is that the image announcement headers could not be sent, because some text was sent before those headers were encountered. And that text that was sent, in turn, was probably just (in this case) the text of the division-by-zero warning itself. Fixing the division-by-zero problem (or whatever the error or warning is in your case) may eliminate the printed error, which may make the header-sending statement happy, which may mean a successful image display.

If, instead, the very first thing you see is the warning about headers, you may be sending blank lines or spaces before the header without being aware of it. Look for any print statements, any included HTML, and (especially) any space at the beginning or end of files that have been included or required. If an included PHP file so much as ends with '?> ' rather than '?>', you may be sending a space's worth of HTML, which would cause text headers to be sent before the image headers are seen.

Symptom: broken image

How exactly this problem displays depends on your browser program — some display a sad, visibly broken image icon, while Firefox may politely inform you that your image can't be displayed because it contains errors. Either way, though, the problem is that your browser cannot read the data in the image format you said you were sending. Some possible causes are:

- The flip side of the previous `Headers already sent` problem: You may be printing random text without being aware of it but, in this case, *after* the image header has already been sent rather than before. This text is interfering with the stream of image data.

- You have misspelled the variable containing your image — for example `imagepng($imag)` where you meant `imagepng($image)`. You are actually calling the convert-and-send function on nothing at all.

- Your convert-and-send function is actually producing a text error rather than a graphic image (possibly because you don't actually have support for that image type compiled into PHP).

- You have actually somehow screwed up your internal `gd` image well before trying to send it off. One very common cause of this is failure to allocate colors in a palette-based image or to use color indexes that haven't been allocated.

In our experience, the best way to debug these sorts of problem is simply to comment out the PHP statement that sends the header announcing an image, and then look at the output as text in your browser. If everything were working perfectly, you would expect to see your binary image data as text, which would mean a lot of strange-character gibberish, possibly starting with a short amount of recognizable text (like `PNG`). If you see a PHP warning or error instead of the image data or in addition to the image data, you can proceed to debug that. If you see nearly nothing, not much of an image is being sent — this may imply problem #2 or #4 above. If you see what looks like a reasonable amount of pure image data, you may need to look at your code very carefully for small amounts of text (like spaces) that you may have introduced.

Summary

If you create on-the-fly graphic images by using PHP, you're creating something completely different from the usual HTML that PHP generates — a completely different format, and a completely different look. Although there are hassles associated with getting the `gd` image library working, after you get past those you have quite a rich set of image-manipulation functions to work with. You can create web pages that are all image, or pages that have `<IMG>` tags that link to dynamic images, or you can start building a library of image files for display later on. Either way you go, you have a richer vocabulary to work with than with pure HTML, and (even if many situations don't require dynamic images) you have another type of tool in your kit.

Part VI

Case Studies

Chapter 43

Developing a Weblog with PHP

S mall standalone PHP applications, such as polls and e-mail forms, are all very useful, but complete content sites are where PHP really shines. Here we give complete instructions for developing the simplest type of standalone site, which is the weblog.

Why Weblogs?

A *weblog* is the simplest kind of dynamic site. It can be thought of as a dynamic version of the personal home page: a content site organized by chronology with frequently updated posts. Most weblogs do not create all their own content in the sense of writing full news stories or producing a trove of artwork; they instead exist to comment on other people's content and events of the day or to provide a venue for personal thoughts and reflections. On the high end of the genre, public weblogs like Slashdot can become extremely popular meeting places for online communities to chew the fat of their common interests.

If you are a newcomer to server-side scripting, we encourage you to immediately start a personal weblog as your first major project. Nothing helps you learn faster than running an actual complete site of your own, where you can try out a range of new techniques and ideas in context. Especially because PHP and other open source technologies grow and change so quickly, it's well-nigh essential to have a preexisting testbed always available to doodle around on.

Weblogs are also just fun and therefore, worthwhile even for those who also use PHP in more serious contexts. There's no pleasure quite like that of conducting an intellectual debate, an argument, or a romance by weblog. Forget movies, pop music, and reality TV — the weblog is the true medium of the age, baby!

715

The Simplest Weblog

The main goal of this section is to introduce you to the layout and display aspects of building a dynamically generated site. In later sections, we will refine our techniques for handling the data-related aspects. At the end of this chapter, you should have the ability to make and maintain a simple data-driven site of your own. Conceptually, even the most complicated dynamic content sites are basically just bigger versions of the concepts you will learn by building a personal weblog.

The easiest weblog is just a PHP template and some included text files. It's limited to local development only — in other words, you won't be able to make entries via HTTP but only by creating text files while logged in to the PHP server as a trusted user (or copying text files to the server via some mechanism like scp (secure copy), which amounts to the same thing). You also won't be able to assign different levels of permissions very effectively, so this style of weblog is most appropriate for a purely personal single-author site.

We decided to use the most basic type of navigation, Previous and Next text links that we'll maintain by hand. This gives you the maximum flexibility to decide how often you want to change the front page of your weblog — we'll do it daily, but you may prefer a weekly, monthly, or irregular changeover, depending on how much you have to write about. We'll also include an old reliable left-side navbar with links to standalone pages, such as About Me and Favorite Things information. A finished weblog page is shown in Figure 43-1.

FIGURE 43-1

A weblog page
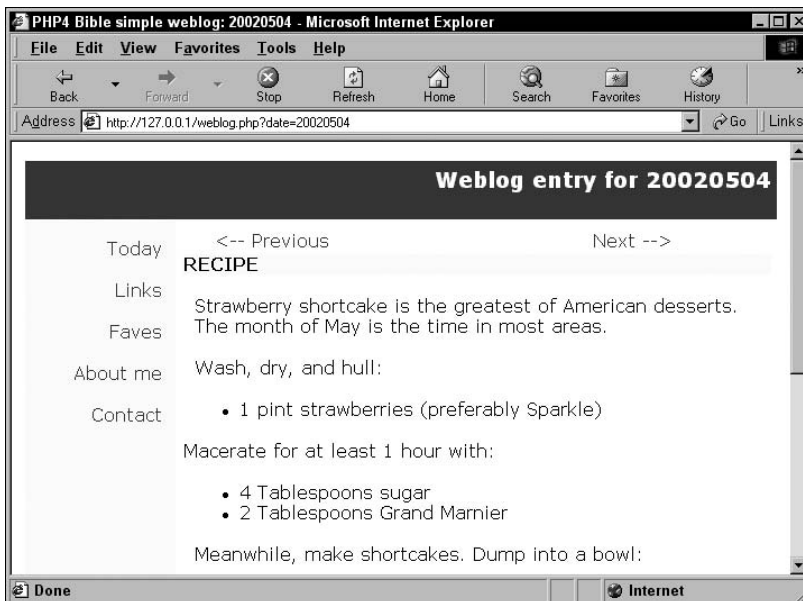
It is assembled from these files:

- `weblog.php`: Main display page template
- `20040101.txt`, `20040102.txt`: Weblog entries (changed daily)
- `default.txt`: Default text entry for days when there is no new content
- `favorites.php`, `links.php`, `aboutme.php`: Semistatic pages (changed infrequently)
- `header.inc`: Header and navigation bar on every page
- `footer.inc`: Footer on every page
- `style.inc`: Internal style sheet

CAUTION You must change the variable `$initial_entry_date` in `weblog.php` to the date of your first entry, or you may start an infinite loop that will eat up all your server cycles! You must also check all the paths to included files and change them to real paths.

Listings 43-1 through 43-7 are the code for a simple weblog. Instead of using a database to store your entries, the data will be stored in text files on your filesystem.

LISTING 43-1

Main weblog template (weblog.php)

```php
<?php

// ------------------------------
// GET YOUR VARIABLES ALL LINED UP
// ------------------------------
// Change this to the date of your first log entry.
$initial_entry_date = 20040101;

// Replace the fake path below with a real one
$DOCUMENT_ROOT = "c:\docs";

$today = date('Ymd');
if (isSet($_GET['date'])) {
  if ($_GET['date'] < $initial_entry_date) {
    // Go to first entry if the specified date is earlier
    // than range
    $date = $initial_entry_date;
  } elseif ($_GET['date'] > $today) {
    // Go to last entry if specified date is later than range
    $date = $today;
  } else {
    $date = $_GET['date'];
  }
} else {
  $date = $today;
}
```

```
$title_msg = $date;
$header_msg = "Weblog entry for $date";

// Assemble the Previous/Next links
$prevdate = $date - 1;
$nextdate = $date + 1;
if ($date == $initial_entry_date) {
  $flipbar = "\n<P CLASS=\"next\">
<A HREF=\"weblog.php?date={$nextdate}\">Next --&#62;</A>
</P>\n";
} elseif ($date == $today) {
  $flipbar = "\n<P CLASS=\"previous\">
<A HREF=\"weblog.php?date={$prevdate}\">&#60;-- Previous</A>
</P>\n";
} else {
  $flipbar = "\n<TABLE BORDER=0><TR>
<TD WIDTH=\"50%\" ALIGN=\"left\">
<SPAN CLASS=\"previous\">
<A HREF=\"weblog.php?date={$prevdate}\">&#60;-- Previous</A>
</SPAN>
</TD><TD WIDTH=\"50%\" ALIGN=\"right\">
<SPAN CLASS=\"next\">
<A HREF=\"weblog.php?date={$nextdate}\">Next --&#62;</A>
</SPAN>
</TD>
</TR></TABLE>\n";
}


// ---------------------
// NOW ASSEMBLE THE PAGE
// ---------------------
include_once('header.inc');

echo $flipbar;
// Include the specified text file, or a default message
// Replace the fake path below with a real one
if (file_exists($DOCUMENT_ROOT."/path/to/entries/$date.txt")) {
  // Replace the fake path below with a real one
  include($DOCUMENT_ROOT."/path/to/entries/$date.txt");
} else {
  include("default.txt");
}
echo $flipbar;

include_once('footer.inc');
?>
```

LISTING 43-2

A dated entry (20000101.txt)

```
<DIV CLASS="topic">HOLIDAY</DIV>
<P>Oh, what a holiday season it has been! I am positively
stuffed with fruitcake. </P>

<P>My New Year's Resolutions are:
<UL>
<LI>Trade in AMC Gremlin.</LI>
<LI>Contribute to OSS project.</LI>
<LI>Take full 2 weeks vacation (dude ranch?).</LI>
<LI>Be less snide.</LI>
</UL>
</P>
```

LISTING 43-3

Default message (default.txt)

```
<P>Sorry, nothing new today!  Check back tomorrow.</P>
```

LISTING 43-4

A static page (favorites.php)

```php
<?php
$title_msg = 'favorites';
$header_msg = 'My favorite things';

include_once('header.inc')
?>

<P>These are a few of my favorite things.</P>

<DIV CLASS="topic">BOOKS</DIV>
<DL>
<DT>Cryptonomicon, by Neal Stephenson</DT>
<DD>The techie masterpiece--it's our life, put in the blender
of a massive inventiveness. Be sure to also download the essay
"In the beginning was the command line" from his site,
www.crytonomicon.com .</DD>
</DL>
```

```
<DIV CLASS="topic">MUSIC</DIV>
<DL>
<DT>Raw Power, by The Stooges</DT>
<DD>See who all those neo-punk bands are copying.</DD>
</DL>
<?php include_once('footer.inc'); ?>
```

LISTING 43-5

Included header file (header.inc)

```
<HTML>
<HEAD>
<TITLE>PHP4 Bible simple weblog:
<?php echo $_GET['date']; ?></TITLE>
<?php include("style.inc"); ?>
</HEAD>

<BODY BGCOLOR="#FFFFFF">
<TABLE BORDER="0" CELLPADDING="5" WIDTH="100%">
<!-- Title box -->
<TR WIDTH="100%" BGCOLOR="#822222">
  <TD WIDTH="100%" ALIGN="right" COLSPAN="2">
    <H1><?php echo $header_msg; ?></H1>
  </TD>
</TR>
<!-- End Title box -->

<!-- Begin main body -->
<TR WIDTH="100%">
  <TD WIDTH="20%" VALIGN="top" BGCOLOR="#FFFECC">
    <!-- Navbar -->
    <P CLASS="sidebar"><A HREF="weblog.php">Today</A></P>
    <P CLASS="sidebar"><A HREF="links.php">Links</A></P>
    <P CLASS="sidebar"><A HREF="favorites.php">Faves</A></P>
    <P CLASS="sidebar"><A HREF="aboutme.php">About me</A></P>
    <P CLASS="sidebar"><A
HREF="mailto:me@localhost">Contact</A></P>
    <!-- End Navbar -->
  </TD>
  <TD WIDTH="80%">
```

Included footer (footer.inc)

```
<!-- End of main body -->
</TD></TR>
</TABLE>
<P CLASS="footer">Copyright Wiley, Inc. 2000 - 2009</P>
</BODY>
</HTML>
```

Included stylesheet (style.inc)

```
<STYLE TYPE="text/css">
<!--
BODY, P, LI    {font-family: verdana, arial, sans-serif;
font-size: 12pt; color: #000000; text-align: left;
margin-left:10px}
H1      {font-family: verdana, arial, sans-serif;
font-size: 14pt; color: #FFFFFF}
A:link, A:visited  {font-family: verdana, arial, sans-serif;
font-size: 12pt; color: #822222; text-decoration:none}
.sidebar  {font-family: verdana, arial, sans-serif;
font-size: 12pt; color: #822222; text-align:right;
margin-top:10; margin-right:7}
.topic    {font-family: verdana, arial, sans-serif;
font-size: 12pt; font-weight: bold; color: #000000;
background: #FFFECC; text-align: left}
.footer  {font-family: verdana, arial, sans-serif;
font-size: 9pt; color: #808080; text-align:right}
.previous  {font-family: verdana, arial, sans-serif;
font-size: 12pt; color: #808080; text-align:left;
margin-left:25; margin-right:100}
.next  {font-family: verdana, arial, sans-serif;
font-size: 12pt; color: #808080; text-align:right;
margin-left:100; margin-right:25}
-->
</STYLE>
```

To use the simple weblog, place all the files in a PHP-enabled directory on your web server. Create a subdirectory for your daily entries (for example, `20000101.txt`, `20020504.txt`); otherwise, you'll quickly end up with dozens of files cluttering up your main directory. The files in this subdirectory need to be writable by you and readable by all.

721

When you're ready to make an entry, log in to your web server, fire up a text editor such as `vi`, and write an HTML-formatted text file for each day you want to post, naming it according to the date convention we've established. Alternatively, instead of logging in to your web server, you can write up your daily text file on a local client copy and then use `scp` to upload it to your web server. Obviously, you can edit this file however many times you like, if you have multiple things to say per day. As long as the files have the correct names, locations, and permissions, this code should run smoothly for you. This type of weblog is self-archiving, so you don't need to do anything special with old entries — they'll just stay around forever if you have a big enough hard disk.

> **NOTE**
>
> If you still use FTP to upload files, please take an hour to learn how to use `scp` instead. A fine command-line Windows client called `pscp` is available for free downloading at `www.chiark.greenend.org.uk/~sgtatham/putty`. **An even easier GUI Windows client called** `WinSCP` **is available for free downloading at** `http://winscp.sourceforge.net/eng`. **FTP should be used only for file downloading, for example from anonymous FTP servers, because it has caused so many security problems when uploading files. One of the reasons to avoid the otherwise fine weblog-publishing applications such as Movable Type and Blogger is that they rely on FTP to write files to your web server.**
>
> We promise that `scp` is just as easy to use, if not easier — instead of typing `ftp myserver.mydomain.com` and `put myfile.php`, **you combine both commands into the simple** `pscp myfile.php me@myserver.mydomain.com:myfile.php`.
>
> If you really can't make up your mind to learn `scp`, it might be safer to use an HTTP-based editing tool as detailed in the next section.

Adding an HTML-Editing Tool

This simple weblog is quite adequate for many purposes, but it has one big disadvantage: You can't write up your daily entries using the web itself. Instead, you must create each entry using a text editor such as `emacs` or Notepad and save it to your web server's `docroot`. This can be a significant issue over time, especially if you are not allowed ssh/FTP access to your server or aren't comfortable with the process. HTTP is the next logical step for many users, and it is probably no less unsafe than using FTP.

This process has one big problem: You need to give read/write permissions to the HTTP user (usually Nobody) in a particular directory. This is an inherently insecure process, and we do not recommend it in the long run. We'll describe the HTTP tools here so that you can become comfortable with the new aspects before moving on to a better solution, which is using a database instead of separate `include()` files for each entry. We'll also try to keep the security problems to a minimum, employing a password and letting you send mail to yourself if an unauthorized person tries to log in.

The files you need for an HTML-based file-writing tool are:

- `login.php`
- `logentry.php`
- `logentry_handler.php`
- `password.inc`

Put `password.inc` in a directory outside the web tree, such as `/home/htmluser`. This will ensure that your passwords cannot be read via the web without being processed by PHP first. The directory must be world-executable and the document must be readable by the `httpd` user (Nobody). If you have root access on this server, you could `chown` it to belong to the `httpd` user; if not, you may have to make the file world-readable, which is a security breach. Be sure to use a password different from your system user password, just in case it's compromised.

Listings 43-8 through 43-10 are the files you need for an HTML form to edit weblog entries.

LISTING 43-8

Weblog entry login screen (login.php)

```
<HTML>
<HEAD>
<TITLE>Weblog login screen</TITLE>
</HEAD>

<P><B>Supply a username and password.</B></P>
<FORM METHOD=POST ACTION="logentry.php">
<P>USERNAME:
<INPUT TYPE=TEXT NAME="test_username" SIZE=20></P>
<P>PASSWORD:
<INPUT TYPE=PASSWORD NAME="test_password" SIZE=20></P>
<P>BLOG ENTRY:<BR>
<TEXTAREA NAME="logtext" COLS=75 ROWS=20WRAP="VIRTUAL">
</TEXTAREA></P>
<P><INPUT TYPE="SUBMIT" VALUE="SUBMIT">
</FORM>
</BODY>
</HTML>
```

LISTING 43-9

Included password file (password.inc)

```
<?php
$username = "logwriter";
$password = "logpass";
?>
```

LISTING 43-10

Weblog data entry script (logentry.php)

```php
<?php
$date = date("Ymd");
include("/home/htmluser/password.inc");

if($_POST['test_username'] == $username &&
  $_POST['test_password'] == $password) {
  $fp = fopen("/entries/$date.txt", "w");
  $try_entry = fwrite($fp, $_POST['logtext']);
  if ($try_entry > -1) {
    print("Weblog entry for $date written to disk.");
  } elseif ($try_entry == -1) {
    print("Weblog entry write failed.");
  }
} else {
  mail("me@localhost", "Weblog snoop alert", "Someone from
$REMOTE_ADDR is trying to get into your weblog entry
handler.");
}
?>
```

Changes and Additions

Things you might want to immediately change, add, or alter in this codebase include:

- Alter colors, styles, layout.

- Allow database connectivity.

- Change frequency of expected updating (weekly, monthly).

- Change to calendar-based navigation rather than Next/Previous links.

- Change to topic-based rather than date-based navigation.

- Stop automatic entry changeover by date.

- Allow future entries in database.

- Allow multiple authors/editors with different permissions.

Besides a personal weblog, you could use this code for any simple, chronological note taking, such as:

- A vacation journal
- A project log
- The story of your vast weight loss through heroic diet and exercise
- A chronicle of your pregnancy and your baby's development

Summary

Although it's handy for small, standalone projects such as polls, PHP's most impressive use is in developing complete data-driven content sites. The easiest such site to develop is the personal weblog. We encourage every PHP user to keep one, if only as a handy testbed for new ideas and techniques.

If you wish, you can store your data in ordinary text files, using PHP to plug these files into a template based on a criterion such as date. This will save a certain amount of formatting-related repetition at the cost of somewhat decreased security. Far better in every way is to keep the data in a database.

The weblog format is very flexible. It can scale up to a major public site like Slashdot, with tens of thousands of contributors and a steady stream of new content upon which to comment. Or you can keep a little secret diary on your own laptop, reading it in a browser window on the sly. The important point is that once you've made a complete data-driven site with PHP, you'll never go back to static web pages.

Chapter 44

A Trivia Game

I n this chapter, we present a full working example of a small PHP application: a web-based trivia game with a twist (the "Certainty Quiz"). The main virtue of the chapter is its completeness: Instead of using code fragments to illustrate talking points, as we do in most other chapters, we're showing everything, soup to nuts. As a result, this is one of the larger examples in the book, weighing in at more than 1300 lines of PHP code.

Concepts Used in This Chapter

The code in this chapter uses a wide variety of techniques, tricks, and technologies that we've presented elsewhere in the book. In particular:

- We make heavy use of the object-oriented features of PHP (Chapter 20).

- We rely on PHP's session mechanism to propagate game data from page to page (Chapter 24).

- We use a back end database (MySQL) to store questions and high scores (Part II).

- We do some behind-the-scenes mathematics, including approximating nth roots (Chapters 9 and 27).

- We use arrays for storing data and for manipulating data returned from the database (Chapters 8 and 21).

- We do a lot of string processing and concatenation to build our display pages, including the heredoc technique for templating pages (Chapters 7 and 21).

- We use the new exception-handling features of PHP5 to catch database and session problems (Chapter 31).

We highlight some of these topics in various sections later in the chapter as we delve into the code.

The Game

Several years ago, a friend asked one of us to try a quiz he'd seen somewhere on the Internet. After I agreed, he told me that he would ask me 10 questions, each of which had a numerical answer (dates, weights, lengths, counts, and so on). The unusual part was that instead of answering with a number, I was to give a lower bound and an upper bound on the answer. I could make the ranges as large as I wanted, and otherwise I had only one instruction: Make sure that you answer nine out of ten questions correctly.

I answered the questions confidently and was surprised at the end to find that my final score was six (or was it four?). At any rate, I did surprisingly badly, but my friend said that everyone else he had tried it on had done even worse. Now, how could anyone lose such an easily winnable game? After all, when asked when Shakespeare was born, I could have said "Sometime between 30,000 B.C. and A.D. 30,000" and been pretty sure that I would be right. What trips people up seems to be some combination of pride and overconfidence. The pride prevents you from giving a ridiculously large range (because then your questioner knows you don't have the foggiest idea when Shakespeare was born); the overconfidence makes you willing to narrow the range beyond your real range of certainty. In the end, the game isn't testing your knowledge — it's testing your knowledge of your own knowledge (or lack of knowledge).

Our version

In this chapter, we implement something like this quiz game, but with some changes to make it more web-friendly. For one thing, rather than having the player type in numbers freely, we present a range of choices that the player narrows down further. For another, we don't rely on pride to make the ranges narrow (because people may end up playing this over the web in the privacy of their own home). Instead we add incentives to the scoring system to make people guess narrowly rather than broadly. Finally, we add some features familiar from online games, such as levels of difficulty and a list of top scorers.

The upshot is a game that, while it may or may not be fun, is certainly frustrating, which for many people is nearly as good.

Sample screens

Figure 44-1 shows the game screen as it may look to a new arrival. There is a welcome message to the right, and a question to the left, with radio buttons for choosing a range of answers.

FIGURE 44-1

Start screen
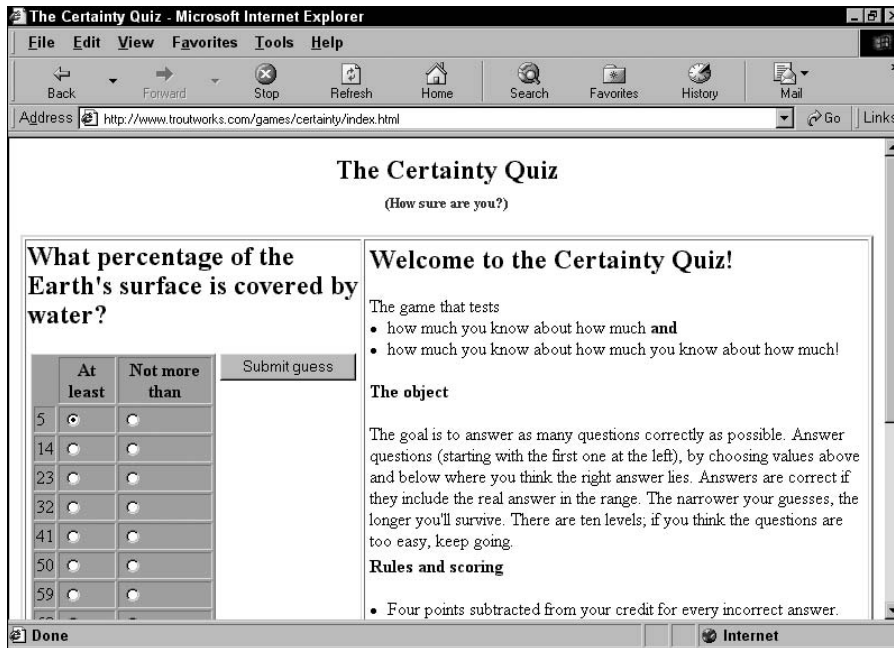
Figure 44-2 shows the screen immediately after the player has answered the first question. Another question is offered on the left, and now the state of the game score is highlighted on the right, showing the correct answers to date, the credit remaining, and the level attained. (See the next section for an explanation of what these things mean.)

Finally, Figure 44-3 shows "Game Over," complete with taunting message and a list of high scorers. (There's a corresponding "Game Won" screen in the unlikely event that the user survived all the questions the game could come up with.)

The rules

The basic play cycle is simple: The player is asked a question that requires a numerical answer, and the player responds by choosing a range of values that should include the answer. The goal is to answer as many questions correctly as possible, while surviving in the meantime. Survival depends on credit, which is accumulated by answering questions correctly within a narrow range and is spent by giving wrong answers or answering questions too broadly.

FIGURE 44-2

Continuing play
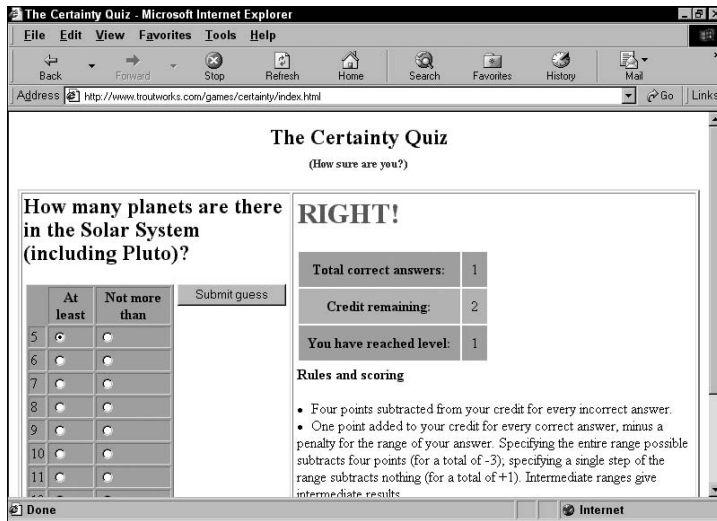
FIGURE 44-3

Game over
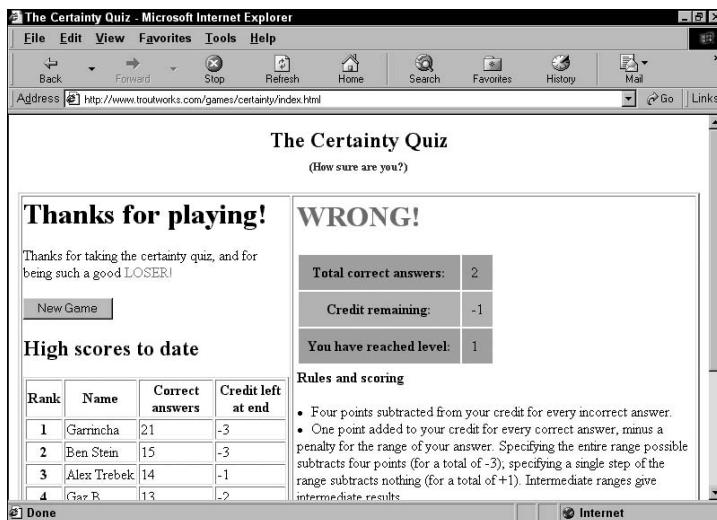
The exact rewards and penalties are easily tweakable in the code, but in this chapter's version they are:

■ **Correct answers**: One point added to credit, minus a penalty for the size of the range specified. The penalty ranges from zero for answers that use only one step of the possible range, up to four points for making the range as wide as possible.

■ **Incorrect answers**: Four points deducted from credit.

Credit starts at five points and can rise only as high as fifteen points. The game is over when credit goes below zero.

It is easy to pass by simply submitting your answer without making a choice, since the radio buttons are set to specify the widest possible answer range unless the player changes them. The penalties are set up so that passing is costly (a total of 1 - 4 = -3 points), but not as costly as guessing wrong (-4 points). The player is better off narrowing the range as much as possible, while still being sure that the real answer is still included.

Playing the game yourself

If you just can't wait to play, we have a playable version of the game available at `www.braingia.org/books/pmb/trivia`.

The Code

The code for this example is almost completely written in an object-oriented style (see Chapter 20 for an introduction to PHP's version of object-oriented programming). Among the classes we define are:

■ `Question`: Each `Question` object includes the text of the question, the correct answer, the lower and upper bounds that are presented to the player, and enough information to display the range of choices that the player can choose from. In addition, `Question` instances track whether or not they are answered correctly.

■ `Game`: There should be one and only one `Game` object in existence at a particular time. Game objects may include up to two question objects (the current question and the previous one), as well as a `GameParameter` instance.

■ `GameParameters`: Contains all the numerical settings that affect how the game behaves and manages some globally available resources such as the database connection.

■ `GameDisplay`: Contains a `Game` instance as a component and does all the work of actually displaying HTML and receiving input. Also contains an instance of `GameText`.

■ `GameText`: A repository for boilerplate HTML that is not dependent on any knowledge of the state of the game. Only this class and `GameDisplay` actually have HTML code in them.

Code files

The code files include definitions for all the classes in the preceding section: `question_class.php`, `game_class.php`, `game_parameters_class.php`, `game_display_class.php`, and `game_text_class.php`. In addition, there are some code files that don't define classes:

- `index.php`: The first file loaded, which handles sessions and post arguments and creates the `GameDisplay` object.

- `certainty_utils.php`: A grab bag of initialization statements (seeding the random-number generator, for example) and math utility functions.

- `entry_form.php`: A form for adding new questions to the database.

- `dbvars.php`: The usual file with definitions for username, password, and host for the database connection.

We now take a tour of the code file listings. Rather than building from the ground up as we sometimes do, in this chapter we work from the top down: first the very first page that is actually loaded, then the code that page depends on, and so on until we bottom out in utility functions and database calls. Finally, at the end of the chapter, we show how to construct the database and populate it with questions.

index.php

Listing 44-1 shows `index.php`, which is the user's entry point. The primary job of this file is to determine where we are in the cycle of play, set up the appropriate PHP objects (either by creating them or by retrieving them from the user's session), and echo out the display code that the objects generate.

Where we are in the course of a game is determined by a combination of session and POST information; as users arrive for the first time, they find neither a current session nor any POST arguments. Successive pages, however, should have both an active session and useful information submitted from the previous page.

LISTING 44-1

index.php

```php
<?php
// Include code files, start up session
include_once("certainty_utils.php");
include_once("game_display_class.php");
session_start();

// Determine state and handle post arguments

try {

  // CASE 1: Player is submitting name for high score list
  if (get_session_value('game') &&
```

```
      get_post_value('HIGHSCORE')) {
    if (get_session_value('game') &&
      get_post_value('HIGHSCORE')) {
      $game_display =
      new GameDisplay(get_session_value('game'));
                  $game_display->handleHighScore();
    }
  }

  // CASE 2: Player is in middle of game that we are tracking
  elseif (get_session_value('game') &&
          !get_post_value('NEW')) {
    $lower = get_post_value('lower');
    $upper = get_post_value('upper');
    $game_display =
      new GameDisplay(get_session_value('game'));
    $game_display->updateWithAnswer($lower, $upper);
  }

  // CASE 3: Player has either just arrived or has
  //    finished a game and asked for a new one.
  elseif (!get_post_value('POSTCHECK') ||
          get_post_value('NEW')) {
    $game_display = new GameDisplay(new Game());
  }

  // CASE 4:  Something is wrong.
  // The page is the result of a POST operation,
  // yet we don't seem to have a live session, so
  // we are not successfully tracking a game.
  // The only thing to do is complain, and ask about
  // cookies.
  else {
    $game_display =
      new GameDisplay(new Game());
    throw (new Exception("We couldn't track your game." .
          "You may have to enable cookies to play"));
  }

  // Construct string that will be displayed as page
  $page_string = $game_display->display();
  // Store game state in session so that next
  //   page can pick it up
    set_session_value('game',
      $game_display->_game);

} // end of try block

catch (Exception $exception) {
  // There is a problem somewhere.  Create
```

```
  //   an error page.
  $exception_msg = $exception->getMessage();
  $display = new GameDisplay(null);
  $page_string =
    $display->makeErrorPage($exception_msg);
  // hope to start fresh next time
  unset_session_value('game');
}

// Actually echo page to browser

  echo($page_string);

?>
```

The object types that index.php cares about are GameDisplay and Game. The Game object contains all the state information that needs to be preserved from page to page about where the user is in the game (score, questions asked already, and so on). The GameDisplay object contains a Game instance and does everything necessary to produce an HTML page from it.

If the user is starting off for the first time, we create a new Game object, relying on the object's constructor to initialize it appropriately. For subsequent pages, though we rely on the automatic object serialization feature of PHP sessions to store the Game object for us. (The actual definitions of get_session_value() and set_session_value() are in certainty_utils.php, but all that is happening is that we stash the Game object in a session. PHP takes care of the serialization that is necessary to read the object into a session and back out again.)

CROSS-REF   **For more on what it means to serialize an object, see Chapter 20; for an explanation of sessions and their workings, see Chapter 24.**

If the user is in the middle of a game, we expect both a Game object stored in the session and a form submission representing either a guess at the answer or a request to be listed on the High Scores page.

Regardless of whether we create a new Game object or retrieve the one from the last page, we create a new GameDisplay object around it and then ask that object for a string that represents the entire HTML page. We store this in a string, ready to echo it out to the browser in the very last code line.

Exceptions

Many things can go wrong with the execution of this game's code. For one thing, of course, there's always a possibility of a code bug that leaves the game in a strange state. In addition, though, the code relies on at least three external "services," any of which might misbehave:

■ The database, which stores the questions and answers

■ The session mechanism, which is in turn probably relying on files on the hard disk

■ Cookies stored on the user's browser, which may refuse them

If any of these services turn out to be unreliable, the game will not be playable. Our goal in this situation should be to fail as gracefully as possible.

In a previous edition of this chapter, the code had to catch all the possible failures and propagate an error up to this page, which would then detect the failure and display an appropriate error string. This time, though, we can use PHP's exception mechanism, which makes it much easier to structure the code. Whenever we encounter a problem that cannot be recovered from, we `throw` the problem, along with a descriptive string. The `catch` statement in `index.php` is the only one in the game's code, and so will receive any of the exceptions that happen as a result of its calls to functions in other code files. In addition, it will catch the exception thrown in Case 4 of `index.php`, which probably indicates that the user's browser is not accepting cookies.

CROSS-REF **See Chapter 31 for an introduction to exceptions and error handling.**

game_display_class.php

Almost all of the look-and-feel action for this game is in `game_display_class.php`, as shown in Listing 44-2.

The code file depends on two other files: `game_class.php` and `game_text_class.php`. The former contains most of the logic for the inner workings of the game, whereas the latter just contains some boilerplate text. The job of the `GameDisplay` class is to extract all the information from the game state necessary to produce actual HTML pages.

The important public functions in the class are:

- The constructor function
- `updateWithAnswer()`, which is called with data from the user's submission of a guess
- `makeErrorPage()`, which returns HTML to display if something has gone wrong
- `display()`, which returns HTML to display when everything has gone right

LISTING 44-2

game_display_class.php

```php
<?php
include_once("game_class.php");
include_once("game_text_class.php");

class GameDisplay
{
  //    presentation
  public $_pageTitle = "The Certainty Quiz";
  public $_blueColor = "#AAAAFF";
  public $_redColor = "#FFAAAA";

  //    contents
```

```
  public $_game = NULL;
  public $_gameText;
  public $_highScorePosted = FALSE;

  // CONSTRUCTOR

  function __construct ($game) {
    $this->_game = $game;
    $this->_gameText = new GameText();
  }

  // PUBLIC FUNCTIONS
  //   accessors
  function getPageTitle() {
    return($this->_pageTitle);
  }
  function getBlueColor() {
    return($this->_blueColor);
  }
  function getRedColor() {
    return($this->_redColor);
  }
  function getGame() {
    return($this->_game);
  }
  function getHighScorePosted() {
    return($this->_highScorePosted);
  }

  function updateWithAnswer ($lower, $upper) {
    $game = $this->getGame();
    $game->updateWithAnswer($lower, $upper);
  }

  function makeErrorPage ($problem_string) {
    // constructs the HTML page to display when
    /// something has gone horribly wrong
    $top_matter_string =
        $this->_makeTopMatter($this->_pageTitle);
    $page_string = <<<EOT
$top_matter_string
</H2></CENTER>
<TABLE BORDER=2>
<TR>
<CENTER>
<H2>We're sorry, but the game is not available
right now.</H2><H4>($problem_string)</H4>
</CENTER>
</TR></TABLE>
```

```
</BODY></HTML>
EOT;
    return($page_string);
  }

function display () {
    // returns entire page as string ---
    // backbone structure of page, plus
    // overridable methods to print components

    // sanity checks
    if (!$this->_game ||
        !is_object($this->_game)) {
      throw new
        Exception("Cannot find valid game object");
    }
    elseif (!$this->_game->getDbConnection()) {
      throw new
        Exception("No database connection");
    }

    // display of apparently valid page
    else {
      $top_matter_string =
        $this->_makeTopMatter($this->_pageTitle);
      $current_question =
        $this->_currentQuestionString();
      $previous_question =
        $this->_previousQuestionString();
      $game_state =
        $this->_gameStateString();
      $introduction =
        $this->_gameText->introduction();
      $rules =
        $this->_gameText->rules();
      if ($this->_game->getGameLost()) {
        $left_side =
          $this->_gameText->gameLostText() .
          $this->_highScoreString();
      }
      elseif ($this->_game->getGameWon()) {
        $left_side =
          $this->_gameText->gameWonText() .
          $this->_highScoreString();
      }
      else {
        $left_side = $current_question;
      }
      if ($this->_game->getPreviousQuestion()) {
```

```
        $right_side =
          "<TABLE><TR><TD>
          $previous_question
          </TD></TR><TR><TD>
          $game_state
          </TD></TR><TR><TD>
          $rules
          </TD></TR></TABLE>";
      }
      else {
        $right_side =
          "<TABLE><TR><TD>
          $introduction
          </TD></TR>
          <TR><TD>
          $rules
          </TD></TR><TR><TD>
          $game_state
          </TD></TR>
          </TABLE>";
      }

// actually construct page
$page_string = <<<EOT
$top_matter_string
</H2></CENTER>
<TABLE BORDER=2>
<TR>
<TD VALIGN=TOP WIDTH=40% >$left_side</TD>
<TD VALIGN=TOP WIDTH=60% >$right_side</TD>
</TR></TABLE>
</BODY></HTML>
EOT;
    return($page_string);
  }
 }

 function handleHighScore () {
   // Handles database update for case where player
   //  has earned high score, and has submitted
   //  a name for the record
   if (!$this->_highScorePosted) {
     $this->_highScorePosted = TRUE;
     if (get_post_value('NICKNAME') &&
         get_post_value('ANSWER_COUNT') &&
         get_post_value('CREDIT') &&
         get_post_value('CHECKSUM') &&
         $this->_checksumChecks(
           get_post_value('ANSWER_COUNT'),
```

```
        get_post_value('CREDIT'),
        get_post_value('CHECKSUM'))) {
    $name = get_post_value('NICKNAME');
    $answer_count = get_post_value('ANSWER_COUNT');
    $credit = get_post_value('CREDIT');
    $connection =
     $this->_game->gameParameters->getDbConnection();
    $safe_name = mysql_real_escape_string($name);
    $safe_ans = mysql_real_escape_string($answer_count);
    $safe_credit = mysql_real_escape_string($credit);
    $query = "insert into high_scores
              (name, answer_count, credit)
              values
              ('$safe_name', '$safe_ans', '$safe_credit')";
    $result = mysql_query($query, $connection);
  }
  else {
    // do nothing--failure to add high score
    // should not be a deal killer
    // really don't even need the else condition
    // but if you do, here's where it would go.
  }
  }
}

// PRIVATE FUNCTIONS

private function _makeTopMatter ($title) {
  // returns HTML fragment that heads both
  //  regular page and error page, containing
  //  HTML head and title
  $return_string = <<<EOT
<HTML><HEAD><TITLE>$title</TITLE></HEAD>
<BODY BGCOLOR=#FFFFFF><CENTER>
<H1>$title<BR>
<FONT SIZE=-1 COLOR=BLUE>
(How sure <B>are</B> you?)</FONT>
</H1></CENTER>
EOT;
  return($return_string);
}

private function _currentQuestionString () {

  return("<H2>" .
         $this->_game->getCurrentQuestionText() .
         "</H2>" .
         "<FORM METHOD=POST ACTION=\"index.php\">" .
         "<INPUT TYPE=HIDDEN NAME=POSTCHECK VALUE=1>" .
```

739

```
            $this->_distractorString(
                $this->_game->getCurrentQuestion()) .
            "</FORM>");
    }

    private function _distractorString ($question) {
      // creates the actual HTML for presentation of
      //  radio-button alternatives for guesses.
      //  Assumes that the array representing the
      //   actual alternatives has been calculated in
      //   advance, retrievable from the question using
      //   getDistractorArray
      $distractor_array = $question->getDistractorArray();
      $distractor_string = "<TABLE><TR VALIGN=TOP><TD>";
      $distractor_string .=
        "<TABLE BORDER=1 BGCOLOR=\"#AAAAFF\"><TR><TH>
          </TH><TH>At least</TH><TH>Not more than</TH>";
      $count = 1;  // 1-based labels are preferable,
                   // so we can just use if ($label) ...
      $total = count($distractor_array);
      foreach ($distractor_array as $distractor) {
        $lower_selected = ($count == 1) ?
          "CHECKED" : "";
        $upper_selected = ($count == $total) ?
          "CHECKED" : "";
        $formatted_distractor =
          ($distractor >= 10000) ?
            number_format($distractor) : $distractor;
        $distractor_string .=
          "<TR><TD>$formatted_distractor</TD>
              <TD><INPUT TYPE=RADIO NAME=\"lower\"
                  VALUE=$count
                  $lower_selected ></TD>\n" .
          "<TD><INPUT TYPE=RADIO NAME=\"upper\"
              VALUE=$count
              $upper_selected ></TD></TR>\n";
        $count++;
      }
      $distractor_string .= "</TABLE>";
      $distractor_string .= "</TD><TD>";
      $distractor_string .=
        "<INPUT NAME=\"Submit guess\" VALUE=\"Submit guess\"
              TYPE=SUBMIT>";
      $distractor_string .= "</TD></TR></TABLE>";

      return($distractor_string);
    }
```

```
private function _previousQuestionString () {
  if (!$this->_game->getPreviousQuestion()) {
    $return_string = "";
  }
  else {
    $return_string =
      $this->_game->previousQuestionCorrect() ?
        $this->_rightString() :
        $this->_wrongString();
  }
  return($return_string);
}

function _rightString () {
  return("<H1><FONT COLOR=GREEN>RIGHT!</FONT></H1>");
}

function _wrongString () {
  return("<H1><FONT COLOR=RED>WRONG!</FONT></H1>");
}

private function _highScoreEligible () {
  // takes a game-ending score, and queries the
  //  DB to see if the player is eligible for the
  //  high score list
  $query = "select name, answer_count, credit
            from high_scores
            order by answer_count desc, credit desc
            limit 10";
  $connection =
    $this->_game->getDbConnection();
  if ($connection && is_resource($connection)) {
    $result = mysql_query($query, $connection);
    $eligible = false;
    if (mysql_num_rows($result) > 9) {
      while ($row = mysql_fetch_assoc($result)) {
        $answer_count = $row['answer_count'];
        $credit = $row['credit'];
        if (($this->_game->getCorrectAnswers()
             > $answer_count) ||
          (($this->_game->getCorrectAnswers()
              == $answer_count) &&
            ($this->_game->_credit > $credit))) {
          $eligible = TRUE;
          break;
        }
      }
    }
```

```
      else {
        $eligible = TRUE;
      }
      return($eligible);
    }
    else {
      throw new
        Exception("Game display has no database connection");
    }
  }

  // Checksum is calculated when posting a score
  //  (comprised of a number of correct answers plus
  //  credit remaining) and the checksum is compared with
  //  the submitted scores.  A first line of defense against
  //  spoofing (unless, of course, the checksum scheme is
  //  published in a book or something).
  private function _checksumChecks ($answer_count, $credit,
                          $checksum) {
    return($checksum ==
          $this->_makeCheckSum($answer_count, $credit));
  }

  private function _makeChecksum ($answer_count, $credit) {
   return ((round($credit)) * 17) *
          ($answer_count * 31);
  }

  private function _postHighScoreString () {
    // The greeting plus HTML form for actually submitting
    //  a name to the high scores list
    global $PHP_SELF;
    $answer_count = $this->_game->getCorrectAnswers();
    $credit = $this->_game->getCredit();
    $checksum = $this->_makeChecksum($answer_count, $credit);
    $result_string =
      "<H2>Congratulations! You have a high score</H2>".
      "Enter your name (or a nickname) for the high ".
      "scores list:".
      "<FORM METHOD=POST ACTION=\"$PHP_SELF\" >".
      "<INPUT NAME=NICKNAME TYPE=TEXT SIZE = 30>".
      "<INPUT NAME=ANSWER_COUNT TYPE=HIDDEN ".
        "VALUE=$answer_count>".
      "<INPUT NAME=CREDIT TYPE=HIDDEN ".
        "VALUE=$credit>".
      "<INPUT NAME=CHECKSUM TYPE=HIDDEN ".
        "VALUE=$checksum>".
      "<INPUT NAME=Submit TYPE=SUBMIT ".
        "VALUE=Submit >".
```

```
      "<INPUT TYPE=HIDDEN NAME=POSTCHECK VALUE=1>" .
      "<INPUT TYPE=HIDDEN NAME=HIGHSCORE VALUE=1>" .
      "<FORM>";
    return($result_string);
  }

  private function _highScoreString () {
    // The table of high scores itself, including
    //  the database interaction necessary to retrieve it
    if ($this->_highScoreEligible() &&
        !$this->_highScorePosted) {
      $result_string = $this->_postHighScoreString();
    }
    else {
      $result_string = "";
    }
    $result_string .=
     "<H2>High scores to date</H2>".
     "<TABLE BORDER=1><TR><TH>Rank</TH>".
     "<TH>Name</TH><TH>Correct answers</TH>".
     "<TH>Credit left at end</TH></TR>";
    $query = "select name, answer_count, credit
              from high_scores
              order by answer_count desc, credit desc
              limit 10";
    $connection =
      $this->_game->gameParameters->getDbConnection();
    if ($connection && is_resource($connection)) {
      $result = mysql_query($query, $connection);
      $rank = 1;
      while ($row = mysql_fetch_assoc($result)) {
        $name = $row['name'];
        $answer_count = $row['answer_count'];
        $credit = (int) ($row['credit']);
        $result_string .=
          "<TR><TH>$rank</TH><TD>$name</TD>".
          "<TD>$answer_count</TD><TD>$credit</TD></TR>";
        $rank++;
      }
      $result_string .= "</TABLE>";
      return($result_string);
    }
    else {
      throw new
        Exception("Game display has no database connection");
    }
  }

  private function _gameStateString () {
```

```
    // The HTML table
    $correct_answers = $this->_game->getCorrectAnswers();
    $credit = round_to_digits($this->_game->getCredit(), 2);
    $level = $this->_game->getLevel();
    return("<TABLE CELLPADDING=10>".
    "<TR BGCOLOR=$this->_blueColor>".
    "<TH>Total correct answers:</TH>".
    "<TD>$correct_answers</TD></TR>".
    "<TR BGCOLOR=$this->_redColor><TH>Credit remaining:</TH>".
    "<TD>$credit</TD></TR>".
    "<TR BGCOLOR=$this->_blueColor>".
    "<TH>You have reached level:</TH>".
    "<TD>$level</TD></TR></TABLE>");
  }
}
?>
```

Note that in the GameDisplay class we use some object-oriented constructs that were new as of PHP5. The constructor function is called __construct(), rather than having the same name as the class. And we have designated the functions that are not intended for external use as private, which will prevent any such use by other classes.

Most of the class's private functions involve querying the Game object for information that it then wraps up in HTML strings. One of the more interesting functions of this type is distractor_ string(), which creates the actual display of alternatives for the answer range. The general division of labor here is:

■ The upper and lower bounds for the answer are specified in the database, as well as how many choices should be displayed and how they should be scaled.

■ The Question object takes this information and creates all the intermediate steps of the answer range as it is constructed.

■ The GameDisplay object queries the game for the current question and then queries that question to discover the upper and lower bounds and the intermediate steps. It then simply wraps those values in HTML to present radio button alternatives, with the maximum answer range preselected.

game_text_class.php

Your humble authors try really hard to make this stuff interesting, but in this case, we must declare defeat. The GameText class just wraps up some boilerplate HTML into member functions so that the GameDisplay class can ask for it. Enough said?

The functions use our favorite technique for creating large chunks of boilerplate, which is the *heredoc* syntax. (See Chapter 7 for more on the uses of heredoc.) Listing 44-3 shows the game _text_class.php.

LISTING 44-3

game_text_class.php

```php
<?php
class GameText
{

function __construct () {
  // no vars, nothing for constructor to do
}

function introduction () {
$intro = <<<EOT
<H2>Welcome to the Certainty Quiz!</H2>
The game that tests
<LI>how much you know about how much <B>and</B> <LI>how much you
know about how much you know about how much!
<H4>The object</H4> The goal is to answer as many questions
correctly as possible.
Answer questions
(starting with the first one at the left), by choosing
values above and below where you think the right answer lies.
Answers are correct if they include the real answer in the
range.
The narrower your guesses, the longer you'll survive.
There are ten levels;  if you think the questions are
too easy, keep going.
EOT;
  return($intro);
  }

function rules () {
$rules = <<<EOT
<H4>Rules and scoring</H4>
<LI>Four points subtracted from your credit for every incorrect
answer.
<LI>One point added to your credit for every correct answer,
minus a penalty for the range of your answer.  Specifying the
entire range possible subtracts four points (for a total of -3);
specifying a single step of the range subtracts nothing
 (for a total of +1).  Intermediate ranges give intermediate re
sults.
<LI>Credit is capped at 15.
<LI>Whenever your credit falls below zero, the game is over.
EOT;
  return($rules);
}

function gameLostText () {
```

```
global $PHP_SELF;
$game_over = <<<EOT
<H1>Thanks for playing!</H1>
Thanks for taking the certainty quiz, and for being
such a good <FONT COLOR=RED>LOSER!</FONT><BR>
<FORM METHOD=POST ACTION="$PHP_SELF">
<INPUT TYPE=SUBMIT NAME=NEW VALUE="New Game">
<INPUT TYPE=HIDDEN NAME=POSTCHECK VALUE=1>
</FORM>
EOT;
return($game_over);
}

function gameWonText () {
global $PHP_SELF;
$game_over = <<<EOT
<H1>You won!</H1>
Thanks for taking the certainty quiz, and for
beating it.  We bow to your superior knowledge
of what you know, and what you don't know.
<FORM METHOD=POST ACTION="$PHP_SELF">
<INPUT TYPE=SUBMIT NAME=NEW VALUE="New Game">
<INPUT TYPE=HIDDEN NAME=POSTCHECK VALUE=1>
</FORM>
EOT;
return($game_over);
}

}
?>
```

game_class.php

In this section, we get to the basic logic of the game. The Game object contains everything worth remembering about the current state of the game, as well as methods for updating it.

Data members

It's worth listing the important pieces of data that the Game object tracks:

- The current question (an instance of class Question)
- The previous question, if any (an instance of class Question)
- The questions that have been asked at this level (an array of database IDs)
- The questions that could still be asked at this level (an array of database IDs)
- The game's numerical defaults (an instance of class GameParameters)
- Numerical variables that track the game's state (credit, questions answered, and so on)

Public functions

As with the `GameDisplay` class, let's list the functions that the `Game` class exposes to callers:

- The constructor function.

- Various accessor functions for member data.

- `updateWithAnswer()` takes the player's upper and lower guesses and updates the game's state accordingly, including both update of scores and setting up the next question to be asked.

Database interaction

The actual questions and answers that the game displays are retrieved from a backend MySQL database. There are two main types of interaction with that database:

- Whenever the player moves to a new level (including the first one), the `Game` object retrieves the IDs of all questions that may be asked at that level and scrambles their ordering. This randomized list is propagated along with the `Game` object from page to page within a particular level of the game.

- Whenever a new question is actually ready to be asked, the `Game` object pops its database ID off the list constructed and then queries the question database to retrieve all the rest of the question's information (the text of the question, the correct answer, the range of possible values to present, and so on).

Listing 44-4 shows `game_class.php`.

LISTING 44-4

game_class.php

```php
<?php
include_once("certainty_utils.php");
include_once("game_parameters_class.php");
include_once("question_class.php");

class Game
{

  public $currentQuestion = NULL;
  public $previousQuestion = NULL;
  public $gameParameters;

  public $_dbConnection = NULL;
  public $_credit = 0.0;
  public $_level;
  public array($_questionIdsAtLevel);  // an array of ids
```

```php
public $_questionsAskedAtLevel = 0; // a count
public $_totalQuestions = 0;
public $_correctAnswers = 0;
public $_gameLost = FALSE;
public $_gameWon = FALSE;

// CONSTRUCTOR
function __construct () {
  $this->gameParameters = new GameParameters();
  $this->_dbConnection =
    $this->gameParameters->getDbConnection();
  if (!$this->_dbConnection) {
     throw new Exception("No database connection");
  }
  else {
    $this->_correctAnswers = 0;
    $this->_level =
      $this->gameParameters->getStartingLevel();
    $this->_credit =
      $this->gameParameters->getStartingCredit();
    // make a list of questions to be asked at the
    //   starting level
    $this->_setupQuestionIds();
    // actually retrieve the first question
    $this->_installQuestion();
  }
}

// PUBLIC FUNCTIONS
// accessors
function getGameParameters()
  {return($this->gameParameters);}

function getCurrentQuestion()
  {return($this->currentQuestion);}

function getPreviousQuestion()
  {return($this->previousQuestion);}

function getCredit() {return($this->_credit);}

function getLevel() {return($this->_level);}

function getQuestionsAskedAtLevel()
  {return($this->_questionsAskedAtLevel);}

function getTotalQuestions()
  {return($this->_totalQuestions);}
```

```
function getCorrectAnswers()
  {return($this->_correctAnswers);}

function getGameLost()
  {return($this->_gameLost);}

function getGameWon()
  {return($this->_gameWon);}

function getCurrentQuestionText() {
  if (!is_object($this->currentQuestion)) {
    print("What is it?<BR>");
    print_r($this->currentQuestion);
  }
  else {
    return($this->currentQuestion->getQuestion());
  }
}

function previousQuestionCorrect() {
  return($this->previousQuestion->getCorrect());
}

function getDbConnection () {
  if (!$this->_dbConnection) {
    $this->_dbConnection =
      $this->gameParameters->getDbConnection();
  }
  return($this->_dbConnection);
}

function updateWithAnswer ($lower, $upper) {
  // The main modifying function for a game object.
  // Takes a player's upper and lower guess, determines
  //  correctness, updates scores, determines if the
  //  player has graduated to the next level, and
  //  swaps in the next question.
  $this->previousQuestion = $this->currentQuestion;
  $this->previousQuestion->updateWithAnswer($lower,
                                            $upper);
  $this->_updateScores();
  $this->_maybeChangeLevel();
  if (!($this->_gameLost || $this->_gameWon)) {
    $this->_installQuestion();
  }
}

// PRIVATE FUNCTIONS
function _installQuestion () {
```

```php
      // actually retrieve a question from the database
      //  and create a corresponding instance of Question
      if (count($this->_questionIdsAtLevel) > 0) {
        // pop a question off the randomized list
        $question_id =
          array_pop($this->_questionIdsAtLevel);
        $query =
          "select id, question, answer,
           upper_limit, lower_limit, scaling_type
           from question
           where id = $question_id";
        if (!$this->_dbConnection) {
          $this->_dbConnection =
            $this->gameParameters->getDbConnection();
        }
        if ($this->_dbConnection &&
            is_resource($this->_dbConnection)) {
            $result = mysql_query($query,
                              $this->_dbConnection);
          if ($row = mysql_fetch_assoc($result)) {
            $this->currentQuestion =
                new Question(
                  $row['id'],
                  $row['question'],
                  $row['answer'],
                  $row['lower_limit'],
                  $row['upper_limit'],
                  10,
                  $row['scaling_type']);
            $this->_questionsAskedAtLevel++;
          }
          else {
           throw new
      Exception("Problem retrieving question from database");
          }
        }
        else {
          throw new
            Exception("Problem querying question database");
        }
      }
      else {
        throw new
          Exception("Could not find any questions to ask");
      }
    }

    function _setupQuestionIds () {
```

```php
    $this->_questionIdsAtLevel =
      $this->_getQuestionIdsAtLevel($this->_level);
}

function _getQuestionIdsAtLevel ($level) {
  // to be used at time of graduation to a new level -
  // retrieves the new ids (only) of all questions at
  // the level, and shuffles them into a random order.
  $return_array = array();
  $query = "select id from question
            where level = $level";
  $this->getDbConnection();
  if (!$this->_dbConnection) {
    throw new
        Exception("No database connection");
  }
  else {
    $result = mysql_query($query,
                          $this->_dbConnection);
    while ($row = mysql_fetch_assoc($result)) {
      array_push($return_array, $row['id']);
    }
  }
  // randomize the order of the questions
  $return_array = create_randomized_array($return_array);
  return($return_array);
}

public function _updateScores () {
  // Change the current score based both on
  //  whether the player got the answer right and on
  //  the spread between the player's upper and lower
  //  guess.  Calculations depend on settings from
  //  the GameParameters class.
  if ($this->previousQuestion->rightAnswer()) {
    $this->_correctAnswers =
      $this->_correctAnswers + 1;
    $this->_credit +=
      $this->gameParameters->getRightAnswerCredit() -
      ($this->previousQuestion->getAnswerSpread() *
       $this->gameParameters->getAnswerSpreadDebit());
  }
  else {
    $new_credit =
    $this->_credit =
      $this->_credit -
      $this->gameParameters->getWrongAnswerDebit();
  }
```

```
      // enforce cap on credit
      $this->_credit =
        min($this->_credit,
            $this->gameParameters->getMaximumCredit());
    }

    function _maybeChangeLevel () {
      if ($this->_credit < 0.0) {
        $this->_gameLost = TRUE;
      }
      else {
        $params = $this->gameParameters;
        $current_level = $this->_level;
        if ($current_level >
            $params->getMaximumLevel()) {
          $this->_gameWon = TRUE;
        }
        else {
          // find out if questions remain to be
          //  asked at this level
          if (($this->_questionsAskedAtLevel >=
               $params->getQuestionsPerLevel($current_level)) ||
              (count($this->_questionIdsAtLevel) == 0)) {
              // either we have asked the limit of
              // questions per level, OR we have simply run out
              $this->_level++;
              $this->_questionsAskedAtLevel = 0;
              $this->_setupQuestionIds();
              // note recursive call --- it's possible
              //  that no questions were found, and we have
              //  to keep going
              $this->_maybeChangeLevel();
          }
        }
      }
    }

    function __sleep () {
      // make sure to serialize all fields except
      //  the database connection (has to be recreated)
      //  and the previous question (no point).
      return(array(
             'gameParameters',
             'currentQuestion',
             '_credit',
             '_level',
             '_questionIdsAtLevel',
             '_questionsAskedAtLevel',
             '_correctAnswers',
             '_totalQuestions',
```

752

```
                    '_gameLost',
                    '_gameWon'));
        }
    }
?>
```

Handling an answer

Here are the steps that a Game object goes through in dealing with a guess range submitted by a player (in the function updateWithAnswer):

1. Move the current question object to the previous question slot.

2. Update the (now previous) question with the upper and lower ranges of the guess (which are still in terms of step numbers from the form rather than actual values).

3. Query the previous question to discover the range of the guess and whether the question was correctly answered. Update all scores (credit, correct answers) appropriately.

4. Decide whether to promote the player to a new level now. If so, retrieve the database IDs of all questions that may be asked at that level. Randomize the order of the question list.

5. If the game has not yet ended, grab a new question ID from the randomized list and use it to ask the database for a new question. Turn that data into a Question object and make it the current question.

Serialization and sleep()

The sleep() function is called to do cleanup whenever an object is serialized and also returns a list of all the member variables that should be recorded in a serialization. The Game class makes use of only the latter capability — all member variables except the previous question and the database connection itself are retained as the object is stored in the session for the next page's use.

game_parameters_class.php

The single instance of the GameParameters class, shown in Listing 44-5, packages up all the default numbers that we may want to customize in making a new version of the game (the penalties and rewards, the number of levels, the starting and maximum credit, and so on). In addition, this object manages global access to the database connection.

LISTING 44-5

game_parameters_class.php

```php
<?php
include_once("certainty_utils.php");
include_once("dbvars.php");
```

```
class GameParameters {

  var $_dbConnection = NULL;
  var $_startingLevel = 1;
  var $_maximumLevel = 10;
  var $_startingCredit = 5.0;
  var $_maximumCredit = 15.0;
  var $_questionsPerLevel = 3;
  var $_rightAnswerCredit = 1.0;
  var $_wrongAnswerDebit = 4.0;
  var $_answerSpreadDebit = 4.0;

  // CONSTRUCTOR
  function GameParameters () {
    // all fields set by default values
  }

  // PUBLIC FUNCTIONS
  // accessors

  function getStartingLevel () {
    return($this->_startingLevel);
  }

  function getMaximumLevel () {
    return($this->_maximumLevel);
  }

  function getStartingCredit () {
    return($this->_startingCredit);
  }

  function getMaximumCredit () {
    return($this->_maximumCredit);
  }

  function getRightAnswerCredit () {
    return($this->_rightAnswerCredit);
  }

  function getWrongAnswerDebit () {
    return($this->_wrongAnswerDebit);
  }

  function getAnswerSpreadDebit () {
    return($this->_answerSpreadDebit);
  }
```

```
  function getQuestionsPerLevel () {
    return($this->_questionsPerLevel);
  }

  function getDbConnection () {
    global $host, $user, $pass, $db;  // from dbvars.inc
    if ($this->_dbConnection &&
        is_resource($this->dbConnection)) {
      return($_dbConnection);
    }
    else {
      // suppress warnings about connection,
      // will handle at higher level if failed
      $connection =
        @mysql_connect($host, $user, $pass);
      if ($connection &&
          mysql_select_db($db, $connection)) {
        return($connection);
      }
      else {
        return(FALSE);
      }
    }
  }

  function __sleep () {
    return(array('_startingLevel',
                 '_startingCredit',
                 '_rightAnswerCredit',
                 '_wrongAnswerDebit',
                 '_answerSpreadDebit',
                 '_questionsPerLevel'));
  }
}
?>
```

certainty_utils.php

This code file, shown in Listing 44-6, is a grab-bag for capabilities and definitions that do not fit neatly into a particular class and that are used in more than one other code file.

Everything in certainty_utils.php fits into one of a few categories:

- Initial declarations (seeding the random number generator, setting the error-reporting level).
- Abstraction functions for session and post variables.
- Utility functions for calculating intermediate answer values and for randomizing question lists.

755

LISTING 44-6

certainty_utils.php

```php
<?php

// Definitions and utility functions for the
//  Certainty Quiz game
error_reporting(E_ALL);
// enumeration constants for the scaling of distractors
define("CERTAINTY_LINEAR", 1);
define("CERTAINTY_GEOMETRIC", 2);

// Seed the random number generator
srand((double) microtime() * 1000000);

// A hack to retrieve the value of POST values
//   without explicitly checking PHP versions.
function get_post_value ($var_name) {
  global $HTTP_POST_VARS;
  if (IsSet($_POST) &&
      IsSet($_POST[$var_name])) {
    return($_POST[$var_name]);
  }
  elseif (IsSet($HTTP_POST_VARS) &&
          IsSet($HTTP_POST_VARS[$var_name])) {
    return($HTTP_POST_VARS[$var_name]);
  }
  else {
    return(FALSE);
  }
}

function get_session_value ($var_name) {
  global $HTTP_SESSION_VARS;
  if (IsSet($_SESSION) &&
      IsSet($_SESSION[$var_name])) {
    return($_SESSION[$var_name]);
  }
  elseif (IsSet($HTTP_SESSION_VARS) &&
          IsSet($HTTP_SESSION_VARS[$var_name])) {
    return($HTTP_SESSION_VARS[$var_name]);
  }
  else {
    return(FALSE);
  }
}

function set_session_value ($var_name, $value) {
```

```php
  global $HTTP_SESSION_VARS;
  if (IsSet($_SESSION)) {
    $_SESSION[$var_name] = $value;
    $HTTP_SESSION_VARS[$var_name] = $value;
  }
  else {
    $HTTP_SESSION_VARS[$var_name] = $value;
  }
}

function unset_session_value ($var_name) {
  if (IsSet($_SESSION[$var_name])) {
    $_SESSION[$var_name] = '';
    unset($_SESSION[$var_name]);
  }
}

// Numerical functions

function round_to_digits ($number, $digits) {
  if ($number < 0) {
    return(- round_to_digits(- $number, $digits));
  }
  else if ($number == 0) {
    return($number);
  }
  else {
    $tens =
      floor(log10($number));
    $divisor = pow(10, ($tens - $digits));
    $significant = (1.0 * $number) /
      $divisor;
    $rounded = round($significant);
    return($rounded * $divisor);
  }
}

function nth_root_initial($product, $n)
{
  $estimate = sqrt($product);
  $roots = 2;
  while ($roots < $n) {
    $estimate = sqrt($estimate);
    $roots = $roots * 2;
  }
  return($estimate);
}

function nth_root ($product, $n) {
```

```
    if (($product <= 1) ||
        ($n < 2)) {
      die("Arguments to nth_root should be ".
          "product (greater than 1) and " .
          "n (greater than 1)");
    }
    $initial_estimate =
      nth_root_initial($product, $n);
    return(nth_root_aux($product, $n,
             $initial_estimate,
             20000,
             0.0001));
  }

  function nth_root_aux ($product, $n,
                         $guess,
                         $iterations_left,
                         $desired_difference) {
    if ($iterations_left <= 0) {
      return($guess);
    }
    else {
      $guessed_product = pow($guess, $n);
      if (abs($guessed_product - $product)
          < $desired_difference) {
        return($guess);
      }
      else {
        $new_guess =
          $guess -
            ((pow($guess, $n) - $product) /
              ($n * pow($guess, $n-1)));
        return(nth_root_aux($product, $n,
                            $new_guess,
                            $iterations_left - 1,
                            $desired_difference));
      }
    }
  }

  function create_randomized_array ($in_array) {
    // Assumes input is simple list, with keys
    //  equal to 0,...,n
    // Returns similar list, with keys as in input
    //  but values in randomized order
    // Assumes prior call to srand()
    $in_array_length = count($in_array);
    $working_array = array();
    for ($i = 0; $i < $in_array_length; $i++) {
      $rand_value = rand();
```

```
        $working_array[$i] = $rand_value;
    }
    asort($working_array);  // orders by random value
    $return_array = array();
    $working_keys = array_keys($working_array);
    foreach ($working_keys as $int_key) {
        array_push($return_array,
            $in_array[$int_key]);
    }
    return($return_array);
}
?>
```

The functions in `certainty_utils.php` take care of figuring out all the intermediate guesses between the lowest value offered to the user and the highest value. In addition, there's a scaling option, which determines whether the intermediate values grow linearly or geometrically. (If you think that the number "between" 10 and 1000 is 100, you are scaling geometrically; if you think the number between 10 and 1000 is 505, you are scaling linearly.) The functions for finding nth roots are used in doing the geometric scaling.

The `create_randomized_array()` function is what we use to scramble the order of questions within a level.

question_class.php

Finally, we get down to the actual questions that are pulled from our database of questions to ask. The definition of the `Question` class is shown in Listing 44-7. The public functions here are:

- The constructor, which is given the question, correct answer, the upper and lower bounds, the number of steps in the guesses, and the type of scaling (linear or geometric).
- Various accessor functions, such as `getAnswer()`, `getQuestion()`, `getScalingType()`.
- `updateWithAnswer()`, which bottoms out here by actually translating the Web form's step numbers to values for the guesses, comparing those guesses to the real answer.
- `getAnswerSpread()`, which returns a measure of how narrow the guess was.

LISTING 44-7

question_class.php

```php
<?php
include_once("certainty_utils.php");

class Question
{
  // PRIVATE VARIABLES
  private $_id; // ID in database
```

759

```php
private $_question; // text of question
private $_answer; // correct numeric answer
private $_lowerLimit; // smallest value in distractors
private $_upperLimit; // largest value in distractors
private $_distractorCount; // number of dist. presented
private $_scalingType;  // representing linear vs. geometric
private $_distractorArray; // contains all dist presented
private $_lowerGuess = NULL; // player's lower bound
private $_upperGuess = NULL; // player's upper bound
private $_correct = NULL;  // TRUE or FALSE after guess

// CONSTRUCTOR
function __construct($id, $question,
                      $answer,
                      $lower_limit,
                      $upper_limit,
                      $distractor_count,
                      $scaling_type) {
  $this->_id = $id;
  $this->_question = $question;
  $this->_answer = $answer;
  $this->_lowerLimit = $lower_limit;
  $this->_upperLimit = $upper_limit;
  $this->_distractorCount = $distractor_count;
  $this->_scalingType = $scaling_type;
  $this->_distractorArray =
      $this->_makeDistractors($lower_limit,
                  $upper_limit,
                  $distractor_count,
                  $scaling_type);
}

// PUBLIC FUNCTIONS

// accessors

function getId () {return($this->_id);}
function getQuestion () {return($this->_question);}
function getAnswer () {return($this->_answer);}
function getCorrect() {return($this->_correct);}
function rightAnswer() {return($this->_correct);}
function getDistractorCount() {return($this->_correct);}
function getScalingType() {return($this->_scalingType);}
function getDistractorArray()
  {return($this->_distractorArray);}
function getLowerGuess() {return($this->_lowerGuess);}
function getUpperGuess() {return($this->_upperGuess);}

function getAnswerSpread () {
  $answer_range = count($this->_distractorArray) - 1;
```

```
    if (IsSet($this->_lowerGuess) &&
        IsSet($this->_upperGuess)) {
      $lower = $this->_lowerGuess;
      $upper = $this->_upperGuess;
      if ($upper < $lower) {
        throw new Exception("Problem in range of answers");
      }
      else {
        $spread =
                (max($upper - $lower, 1) - 1)
                / ($answer_range - 1);
        return($spread);
      }
    }
    else {
      throw new Exception("Answer variables not set");
    }
}

function updateWithAnswer($lower, $upper) {
  // takes a lower and upper guess from player, and
  //  determines if the guesses bound the right answer
  $this->_lowerGuess = $lower;
  $this->_upperGuess = $upper;
  $upper_value = NULL;
  $lower_value = NULL;
  $count = 1;
  foreach ($this->_distractorArray as $distractor) {
    if ($count == $lower) {
      $lower_value = $distractor;
    }
    if ($count == $upper) {
      $upper_value = $distractor;
    }
    $count++;
  }
  if (IsSet($lower_value) && IsSet($upper_value)) {
    $answer = $this->_answer;
    $lower_value_lowered = $lower_value -
      max(0.0001, abs($lower_value / 1000000.0));
    $upper_value_raised = $upper_value +
      max(0.0001, abs($upper_value / 1000000.0));
    if (($lower_value_lowered <= $this->_answer) &&
        ($upper_value_raised >= $this->_answer)) {
      $this->_correct = TRUE;
    }
    else {
      $this->_correct = FALSE;
    }
  }
```

```
    else {
      $this->_correct = NULL;
    }
}

// PRIVATE FUNCTIONS
private function _makeDistractors ($lower, $upper,
                                  $distractor_count,
                                  $linear_or_geometric)
// Create the array of intermediate values between
//   the upper bound and the lower bound on guesses
//   that the player can choose from.  Depending on
//   a flag in each row of the question database,
//   the scaling of possible answers ("distractors")
//   can be linear (10, 20, 30 ...) or geometric
//   (10, 20, 40, 80 ...)
// Code for construction of geometric distractors can
//    blow up for some arguments, so arguments are
//    checked before calls to make_distractors_geometric
//    are allowed.  Failures default back to linear.
{
  if (($linear_or_geometric == CERTAINTY_GEOMETRIC) &&
      ($this->safeGeometricArguments($upper, $lower))) {
      return($this->_makeDistractorsGeometric(
        $lower, $upper, $distractor_count));
  }
  else {
    return($this->_makeDistractorsLinear(
      $lower, $upper, $distractor_count));
  }
}

private function safeGeometricArguments ($upper, $lower) {
 // should probably really also include the number
 // of distractors as an argument.  Only tested for
 // # of distractors approx 10.
 return (($upper > 0) && ($lower > 0) &&
         ($upper > $lower) &&
         (($upper / $lower) < 10000000000));
}

private function _makeDistractorsLinear
        ($lower, $upper, $distractor_count)
{

  $return_array = array();
  array_push($return_array, round_to_digits($lower, 3));
  $current = $lower;
  $increment = (($upper - $lower) / $distractor_count);
  // add in all the intermediate values
```

```
    for ($x = 1; $x < $distractor_count; $x++) {
      array_push($return_array,
                  round_to_digits($lower +
                          ($x * $increment),
                               3));
    }
    array_push($return_array, round_to_digits($upper, 3));
    return($return_array);
  }

  private function _makeDistractorsGeometric
          ($lower, $upper, $distractor_count)
  {
    if (($lower >= $upper) ||
        ($distractor_count < 2)) {
      die("Args to _makeDistractorsGeometric should be " .
          "1) a lower limit, 2) an upper limit, " .
          "3) a count (>= 2) of divisions between them.<BR>" .
  "Args were 1) $lower, 2) $upper, 3) $distractor_count<BR>");
    }
    $return_array = array();
    array_push($return_array, round_to_digits($lower, 3));
    $limit_ratio = $upper / $lower;
    $root = nth_root($limit_ratio, $distractor_count);
    $current = $lower;
    // add in the intermediate values
    for ($x = 1; $x < $distractor_count; $x++) {
      $distractor = round_to_digits(
                      $lower * pow($root, $x),
                      3);
      array_push($return_array,
                  $distractor);
    }
    array_push($return_array, round_to_digits($upper, 3));
    return($return_array);
  }
}
?>
```

dbvars.php

When we actually query the database, we need to have access information. The file shown in
Listing 44-8 is loaded by GameParameters.php and sets up the variables necessary for making a
MySQL connection. Note that the current values are dummies and will *not work* on your system! You
need to fill in the correct values for your own MySQL configuration. If your web server is connected
to the Internet, it's also a good idea to move this file somewhere outside the web server document
tree and to change the reference in GameParameters.php to point to its new location.

LISTING 44-8

dbvars.php

```php
<?php
$host = "YOUR_HOSTNAME";
$user = "YOUR_MYSQL USERNAME";
$pass = "YOUR_MYSQL_PASSWORD";
$db = "certainty";
?>
```

Creating the database

The trivia game is fueled by a database of questions. So far, we have said nothing about how to create such a database.

Table definitions

Listing 44-9 shows a MySQL dump file of all the table definitions used in the code, along with a few sample entries. Before loading it, you need to create a database called `certainty`; after that is done, you should be able to simply `cat` or pipe the contents of this file to the `mysql` command.

Note that the question table includes several fields that are not actually used in the current code. One of them is `attribution`, useful for recording the book or web site that served as the authority for the answer. Another is `include`, which we intended for filtering out questions in development that were not ready to be displayed. A third is `subjectID`, which we use to tag questions according to subject area (Science, Geography, History, and so on). although that association is not actually displayed anywhere. A final as-yet unused column is `unitID`, which could be used to record the unit (kilometers, years, furlongs, bushels) of the answer in case the unit affects how guesses should be displayed.

LISTING 44-9

Table definitions

```
# MySQL dump 7.1
#
# Host: [host deleted]    Database: certainty
#--------------------------------------------------------
# Server version   3.22.32

#
# Table structure for table 'high_scores'
#
CREATE TABLE high_scores (
  id int(11) NOT NULL auto_increment,
```

```
  name varchar(30),
  answer_count int(11),
  credit double(16,4),
  PRIMARY KEY (id)
);

#
# Dumping data for table 'high_scores'
#

INSERT INTO high_scores VALUES (8,'Ben Stein',15,-3.0000);

#
# Table structure for table 'question'
#
CREATE TABLE question (
  ID int(11) NOT NULL auto_increment,
  answer double(16,4),
  unitID int(11),
  level int(11),
  subjectID int(11),
  include tinyint(4),
  upper_limit double(16,4),
  lower_limit double(16,4),
  scaling_type tinyint(4),
  question varchar(255),
  attribution varchar(255),
  PRIMARY KEY (ID)
);

#
# Dumping data for table 'question'
#

INSERT INTO question VALUES
(1,5283755345.0000,1,1,1,1,200000000000.0000,1000000.0000,2,
'What was the human population of the world
in the middle of 1990?',
'http://www.census.gov/ipc/www/worldpop.html');
INSERT INTO question VALUES (2,70.0000,1,1,1,1,95.0000,5.0000,1,
'What percentage of the Earth\'s surface is covered by water?',
'http:/www.sciencenet.org.uk/database/Geography/
Original/g00057d.html');
INSERT INTO question VALUES
(4,1969.0000,NULL,1,2,NULL,2000.0000,1950.0000,1,
'In what year did human beings first walk on the moon?','');
#
# Table structure for table 'subject'
#
CREATE TABLE subject (
```

```
    id int(11) NOT NULL auto_increment,
    subject varchar(255),
    PRIMARY KEY (id)
);

#
# Dumping data for table 'subject'
#

INSERT INTO subject VALUES (1,'Geography');
INSERT INTO subject VALUES (2,'History');
INSERT INTO subject VALUES (3,'Science');
INSERT INTO subject VALUES (4,'Mathematics');
INSERT INTO subject VALUES (5,'Miscellaneous');
```

The MySQL dump includes only three sample questions. You can always add more through a direct interaction with MySQL, but it's more convenient to do it via a Web form.

entry_form.php

Listing 44-10 shows a bare-bones Web form for entering more questions into the question database. This simply takes typed input (except for a pull-down association with the subject table) and trusts the results. Note that there is neither security nor error-checking here — this is intended for use only by the game creator, and if misuse is a concern you should probably add password protection or some other kind of authentication. (See Chapter 44 for more on creating authentication systems.)

LISTING 44-10

entry_form.php

```php
<?php
include_once("certainty_utils.php");
include_once("game_parameters_class.php");
$params = new GameParameters();
$connection = $params->getDbConnection();

if (get_post_value('POSTCHECK')) {
  handleEntryForm();
}
displayEntryForm();

function handleEntryForm () {
  $question = mysql_real_escape_string(get_post_value('QUESTION'));
  $answer = mysql_real_escape_string(get_post_value('ANSWER'));
  $lower_limit = mysql_real_escape_string(get_post_value('LOWER_LIMIT'));
  $upper_limit = mysql_real_escape_string(get_post_value('UPPER_LIMIT'));
  $level = mysql_real_escape_string(get_post_value('LEVEL'));
```

```
    $subject = mysql_real_escape_string(get_post_value('SUBJECT'));
    $scaling_type = mysql_real_escape_string(get_post_value('SCALING_TYPE'));
    $attribution = mysql_real_escape_string(get_post_value('ATTRIBUTION'));
    if ($upper_limit > $lower_limit) {
      $query =
        "insert into question
         (question, answer, lower_limit, upper_limit,
          level, subjectID, scaling_type,
          attribution)
         values
         ('$question', $answer, $lower_limit, $upper_limit,
          $level, $subject, $scaling_type,
          '$attribution')";
      $result = mysql_query($query);
      if ($result) {
        print("Entry was successful<BR>");
      }
      else {
        print("Entry was not successful<BR>");
      }
    }
    else {
      print("Upper limit must be greater than lower<BR>");
    }
}

function displayEntryForm () {
global $PHP_SELF;
$linear = CERTAINTY_LINEAR;
$geometric = CERTAINTY_GEOMETRIC;
$subject_string = make_subject_string();
$form_string = <<<EOT
<FORM METHOD=POST TARGET="entry_form.php" >
Question:
<INPUT TYPE=TEXT NAME=QUESTION SIZE=60 ><BR>
Answer:
<INPUT TYPE=TEXT NAME=ANSWER><BR>
Lower:
<INPUT TYPE=TEXT NAME=LOWER_LIMIT><BR>
Upper:
<INPUT TYPE=TEXT NAME=UPPER_LIMIT><BR>
Level:
<INPUT TYPE=TEXT NAME=LEVEL><BR>
Subject:
$subject_string<BR>
Scaling type:
<SELECT NAME=SCALING_TYPE>
  <OPTION VALUE=$linear>Linear
  <OPTION VALUE=$geometric>Geometric
</SELECT><BR>
```

```
Attribution:
<INPUT TYPE=TEXT NAME=ATTRIBUTION><BR>
<INPUT TYPE=SUBMIT NAME=SUBMIT VALUE=SUBMIT>
<INPUT TYPE=HIDDEN NAME=POSTCHECK VALUE=1>
</FORM>
EOT;
echo $form_string;
}

function make_subject_string () {
  $result_string = "<SELECT NAME=SUBJECT>";
  $query = "select id, subject from subject order by id";
  $result = mysql_query($query);
  while ($row = mysql_fetch_row($result)) {
    $id = $row[0];
    $display = $row[1];
    $result_string .= "<OPTION VALUE=$id>$display";
  }
  $result_string .= "</SELECT>";
  return($result_string);
}
?>
```

General Design Considerations

What follows is a brief list of issues that we were forced to consider while writing the code in this chapter.

Separation of code and display

The question of separating code and display is a vexing one, especially in situations where you have different personnel assigned to maintaining logic and appearance. Our own view on this is that perfect separation of code and display is like a perfect vacuum — you can get asymptotically closer to the ideal as you expend infinite effort.

For large web sites employing many people, some pretty good techniques exist for making a strong separation, including templating systems and database storage of graphics and display text. For this relatively small and informal example, we were satisfied by simply segregating all HTML into two display-oriented classes, leaving the remainder of the code focused on logic and data.

Persistence of data

There are several kinds of data in this game that survive longer than the execution time of a page. We chose to use PHP's session mechanism for all the data particular to a particular game invocation and a backend database for everything else (questions, answers, and high score lists).

For reasons of efficiency, we didn't want to store too much data via the session mechanism. So we separated out the most important data (that could not be easily recreated) into the `Game` class and stored only an instance of that class. Everything else (question text, boilerplate HTML text, high scores, and so on) was either embedded in code files or easily retrievable from the database.

Exception handling

We used the new (as of PHP5) exception mechanism to bail out whenever we encountered a problem that the code could not recover from. Failures to recover session info, failures to find cookies, and database interaction problems were all grounds for giving up. In general, when we threw an exception, we sent a string suitable for display to a user and then caught all thrown exceptions at the point of display. This has the disadvantage of not providing a lot of rich debugging information (particularly since several different code paths can throw a "No database connection" exception) but has the advantage that we can tell the user something reasonable and fairly cosmetic, while giving the developer a hint.

Summary

The Certainty Quiz is a small, simple, self-contained PHP application that you should be able to install and enjoy in the privacy of your own home (after connecting it to your favorite web server and a MySQL database, of course). Although small, the code relies on database interaction, OOP features, use of sessions, string processing, object serialization, exception handling, and nontrivial arithmetic to achieve its effects. Although PHP has many capabilities that we didn't come close to touching on, this chapter uses a fair cross-section of its most popular features — if you understand everything in this example, you are well on your way to exploiting the power of PHP.

Chapter 45

Data Visualization with Venn Diagrams

I n this chapter's case study, we show one way to use PHP to combine MySQL databases with graphic images. We build a complete system that starts with a database and uses the gd library to produce a kind of visualization of the data. The portions of the book we draw on for this are:

- **Part II**: We use PHP to interrogate a MySQL database.
- **Chapter 42 (Graphics)**: Our end product is an image produced with the gd library.
- **Chapter 27 (Mathematics)**: We need a bit of trigonometry as we create the images.

Scaled Venn diagrams

The visualization we have in mind is something like the Venn diagram. If you've ever been in an academic setting where set intersection was being discussed, then you've probably seen these diagrams — they're the circles that may or may not have overlapping portions representing intersections.

We say "something like" the Venn diagram, because scale has no significance in a traditional Venn diagram. If you want to illustrate the fact that there are people who use both BeOS and Windows, then you might draw two circles of equal size (representing Windows users and BeOS users) that happen to have a region of overlap. In our version, which you might call a *scaled* or *proportional* Venn diagram, the sizes of both circles and intersections matter; the Windows/BeOS example would become one large circle and one much smaller circle, with an overlap area proportional to the number of people in both sets. (To see an example of this kind of diagram, please skip ahead to Figure 45-5.)

The task

The job of our code is to start with a database, provide a way to query that database about sets and their overlap, and then display the results as a scaled Venn diagram, generated by using the gd library. As a sample database, we use the pseudosurvey data set that we used in the "HTML Graphics" section of Chapter 42.

If we're going to offer a way to query the database, then it may as well be via a Web form. So the end-to-end view of our task is that we start with a Web form and end up with a picture to display. Let's start the design by enumerating the things that need to happen for this to come about. We'll need to:

1. Generate (or at least present) the Web form itself.

2. Receive the submitted form data and transform it into appropriate SQL queries for submission to the database.

3. Receive results from the SQL queries.

4. Use the SQL results to decide on the locations and sizes of all the elements in our graphic.

5. Actually generate the graphic and send it back to the user.

> **NOTE**   **All of the code in this chapter should work with either PHP5 or PHP6, but it assumes that your PHP installation has access to the** gd **image library and is configured to produce PNG images. Any version of** gd **later than 1.8, bundled or unbundled, should be OK. (See Chapter 42 for details of configuration and installation of** gd**.)**

Outline of the code

Our system contains the following code files:

- `visualization_form.php`: This is essentially a hardcoded form that enables the user to choose two different restrictions on the data in our table. The restrictions chosen map directly to `where` clauses loaded from an auxiliary file called `query_clauses.php`.

- `db_visualization.php`: This code handles the form data sent by `visualization_form.php` and builds three SQL statements: one with only the first `where` clause, one with the second `where` clause, and a third with both clauses joined by an `and`. It collects the resulting three counts and displays the numbers in a graphic by calling functions loaded from `venn.php`.

- `venn.php`: This actually produces the Venn diagram graphic and ships it back to the user. Its primary function takes as input the three amounts (the sizes of the two sets and their intersection), decides the sizes and locations of corresponding circles, and does all the drawing and shading necessary. For the complicated case of sets that actually have an overlapping area, it uses functions loaded from `trig.php` to calculate areas.

- `trig.php`: This code actually calculates the intersection area whenever circles overlap.

We discuss these code files in reverse order, from the bottom up. By the way, although we like this example, we don't want to give the impression that you need to do trigonometry to do computer graphics in PHP, or even vector graphics in PHP. If you want to understand every bit of this example, then you need to go through the trig, but we encourage those who don't care to skip the next section ("Necessary trigonometry"). The core of the graphics code itself is in `venn.php`, and that example code really is important to understand if you want to do `gd`-based graphics in PHP.

Necessary Trigonometry

Let's get the math out of the way first. Unavoidably, because we're talking about circles and areas, we're going to be talking about trigonometry. (As we've said, though, if you're not interested and are willing to trust us that we have code to calculate the area of circle intersections, please do skip ahead to the section "Planning the display," later in this chapter.)

The eventual task for our system is to start with three quantities (items in set A, items in set B, and items in the intersection) and produce a diagram containing two circles, with areas proportional to the set sizes, and positioned so that the area of overlap is proportional to the size of the intersection. For this section, we go in the other direction and calculate intersection area from given circles. Our starting information will be the radii of the two circles and the distance between their centers.

With reference to Figure 45-1, say that our circles have centers at points C and D, respectively, and that we know the radius of the circle on the left (segment CA or segment CB) and the radius of the circle on the right (DA or DB). What we'd like to know is the size of that odd lens-shaped object in the middle.

FIGURE 45-1

Area of intersection

Area of intersecting circles

The lens-shaped intersection area is split into two "halves" by segment AB (not quite halves because the circle sizes may be different), and we can calculate the area of each half independently. The crucial thing to notice is that the area of each of these half-lenses is the area that you get after you subtract the area of a triangle from the area of a pizza-slice-shaped sector of a circle. The right-hand lens half, for example, has an area equal to the sector of the left-hand circle determined by angle ACB, minus the area of the triangle ACB.

So if we can calculate the areas of sectors and triangles, we are nearly done. The area of a sector is straightforward — it's just the area of the circle multiplied by the fraction of that circle that the angle of the sector sweeps over.

It takes a little more work and trigonometry to get the areas of the triangles. In our code, we make the job more straightforward by drawing a line from point C to point D and considering only the half of the diagram above that line — then at the end, we multiply by two to get the real area. If we say that the intersection of segments AB and CD is point E, what we eventually care about is the area of triangles CAE and DAE. We start by calculating the angles of triangle CDA (whose side lengths are known to us) and, by using that information, determining the lengths of CE, DE, and AE. After we know these lengths, we know the bases and heights, and the areas of the right triangles CAE and DAE are just ½ (*base* × *height*).

Listing 45-1 shows code to do this kind of area calculation. Its main "public" function is `circle_intersection_area()`, which expects as arguments the radii of two circles and the distance between them. The simplest case is where the distance is greater than the sum of the radii: The circles do not touch; there is no intersection, and the answer is zero.

LISTING 45-1

trig.php

```php
<?php

function angle_given_sides ($opposite, $other_1, $other_2) {
  if (($opposite <= 0) ||
      ($other_1 <= 0) ||
      ($other_2 <= 0) ||
      ($opposite >= ($other_1 + $other_2)) ||
      ($other_1 >= ($opposite + $other_2)) ||
      ($other_2 >= ($other_1 + $opposite))) {
    die("Triangle with impossible side lengths in ".
        "angle_given_sides: $opposite, $other_1, $other_2");
  }
  else {
    $numerator =
      ((($other_1 * $other_1) +
        ($other_2 * $other_2)) -
        ($opposite * $opposite));
    $denominator = 2 * $other_1 * $other_2;
    return(acos($numerator / $denominator));
```

```
  }
}

function area_to_radius ($area) {
  return (sqrt ($area / M_PI));
}

function circle_intersection_area ($radius_left,
                                   $radius_right,
                                   $distance) {
  if ($radius_right + $radius_left <= $distance) {
    return(0);
  }
  else {
    // first, we find the angle measures of a triangle
    //    formed by the two radii and the distance
    //    between them
    $left_sector_angle =
      angle_given_sides($radius_right, $radius_left,
                        $distance);
    $right_sector_angle =
      angle_given_sides($radius_left, $radius_right,
                        $distance);

    // test for obtuseness --- the sector angle can
    // be obtuse, but the triangle angle should not
    // be.  Also save the result as a sign for the
    // eventual area calculation

    if ($left_sector_angle < M_PI / 2) {
      $left_triangle_angle = $left_sector_angle;
      $left_triangle_sign = 1;
    }
    else {
      $left_triangle_angle = M_PI - $left_sector_angle;
      $left_triangle_sign = -1;
    }
    if ($right_sector_angle < M_PI / 2) {
      $right_triangle_angle = $right_sector_angle;
      $right_triangle_sign = 1;
    }
    else {
      $right_triangle_angle = M_PI - $right_sector_angle;
      $right_triangle_sign = -1;
    }

    // next, find the height of that triangle, assuming
    //   the distance is the base

    $height = ($radius_left / sin(M_PI_2)) *
```

```
                sin($left_triangle_angle);
    $base_left = ($radius_left / sin(M_PI_2)) *
                  sin(M_PI_2 - $left_triangle_angle);
    $base_right = ($radius_right / sin(M_PI_2)) *
                  sin(M_PI_2 - $right_triangle_angle);

    // finally find triangle and sector areas, and
    // subtract (or add) appropriately to get the
    // intersection area.  Multiply by 2 to reflect
    // areas on both sides of the segment connecting
    // the circle centers

    $left_triangle_area = $base_left * $height / 2;
    $right_triangle_area = $base_right * $height / 2;
    $left_sector_area =
       ($left_sector_angle / (2 * M_PI)) *
       (M_PI * $radius_left * $radius_left);
    $right_sector_area =
       ($right_sector_angle / (2 * M_PI)) *
       (M_PI * $radius_right * $radius_right);

    $intersection_area = 2 *
       (($left_sector_area -
           ($left_triangle_sign * $left_triangle_area)) +
        ($right_sector_area -
           ($right_triangle_sign * $right_triangle_area)));

    return($intersection_area);
  }
}

?>
```

Note that all the angle calculations are in radians, rather than degrees. In radians, a right angle is pi/2, and a complete revolution around a circle is 2 pi. We tend to use PHP constants for these values whenever we can, in particular M_PI (the value of pi), and M_PI_2 (pi/2).

There's one final wrinkle that we've ignored in our discussion so far but that we had to deal with in the code. The problem is that it's possible for either angle ACE or angle ADE (as we call them in Figure 45-1) to be obtuse — that is, more than 90 degrees in size. To see this, look at that diagram and imagine what happens as you make the circle on the right smaller, and move its center D progressively closer to C. At some point D actually moves to the left of segment AB. In this case, the circle intersection area to the left of AB is actually the sum of a sector and a triangle rather than a difference. The sector determined by DA and DB sweeps out more than half of the circle centered at D, and the remaining portion we want to include is the area of the triangle ADB. We handle this in the code by testing if the angles are obtuse and multiplying the triangle areas by either 1 or -1, depending on the result of the test.

Planning the Display

Now we pop up a couple of levels and think about actually generating a diagram. We assume that we have as input three numbers (size of set 1, size of set 2, and size of intersection), along with some textual labels. We want to scale and locate these circles so that everything has the right area, labels are associated with the right circles, and everything fits within the size of the diagram we're creating.

Simplifying assumptions

We start off with some totally arbitrary decisions that, after being made, simplify everything. We decree that:

- All the images that we generate are the same size, and that size is 300 pixels high and 600 pixels wide.

- The centers of the circles are always on the same horizontal line. This means that their y-coordinate is decided in advance, and we change the area of intersection just by changing the x-coordinates.

- The circles always fit within the top two-thirds of the diagram (reserving the lower third for labels). So put the y-coordinate of the centers one-third of the way down the image from the top. And because the circles may not intersect at all, they shouldn't be larger than half of the width of the image, so we have room to display two of them. We also make sure that the circles are no greater than 90 percent of the room available given everything we've said so far, so that they don't touch the image borders. Finally, we decide that, regardless of the actual numbers as input, the larger of the two circles is as large as it can be. (Scale is consistent within the diagram, but not between diagrams.)

Determining size and scale

Now we have nearly all the information needed to create a visualization, and the pieces we are lacking, of course, depend on the input values we are going to receive. We use the sizes of the actual sets to determine the radii of the circles for display. We want the larger of the two set counts to correspond to the largest circle we can afford to display, and then scale everything else appropriately. (We do all this in the code in Listing 45-2 [`venn.php`] — you may want to look ahead to that code as we lay out what we need to do in it.)

It's actually easiest for us to calculate the largest radius we can afford: Given the constraints we've already listed, the larger radius should be 90 percent of ¼ of the image width, or 90 percent of ⅓ of the image height, whichever is smaller. So we calculate this maximum radius, assume that the larger set size is proportional to the area of a circle with this radius, and come up with a general conversion for mapping from input numbers to area as measured in pixels. We use this to decide on the areas of the circles and of the intersection area we want.

What numbers should we know at this point? We know:

- The radius of the bigger circle. (It's the largest radius that fits our constraints.)
- The area of the bigger circle (calculated as πr^2).
- The radius of the smaller circle (from the ratio of the input set sizes treated as area ratios and then mapped back to a radius).
- The area of the smaller circle (calculated).
- The area of intersection (scaled the same way as other areas, from the input numbers).
- The y-coordinate of the circle centers. (We decreed that it be the line that's one-third of the way down the image.)

What are we missing before we can display our circles? The only thing that we're missing is the x-coordinates of the centers.

The easy cases

Where we decide to put the circle centers depends on the extent to which our sets overlap. There are some cases that we can dispense with, that don't need all this trigonometry we've been spending our time on. Those are:

- **No items are in the intersection**: In this case, we don't want the circles to touch at all. We simply locate the centers at default locations in the middles of the two halves of the diagram. Because of the way that we limited the maximum radius, the circles are completely separated.
- **One set is completely contained in the other — that is, one of the sets has the same size as the intersection**: For this case, we just choose to put the center of the larger circle in the middle of the diagram and the middle of the other circle offset a bit from it but not so much that any of the smaller circle is outside the larger one.
- **The two sets are the same (and all three input numbers are the same)**: For this, we just draw one circle with an x-coordinate right in the middle of the picture.

The hard case

Now the hard one: If the sets only partially overlap, where should we put the circle centers? At this point, we have some math in our pocket from the "Necessary trigonometry" section: Given two circles and the distance between their centers, we can figure out the area of overlap. Unfortunately, this is not the direction we need the calculation go in — we start with the desired area of intersection, and we must work backwards to the desired locations of the circle centers.

Now if we were good and diligent mathematicians but lazy programmers, we would just invert the trigonometric equations we used in `trig.php`, to solve for center distance rather than for intersection area. As it is, though, we're enthusiastic programmers, and if we're any kind of mathematicians at all we're definitely the lazy kind. So what we're going to do instead is search for the answer. The

function `find_circle_centers()` in Listing 45-2 implements a binary search for the answer: It starts with a middling distance, asks our trigonometry code what the resulting area would be, and successively refines the distance to zero in on the desired area. (The rest of the code in Listing 45-2 is discussed in the next section.)

LISTING 45-2

venn.php

```php
<?php
include_once("trig.php");

$IMAGE_WIDTH = 600;
$IMAGE_HEIGHT = 300;
$CENTER_FINDING_ITERATIONS = 20;

function imagecircle ($image, $center_x, $center_y,
                $radius, $color)
{
  $diameter = $radius * 2;
  imagearc($image, $center_x, $center_y,
     $diameter, $diameter, 0, 360,
          $color);
}

function venn_visualization
  ($left_amount, $left_name,
   $right_amount, $right_name,
   $intersection_amount)
{
  global $IMAGE_HEIGHT, $IMAGE_WIDTH,
         $CENTER_FINDING_ITERATIONS;
  // --- create the image and allocate colors
  $image = imagecreate($IMAGE_WIDTH, $IMAGE_HEIGHT)
    or die("Could not create image");
  $background_color = ImageColorAllocate($image, 255,255,255);
  $left_color = ImageColorAllocate($image, 100, 100, 200);
  $right_color = ImageColorAllocate($image, 200, 100, 100);
  $intersection_color =
    ImageColorAllocate($image, 225, 225, 225);
  $black_color = ImageColorAllocate($image, 0,0,0);

  // --- decide how big the circles should be
  $max_radius = min((($IMAGE_HEIGHT * 0.9) / 3),
             (($IMAGE_WIDTH * 0.9) / 4));
  $center_y = $IMAGE_HEIGHT / 3.0;
  $default_center_x_left = $IMAGE_WIDTH / 4.0;
  $default_center_x_right = (3 * $IMAGE_WIDTH) / 4.0;
```

```
$middle_x = $IMAGE_WIDTH / 2.0;
$radius_left_side_raw =
  area_to_radius($left_amount);
$radius_right_side_raw =
  area_to_radius($right_amount);
$intersection_radius_raw =
  area_to_radius($intersection_amount);
$scale_factor = $max_radius /
                (max($radius_left_side_raw,
                     $radius_right_side_raw));
$radius_left_side = $radius_left_side_raw * $scale_factor;
$radius_right_side = $radius_right_side_raw * $scale_factor;
// (it's convenient to pretend that the intersection area
//  has a radius (although it's not circular) just so we can
//  calculate things the same way as the circles)
$intersection_radius =
  $intersection_radius_raw * $scale_factor;
$area_left_side = M_PI *
                  $radius_left_side * $radius_left_side;
$area_right_side = M_PI *
                  $radius_right_side * $radius_right_side;
$intersection_area = M_PI *
                  $intersection_radius * $intersection_radius;

// We now have all necessary info except where to locate the
//  centers of the circles.
// Four cases:
// 1) no intersection, 2) partial intersection
// 3) left is strict subset of right,
// 4) right is subset of left.

if ($intersection_amount == 0) {
  // No intersection
  $center_x_left = $default_center_x_left;
  $center_x_right = $default_center_x_right;
  $left_fill_x = $center_x_left;
  $right_fill_x = $center_x_right;
  $intersection_fill_x = -1;
}
else if (($intersection_area < $area_left_side) &&
         ($intersection_area < $area_right_side)) {

  // The complicated case --- we must decide where the
  //   circle centers should be so that the overlap is
  //   proportional to the set intersection
  // First, we call a function that decides how far apart
  //   the circle centers need to be.
  $center_distance =
    find_center_distance($radius_left_side,
```

```
                         $radius_right_side,
                         $intersection_area,
                         $CENTER_FINDING_ITERATIONS);

  // Once we know the distance, we place the circle centers
  //   approximately in the middle of the image
  $center_x_left = $middle_x  // left/right middle of image
                    - ($center_distance *
                        ($radius_left_side /
                          ($radius_left_side +
                            $radius_right_side)));
  $center_x_right = $middle_x  // left/right middle of image
                    + ($center_distance *
                        ($radius_right_side /
                          ($radius_left_side +
                            $radius_right_side)));

  // we have decided the sizes and centers of the circles.
  //   Now, we must determine good points to start a
  //   "flood fill" coloring of the three different regions
  $left_fill_x =
    (($center_x_left - $radius_left_side) +
     ($center_x_right - $radius_right_side))
     / 2.0;
  $right_fill_x =
    (($center_x_left + $radius_left_side) +
     ($center_x_right + $radius_right_side))
     / 2.0;
  $intersection_fill_x =
      (($center_x_right - $radius_right_side) +
       ($center_x_left + $radius_left_side))
       / 2.0;
}
else if (($intersection_area == $area_left_side) &&
        ($intersection_area < $area_right_side)) {
  // The right set completely contains the left set
  //   We need to place the left circle somewhere
  //   inside the right circle.
  $center_x_right = $middle_x;
  $center_x_left = $middle_x -
    ($radius_right_side - $radius_left_side) / 2;
  $left_fill_x = -1;
  $right_fill_x =
    (($center_x_left + $radius_left_side) +
     ($center_x_right + $radius_right_side))
     / 2.0;
  $intersection_fill_x = $center_x_left;
}
else if ($intersection_area == $area_right_side) {
```

```php
    $center_x_left = $middle_x;
    $center_x_right = $middle_x +
        ($radius_left_side - $radius_right_side) / 2;
    $right_fill_x = -1;
    $left_fill_x =
      (($center_x_left - $radius_left_side) +
       ($center_x_right - $radius_right_side))
       / 2.0;
    $intersection_fill_x = $center_x_right;
  }

  // now, actually draw and fill regions
  imagecircle($image, $center_x_left, $center_y,
              $radius_left_side, $black_color);
  imagecircle($image, $center_x_right, $center_y,
            $radius_right_side, $black_color);
  if ($left_fill_x > 0) {
    imagefill($image, $left_fill_x,
              $center_y, $left_color);
  }
  if ($right_fill_x > 0) {
    imagefill($image, $right_fill_x,
              $center_y, $right_color);
  }
  if ($intersection_fill_x > 0 ) {
     imagefill($image, $intersection_fill_x,
               $center_y, $intersection_color);
  }
  $left_hand_text = "$left_name ($left_amount)";
  $right_hand_text = "$right_name ($right_amount)";
  $intersection_text = "Intersection: $intersection_amount";
  left_label($image, $left_hand_text, $left_color);
  right_label($image, $right_hand_text, $right_color);
  intersection_label($image, $intersection_text, $black_color);

  // send off the image
  header("Content-type: image/png");
  imagepng($image);
  imagedestroy($image);
}

function left_label ($image, $label_string, $color) {
  global $IMAGE_WIDTH, $IMAGE_HEIGHT;
  imagestring($image, 5,
              ($IMAGE_WIDTH / 4.0 -
                (imagefontwidth(5) * strlen($label_string))
                 / 2),
              $IMAGE_HEIGHT - 55.0,
              $label_string, $color);
```

```
}

function right_label ($image, $label_string, $color) {
  global $IMAGE_WIDTH, $IMAGE_HEIGHT;
  imagestring($image, 5,
              ($IMAGE_WIDTH * 3 / 4.0 -
                (imagefontwidth(5) * strlen($label_string))
                 / 2),
              $IMAGE_HEIGHT - 55.0,
              $label_string, $color);
}

function intersection_label ($image, $label_string, $color) {
  global $IMAGE_WIDTH, $IMAGE_HEIGHT;
  imagestring($image, 2,
              ($IMAGE_WIDTH / 2.0  -
                (imagefontwidth(2) * strlen($label_string))
                 / 2),
              $IMAGE_HEIGHT - 30.0,
              $label_string, $color);
}

function find_center_distance ($r1, $r2, $desired_area,
                                    $iterations) {
  // The greatest possible distance is r1 + r2, and
  //  the smallest is abs(r1 - r2) Let's start in the middle.
  $distance_guess = (($r1 + $r2) + abs($r1 - $r2)) / 2.0;
  $distance_increment = (($r1 + $r2) - abs($r1 - $r2)) / 4.0;
  for ($x = 0; $x < $iterations; $x++) {
    $calculated_area =
      circle_intersection_area($r1, $r2, $distance_guess);
    if ($calculated_area < $desired_area) {
      // move centers closer
      $distance_guess -= $distance_increment;
      $distance_increment *= 0.5;
    }
    else if ($calculated_area > $desired_area) {
      // move centers apart
      $distance_guess += $distance_increment;
      $distance_increment *= 0.5;
    }
    else {
      // unlikely, but ya never know
      break;
    }
  }
  return($distance_guess);
}
?>
```

Display

Now we know exactly where we want to put our circles, and how large they should be. What remains is the graphics code to actually make the display happen. This is also in Listing 45-2 (venn.php).

To produce the graphic, we go through the following steps by using the gd library (which is covered in 42):

1. We create an image by using ImageCreate(). (At this point, the image is not any particular image format, such as PNG or JPEG, but just an internal gd image.)

2. We allocate colors within the image by using ImageColorAllocate(). We care about five colors: the background color (white), a color for borders and regular text (black), a color for the interior of the left-hand circle (which we decide is bluish), a color for the interior of the right-hand circle (reddish), and a color for the intersection (gray). All these colors are specified by using a red-green-blue scale of 0 to 255.

3. We draw the circles in black by using a function of our own, imagecircle(), that takes as arguments the image, the radius, the location of the centers, and a color. (See the "Notes on circles" section about drawing circles in gd.)

4. Now, we want to fill in the three areas (the intersection and the two non-intersection portions of the circles) with the appropriate colors. We use ImageFill() for this, which *flood-fills* outward from a specified point until the fill encounters previously drawn lines. Choosing the starting points for the fills is somewhat tricky because it depends on the different intersection cases. In general, though, we start with a y-coordinate that is the same as the circle centers and calculate an x-coordinate that's right in the middle of the area we are trying to color.

5. We use ImageString() to draw the appropriate labels for each circle, centering each one in the middle of the lower third of the image and in the middle of the appropriate left or right half. We also create and display a count label for the intersection and display it in the middle of the image.

6. Now, we have a complete gd image, and what remains is to ship it off to the user. We send an HTTP header advising the browser that a PNG image is on the way. Then we use ImagePng() to convert the gd image to PNG and send it off.

7. Finally, we call ImageDestroy() to free any resources associated with the temporary image we created. More recent versions of PHP should be handling this already, assuming that the image is of type resource, but either way calling ImageDestroy() does no harm.

Notes on circles

One thing that puzzles people sometimes, if confronted with the gd functions, is that there seems to be no way to draw a circle (or at least there is no function name with circle in its name). This is because there are at least two functions that generalize circle-drawing: imageellipse() (available only with gd 2.0.1 and later) and imagearc. The former draws an ellipse (which can be a circle

if the width and height are the same), and the latter draws a circular arc portion (which can be a circle if you specify a full 360 degrees of arc). In our code, we chose the latter because we wanted to remain compatible with earlier versions of `gd`.

Notes on centering text

As we wrote textual labels in the image code, we actually didn't bother centering the text around any horizontal axis, but we did do some left-right centering. We simply used a built-in numbered font from `gd`, calculated the width of our text as displayed by that font (by using `imagefontwidth()`) and ensured that the left-hand starting point for the text was our desired center minus half the width of the text. This was easy, in part because the built-in fonts we used were monospace, and so `imagefontwidth()` was able to calculate width by referring only to the length of the string. Things get slightly more complicated if you're using a variable-width font — any calculation of string width then needs to know the actual string that is printed, not just the number of characters in it.

Visualizing a Database

We can now produce these Venn-like diagrams on demand, given some numbers and text to start with. Our final task is to hook this up appropriately to a database via a Web form.

> **NOTE** For this application, we assume exactly the same sample MySQL table (`programmers`) as we used in Chapter 42. See Listing 42-1 for a description of the data.

The goal is to let the user choose exactly two restrictions on our database's table, extract counts corresponding to how many rows survive each restriction, count how many rows survive both restrictions, and then pass the results off to our diagramming code.

Listing 45-3 shows code for a form designed around our particular database, mostly just hardcoded HTML. It loads an auxiliary file called `query_clauses.php` (shown in Listing 45-4), which is extremely hardcoded. This file lists and numbers all the restrictions that we want to offer to users, in the form of both an SQL `where` clause and in an English translation.

LISTING 45-3

visualization_form.php

```
<HTML><HEAD><TITLE>DB Visualization</TITLE></HEAD>
<BODY>

<B>Choose one from each column, and we'll<B>
display the intersection from the survey data:<BR>
<FORM METHOD=POST ACTION="db_visualization.php"
     TARGET=_new >
<TABLE>
```

```php
<?php
include("query_clauses.php");
for ($x = 0; $x < count($QUERY_CLAUSES); $x++) {
  print("<TR><TD><INPUT
            TYPE=RADIO NAME=\"left_clause\"
            VALUE=$x>".
            $QUERY_DESCRIPTION[$x] ."</TD>
        <TD><INPUT
            TYPE=RADIO NAME=\"right_clause\"
            VALUE=$x>".
            $QUERY_DESCRIPTION[$x] ."</TD></TR>");
}
?>

</TABLE>
<INPUT TYPE=HIDDEN NAME="table" VALUE="programmers">
<INPUT TYPE=SUBMIT NAME=SUBMIT>
</FORM>

</BODY>
</HTML>
```

Notice that this form is not self-submitting. For this example, we've chosen to completely separate PHP-generated HTML pages from PHP-generated PNG pages and avoid the complexity of embedding images in HTML. We've also chosen a _new target type for the form submission so that the image appears in a new browser window.

LISTING 45-4

query_clauses.php

```php
<?php

$QUERY_CLAUSES = array();
$QUERY_DESCRIPTION = array();

$QUERY_CLAUSES[0] = "sex = 'F'";
$QUERY_DESCRIPTION[0] = "Female";

$QUERY_CLAUSES[1] = "sex = 'M'";
$QUERY_DESCRIPTION[1] = "Male";

$QUERY_CLAUSES[2] = "language = 'PHP'";
$QUERY_DESCRIPTION[2] = "likes PHP";
```

```
$QUERY_CLAUSES[3] = "language = 'Java'";
$QUERY_DESCRIPTION[3] = "likes Java";

$QUERY_CLAUSES[4] = "language = 'Lisp'";
$QUERY_DESCRIPTION[4] = "likes Lisp";

$QUERY_CLAUSES[5] = "language = 'C'";
$QUERY_DESCRIPTION[5] = "likes C";

$QUERY_CLAUSES[6] = "language = 'Perl'";
$QUERY_DESCRIPTION[6] = "likes Perl";

$QUERY_CLAUSES[7] = "os = 'Linux'";
$QUERY_DESCRIPTION[7] = "uses Linux";

$QUERY_CLAUSES[8] = "os = 'Solaris'";
$QUERY_DESCRIPTION[8] = "uses Solaris";

$QUERY_CLAUSES[9] = "os = 'MacOS'";
$QUERY_DESCRIPTION[9] = "uses MacOS";

$QUERY_CLAUSES[10] = "os = 'Windows'";
$QUERY_DESCRIPTION[10] = "uses Windows";

$QUERY_CLAUSES[11] = "age < 20";
$QUERY_DESCRIPTION[11] = "is less than 20 years old";

$QUERY_CLAUSES[12] = "age > 30";
$QUERY_DESCRIPTION[12] = "is over 30 years old";

$QUERY_CLAUSES[13] = "continent = 'North America'";
$QUERY_DESCRIPTION[13] = "lives in North America";

$QUERY_CLAUSES[14] = "continent = 'South America'";
$QUERY_DESCRIPTION[14] = "lives in South America";

$QUERY_CLAUSES[15] = "continent = 'Antarctica'";
$QUERY_DESCRIPTION[15] = "lives in Antarctica";

$QUERY_CLAUSES[16] = "continent = 'Asia'";
$QUERY_DESCRIPTION[16] = "lives in Asia";

$QUERY_CLAUSES[17] = "continent = 'Europe'";
$QUERY_DESCRIPTION[17] = "lives in Europe";
?>
```

A screenshot of the form itself is shown in Figure 45-2.

DB visualization Web form
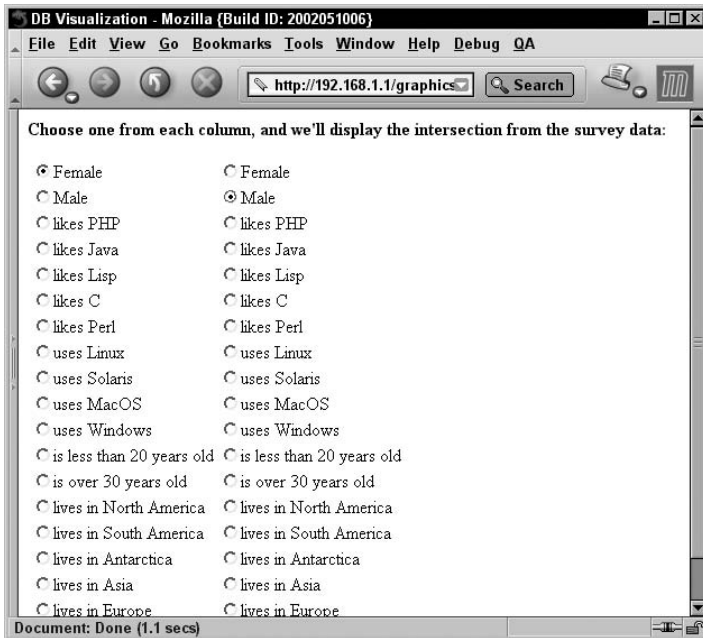
One last code file and we're done. We have our image creation code and a form for requesting an image. The last piece of the puzzle is code to handle the form submission, perform the appropriate counts on the database, and call the image code. This code is shown in Listing 45-5.

LISTING 45-5

db_visualization.php

```php
<?php
include_once("dbconnect.php");
include_once("query_clauses.php");
include_once("venn.php");

if (IsSet($_POST['table']) &&
    IsSet($_POST['left_clause']) &&
    IsSet($_POST['right_clause'])) {
  $table = $_POST['table'];
  $left_clause_id = $_POST['left_clause'];
  $right_clause_id = $_POST['right_clause'];

  $left_clause = $QUERY_CLAUSES[$left_clause_id];
```

```
    $right_clause = $QUERY_CLAUSES[$right_clause_id];

    visualize_intersection ($table, $left_clause,
                                    $right_clause);
}
else {
  print("Form submission not handled correctly.<BR>".
        "Did you choose all options?");
}

function visualize_intersection ($table, $left_clause,
                                    $right_clause)
{

  $left_query = "select count(*) from $table
                  where $left_clause";
  $right_query = "select count(*) from $table
                  where $right_clause";
  $intersection_query =
                  "select count(*) from $table
                  where $left_clause and $right_clause";

  $result = mysql_query($left_query)
    or die("Query was $left_query:" . mysql_error());
  $row = mysql_fetch_row($result);
  $left_count = $row[0];

  $result = mysql_query($right_query)
    or die(mysql_error());
  $row = mysql_fetch_row($result);
  $right_count = $row[0];

  $result = mysql_query($intersection_query)
    or die(mysql_error());
  $row = mysql_fetch_row($result);
  $intersection_count = $row[0];

  venn_visualization($left_count, $left_clause,
                      $right_count, $right_clause,
                      $intersection_count);
}
?>
```

The submission form passes in index numbers of SQL clauses, rather than the clauses themselves, so we don't need to worry about escape characters in the submission. The form-handling code includes the same `query_clauses.php` file, so the index numbers should always agree. The form handler collects the two clauses, creates three SQL statements out of them, executes the statements to get counts, and uses the results as arguments to the `venn_visualization()` function.

Note that Listing 45-5 refers to one auxiliary code file we haven't mentioned yet: `dbconnect.php`. We assume that this file contains (or refers to a file containing) your MySQL username and password, and also makes a call to `mysql_connect()` to create a global DB connection for the rest of the script. Something like the following should suffice:

```php
<?php
$user = 'USER';
$pass = 'PASS';
$db = 'venn';
mysql_connect('localhost', $user, $pass)
    or die("Couldn't make DB connection:" . mysql_error());
mysql_select_db($db);
?>
```

This assumes that your username and password are replaced appropriately, and that you have a MySQL database called `venn`, which contains a table called `programmers`, as described in Chapter 42.

Trying it out

Now that our system is complete, let's give it a spin. Bringing up `visualization_form.php`, we choose the first option from the left-hand column (Female), and the second option from the right-hand column (Male), and submit. The result is shown in Figure 45-3 — two separate circles because, given the database schema, it's impossible for anyone to be both male and female. (Please, no e-mails about the narrowness of our views — it's just an example!)

FIGURE 45-3

No intersection
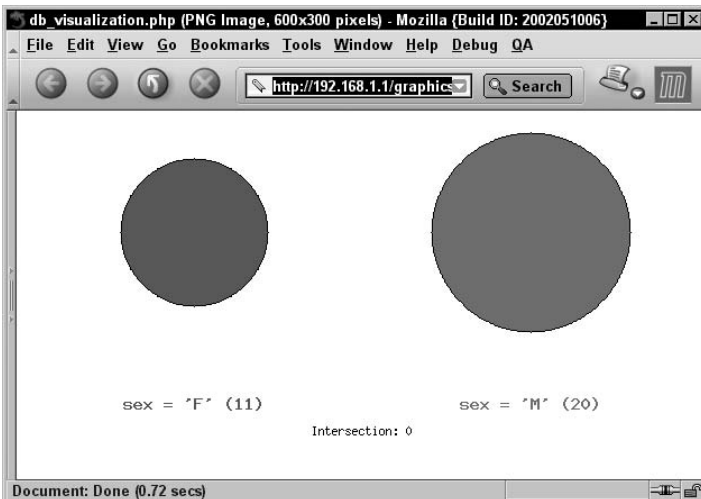
On a color monitor the left-hand circle is bluish, while the right-hand circle is reddish. Because this is a grayscale book, though, you probably see two gray or black circles.

Now a different query: `uses Solaris` on the left, and `lives in Antarctica` on the right. (We chose this deliberately, knowing that our lone South Pole correspondent sees only one kind of Sun during the winter.)

The result is shown in Figure 45-4: one small gray circle inside a larger blue one, indicating that all Antarcticans are Solaris users but not vice versa.

FIGURE 45-4

Subset
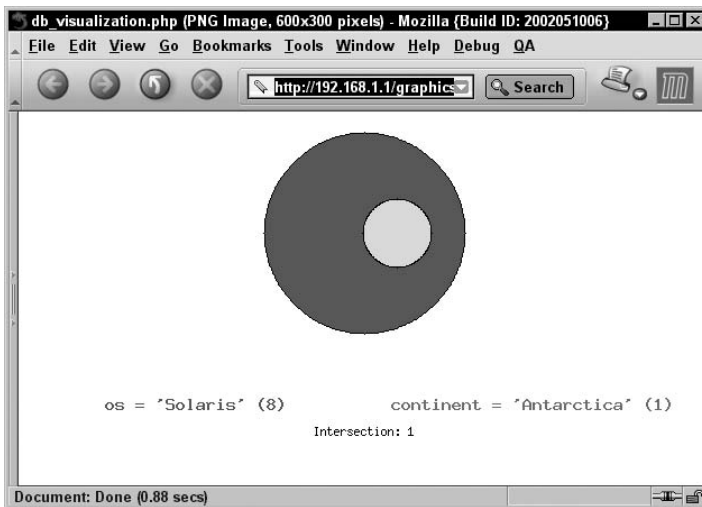
Finally, for the more typical case, let's choose `likes PHP` from the left, and `uses Linux` from the right.

The result, in Figure 45-5, is two mostly overlapping circles. (Again, please no letters — we have no idea if these proportions match the world as it is or just the world as we would like it to be.)

Partial intersection
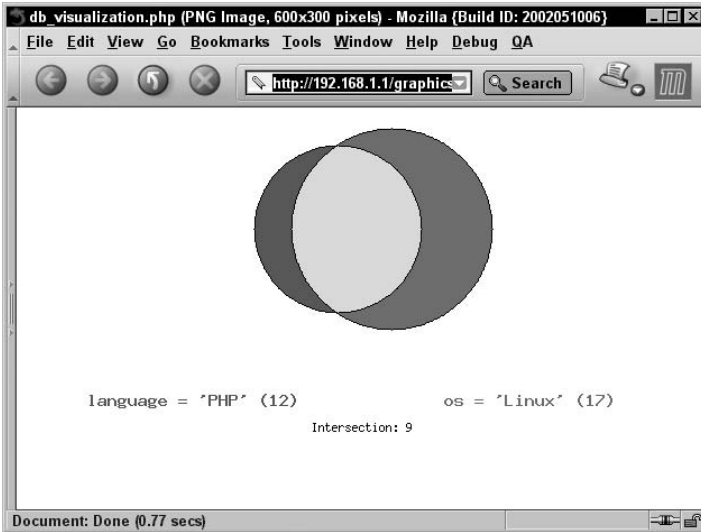
Extensions

This example works nicely, but as always, there are countless ways in which it could be tweaked, improved, and especially extended. Naturally, there is a lot you could do to change the cosmetics of the images or to make the look more configurable.

The weakest part right now in our view is the form submission, which is hardcoded to pertain only to a particular known MySQL table. Much cooler would be code that, armed only with a table (or view) name and the appropriate login, would quiz the database about column names and types and the distribution of values, and then develop such a form on its own.

One extension that may immediately occur to you doesn't work, unfortunately, at least without substantial changes to the display code. There is no good way to involve a third set (and circle) in the diagram, cover all the cases, and be assured that all the shapes can still be circular. If you have any doubt about this, try diagramming the following three sets: people who were born in Europe, people who currently live in Europe, and people who currently live in a continent different from the one they were born in.

Summary

We've shown you a small but complete system for visualizing data in a MySQL database. It allows the user to select aspects to compare, makes corresponding SQL queries, and transforms the results into a scaled Venn diagram, showing how the database records overlap.

In addition to using MySQL techniques from Part II of this book, we drew on the image techniques from Chapter 42, and as little math as we could get away with from Chapter 27.

Appendix A

PHP for C Programmers

I n this appendix, we assume that you have more C (or C++) programming experience than PHP experience and are looking to get up to speed in PHP quickly. First, we'll provide a quick overview of PHP from a C perspective; next, we'll break down the similarities and differences, and finally we'll point out which parts of the book you are likely to benefit from the most.

The simplest way to think of PHP is as interpreted C that you can embed in HTML documents. The language itself is a lot like C, except with untyped variables, a whole lot of web-specific libraries built in, and everything hooked up directly to your favorite web server. The syntax of statements and function definitions should be familiar, except that variables are always preceded by `$`, and functions do not require separate prototypes.

Similarities

In this section, we offer some notes (by no means exhaustive) on ways in which PHP can be expected to be C-like.

Syntax

Broadly speaking, PHP syntax is the same as in C: Code is blank-insensitive, statements are terminated with semicolons, function calls have the same structure (`my_function(expression1, expression2)`), and curly braces (`{` and `}`) make statements into blocks. PHP supports C and C++ style comments (`/* */` as well as `//`), and also Perl and shell-script-style (`#`) ones.

Operators

The assignment operators (=, +=, *=, and so on), the Boolean operators (&&, ||, !), the comparison operators (<, >, <=, >=, ==, !=), and the basic arithmetic operators (+, -, *, /, %) all behave in PHP as they do in C.

Control structures

The basic control structures (`if`, `switch`, `while`, `for`) behave as they do in C, including supporting `break` and `continue`. One notable difference is that `switch` in PHP can accept strings as case identifiers.

Many function names

As you peruse the documentation, you'll see many function names that seem identical to C functions. It's a safe bet that these functions perform the exact same tasks, although they may sometimes take a slightly different form in terms of arguments or the way results are returned. Most string-modifying functions, for example, return new strings as the value of the function rather than modifying a string passed as an argument. Note, however, that function names are not case-sensitive in PHP.

Differences

Although PHP has quite a bit of C ancestry, it also has some other ancestors (Perl, shell scripts), as well as some unique features not at all C-like.

Those dollar signs

All variables are denoted with a leading $. Variables do not need to be declared in advance of assignment, and they have no intrinsic type — the only type a variable has is the type of the last value assigned to it. The PHP version of the C code:

```
double my_number;
my_number = 3.14159;
```

would simply be:

```
$my_number = 3.14159;
```

Types

PHP has only two numerical types: *integer* (corresponding to a long in C) and *double* (corresponding to a double in C).

Strings are of arbitrary length. There is no separate character type. (Functions that might take character arguments in their C analogues typically expect a one-character string in PHP (`ord()`, for example.) Beginning with PHP4, there is also a genuine Boolean type (`TRUE` or `FALSE`). See the following sections for arrays and objects.

Type conversion

Types are not checked at compile time, and type errors do not typically occur at runtime either. Instead, variables and values are automatically converted across types as needed. This is somewhat analogous to the way arithmetic expressions in C will "promote" numerical arguments as needed, but it is extended to the other types as well. (See Chapter 25 for details of the conversion rules.)

Arrays

Arrays have a syntax superficially similar to C's array syntax, but they are implemented completely differently. They are actually associative arrays or hashes (with some additional supporting machinery), and the "index" can be either a number or a string. They do not need to be declared or allocated in advance.

No structure type

There is no *struct* in PHP, partly because the array and object types together make it unnecessary. The elements of a PHP array need not be of a consistent type.

Objects

PHP4 had a very basic OOP syntax, which allowed definition of classes with member data items and member functions. PHP5 introduces a much fuller object model, although in approach and syntax it owes more to Java than to C++. Some highlights: abstract classes, private/protected members, constructors/destructors, and interfaces (but no multiple inheritance as in C++).

No pointers

There are no *pointers* per se in PHP, although the typeless variables play a similar role. PHP does support variable references. You can also emulate function pointers to some extent, in that function names can be stored in variables and called by using the variable rather than a literal name.

No prototypes

Functions do not need to be declared before their implementation is defined, as long as the function definition can be found somewhere in the current code file or included files.

Memory management

The PHP engine is effectively a garbage-collected environment (reference-counted), and in small scripts there is no need to do any deallocation. You should freely allocate new structures — such as new strings and object instances — especially because they will reliably go away when your script terminates. If you need to free memory within a script's execution, call `unset()` on the variable that refers to it, which will release the memory for collection. External resources (such as database result sets) can also be explicitly freed within a script, but doing so is worth it only if the script would use an unacceptable amount of the resource before terminating.

In PHP5, it is possible to define destructors for objects, but there is no `free` or `delete`. Destructors are called when the last reference to an object goes away, before the memory is reclaimed.

Compilation and linking

There is no separate compilation step for PHP scripts — the development cycle is simply edit-reload. Errors and warnings show up in the browser output by default, although there is also an error-logging capability. Typically, there is no dynamic loading of libraries (although such a capability exists) — you decide at PHP configuration time which function families to include in your module, and they are then available to any script.

Permissiveness

As a general matter, PHP is more forgiving than C (especially in its type system) and so will let you get away with new kinds of mistakes. Unexpected results are more common than errors. In particular, under the default error-reporting level, PHP does not warn you if you use a variable that has not yet been assigned (although it does supply reasonable default values rather than garbage). If you would rather be warned, you can set the error-reporting level by evaluating `error_reporting(E_ALL)` early in your script, or set the error-reporting level to `E_ALL` permanently by editing the `php.ini` file.

Guide to the Book

In writing this book, we very intentionally did not assume that the reader had prior knowledge of C. Because PHP resembles C in many aspects, some of the chapters may cover familiar ground. This is especially true of Part I, which is essentially a language introduction.

In Table A-1, we label the chapters of Part I according to how familiar they are likely to be to C programmers. Parts III, IV, and V are more PHP-specific and likely to be novel, but you may also find portions of Part II to be familiar if you have some experience with SQL databases.

Guide to Part I for C Programmers

Chapter	Chapter Title	Verdict?	Notes
1	Why PHP?	Novel	The chapter you need to justify PHP to your boss.
2	Server-Side Scripting Overview	Novel	Important if you have not seen web-scripting languages before.
3	Getting Started with PHP	Novel	Installation, hosting, and so on.
4	Learning PHP Syntax and Variables	Mostly familiar	Skimmable until the section on variables (which really are different in PHP).
5	Learning PHP Control and Functions	Mostly familiar	All the PHP control structures (`if`, `while`, `switch`, `for`) work the same way as in C. Some differences in function behavior, particularly with scoping of variables.
6	Passing Information with PHP	Novel	Specific to web-scripting.
7	Learning PHP String Handling	Mostly familiar	Doubly quoted strings do automatic interpolation of variable values.
8	Learning Arrays	Novel	Deceptively familiar — PHP arrays are syntactically like C arrays but behave totally differently.
9	Learning PHP Number Handling	Familiar	Two numerical types, corresponding to the long and double types. Numerical operators are as in C.
10	PHP Gotchas	Novel	Error messages and stumbling blocks do not have much overlap with C.

A Bonus: Just Look at the Code!

As a final bonus, C programmers are uniquely qualified to benefit from the open source nature of PHP. Although the combination of this book and the online manual should answer almost all your questions, if you have the PHP source available you may be able to gain some extra insight by poking around in it and seeing how things are implemented. Although you would need to be familiar

with lexing/parsing technology to get much out of the parser code itself, many PHP functions are simple wrappers around their similarly named C counterparts, and some others that have no C counterparts are at least implemented in clear and simple C.

It's also easy for C programmers to add new capabilities to the language, whether for their own use or for eventual use by the PHP community. Most PHP programming tasks are addressed by writing PHP (often by defining functions in PHP), but you can also pop the hood and add functionality to the underlying language by adding a new module, written (of course) in C.

Appendix B

PHP for Perl Hackers

I n this appendix, we assume that you know Perl, but not PHP, and are
looking to quickly get up to speed in PHP. The good news is that the
two languages are very similar indeed.

This is by no means a comprehensive guide to how Perl and PHP compare.
Although similar, and sharing some ancestry, they really are distinct lan-
guages with distinct syntaxes and feature sets, and there is no replacement
for getting to know them individually. Our main goal for this appendix is
to save you some time up front — to warn you, for example, that `elsif`
means nothing in PHP (but that `elseif`, however, is significant) rather
than letting you debug your way to that realization.

Similarities

In this section, we discuss some ways in which PHP and Perl are similar.

Compiled scripting languages

First, the obvious: Both Perl and PHP are scripting languages. This means
that (unlike compiled languages such as C) they are not used to produce
native standalone executables in advance of execution, which can then be
run without reference to the language they were written in. Instead, Perl or
PHP source files are both fed to an appropriate engine at execution time.
This does not mean, however, that Perl/PHP code is interpreted line by line
at execution time; in both Perl and PHP, scripts are quickly and automati-
cally compiled at execution time and then executed. But it does mean that
the development cycle for PHP/Perl programmers is edit-execute, rather
than edit-compile-execute, as in C.

Syntax

PHP's basic syntax is very close to Perl's, and both share a lot of syntactic features in common with C. Code is insensitive to whitespace, statements are terminated by semicolons, and curly braces organize multiple statements into a single block. Function calls start with the name of the function, followed by the actual arguments enclosed in parentheses and separated by commas.

Dollar-sign variables

All variables in PHP look like scalar variables in Perl: a name with a dollar sign ($) in front of it. (See the "Differences" section for what happened to @ and %.)

No declaration of variables

As in Perl, you don't need to declare the type of a PHP variable before using it. The following line is legal in both languages, with no prior mention of the variable called $the_answer:

```
$the_answer = 42;
```

Loose typing of variables

As in Perl, variables in PHP have no intrinsic type other than the value they currently hold. This is different from languages such as C and Java, in which once a variable is declared to be for holding, say, strings, you get into trouble if you try to use it to store an integer.

The following sequence of two lines is legal in both Perl and PHP:

```
$the_answer = 42;
$the_answer = "the answer";
```

The variable $the_answer is assigned sequentially to an integer and a string. This would not be legal in more strongly typed languages such as C, Pascal, or Java.

Strings and variable interpolation

Both PHP and Perl do more interpretation of double-quoted strings (`"string"`) than of single-quoted strings (`'string'`). In particular, the value of $ variables is spliced into double-quoted strings at the time the strings are read.

The following code fragment is both legal Perl and legal PHP, and it has the same behavior in both languages.

```
$the_answer = 42;
$the_statement = "the answer is $the_answer";
```

At the end of the second line, the variable $the_statement contains the string:

```
the answer is 42
```

Differences

This section warns you (again, not exhaustively) about some ways that Perl and PHP diverge from each other.

PHP is HTML-embedded

Although it is possible to use PHP for arbitrary tasks by running it from the command line, it is more typically connected to a web server and used for producing web pages. The code for these pages can consist partially (or even completely!) of straight HTML, with fragments of PHP embedded in them to produce the dynamically generated portions.

If you are used to writing CGI scripts in Perl, the main difference in PHP is that you no longer need to explicitly print large blocks of static HTML using `print` or heredoc statements and instead can simply write the HTML itself (outside of the PHP code block). Also, for typical pages, there is no need to explicitly send HTTP headers from PHP code.

No @ or % variables

PHP has one only kind of variable, which starts with a dollar sign ($). Any of the datatypes in the language can be stored in such variables, whether scalar or compound. For example, the expression `$my_array[0]` refers to the first element in an array, while `$my_array` refers to the array itself.

Arrays versus hashes

PHP has a single datatype called an *array* that plays the role of both hashes and arrays/lists in Perl. For all the details, see Chapter 9; the short version is that a PHP array acts like a Perl hash when you supply keys and like a Perl (nonassociative) array when keys are omitted. Values can be extracted by key (as in Perl hashes) or by iteration through the array (as with Perl arrays).

There is, however, a `list()` function in PHP. It's used to extract the contents of an array into a set of separate variables. The `list()` function works like this:

```
$myArray = ('a', 'b','c');
list($var1, $var2, $var3) = $myArray
```

After that runs, `$var1` contains a, `$var2` contains b, and `$var3` contains c.

Specifying arguments to functions

Function calls in PHP look pretty much like subroutine calls in Perl. Function definitions in PHP, on the other hand, typically require some kind of list of formal arguments as in C or Java. For example, although the typical syntax for a two-argument subroutine in Perl might look like:

```
sub two_arg_sub () {
    my ($first_arg, $second_arg) = @_;
```

```
        ...
    }
```

The corresponding PHP function definition would be:

```
    function two_arg_function ($first_arg, $second_arg) {
        ...
    }
```

Although your humble authors try hard not to be partisan, we feel strongly that subroutine arguments in Perl are a bug and that function arguments in PHP are a feature. Two kinds of silliness common in Perl that don't usually arise in PHP are: (1) popping the argument stack at various points in a subroutine (so that it is hard for a reader of the code to figure out what the formal arguments are supposed to be) and (2) compound arguments bleeding into one another because arguments are passed as a list of scalars. PHP arguments arrive intact and without confusion regardless of number and type.

Variable scoping in functions

In Perl, the default scope for variables is global. This means that top-level variables are visible inside subroutines. Often, this leads to promiscuous use of globals across functions.

In PHP, the scope of variables within function definitions is local by default. This means that (with some exceptions) the only variables visible within functions are the formal parameters and variables assigned locally within the function. If you want to refer to a variable from the global context within a function, you must declare it by name in the function definition itself, using the `global` keyword.

The exceptions to that rule in PHP are the so-called superglobals, the most popular of which is `$_POST`. Representing the variables that arrived at a PHP script as a result of an `HTTP POST` operation, `$_POST` and its contents are accessible anywhere, without the need to use the `global` keyword.

For example, if we called a PHP script with this URL:

```
    http://localhost/quoteStooges.php?stooge1=Curly&stooge2=Shemp
```

these calls would be legal anywhere in: `quoteStooges.php`:

```
    lookupQuote($_GET['stooge1']);
    lookupQuote($_GET['stooge2']);
```

No module system as such

In PHP there is no real distinction between normal code files and code files used as imported libraries. To import a PHP code file full of function or class definitions, simply use `include()`, `require()`, or `require_once()`, which have much the same effect as splicing the definitions in at the point of the statement.

Break and continue rather than next and last

Perl has some idiosyncratic language keywords, which are not the same as the corresponding C constructs. In general, if Perl and C disagree about such a name, you will find that PHP follows C rather than Perl. In particular, if you want to skip to the end of a `for` or `while` iteration, use `continue` (not `next`); if you want to break out of the loop altogether, use `break` (not `last`).

No elsif

A minor spelling difference: Perl's `elsif` is PHP's `elseif`. Also, its use is not mandatory. In PHP, the following is legal:

```
if ($boolean_var) {
    # case 1
}
else if ($other_boolean) {
    # case 2
}
```

In Perl, on the other hand, if you for some reason decline to use `elsif`, your other alternative is the more awkward form:

```
if ($boolean_var) {
    # case 1
}
else {
  if ($other_boolean) {
      # case 2
  }
}
```

More kinds of comments

Perl people like the phrase "There's more than one way to do it," and yet they suffer with a really impoverished set of options for comments. In addition to Perl-style (#) single-line comments, PHP offers C-style multiline comments (`/* comment */`) and Java-style single-line comments (`// comment`).

Regular expressions

PHP does not have a built-in syntax specific to regular expressions, but has most of the same functionality in its "Perl-compatible" regular expression functions. See Chapter 22 for the details.

Miscellaneous Tips

Following are answers to a couple of questions that Perl programmers might have on their minds.

What about use of strict "vars"?

Like Perl, PHP allows you to use variables without declaring them or initializing them, and (as in Perl) this capability is a frequent source of bugs. If you would like a declaration like Perl's `use strict "vars"`, try `error_reporting(E_ALL)`, which will at least warn you about the use of any unassigned variables.

Where's CPAN?

PHP doesn't yet have a code repository as comprehensive as the Comprehensive Perl Archive Network (CPAN), but the PEAR project provides a reasonable substitute (`http://pear.php.net`).

Guide to the Book

As in Appendix A (PHP for C Programmers), in this section we offer Perl hackers a quick guide to Part I of the book to give a sense of which chapters are likely to already be familiar.

TABLE B-1

Guide to Part I for Perl Programmers

Chapter	Chapter Title	Verdict?	Notes
1	Why PHP?	Novel	The chapter you need to justify PHP to your boss
2	Server-Side Web Scripting	Possibly novel	Familiar to Perl CGI and mod_perl programmers; important if you have not seen web scripting before
3	Getting Started with PHP	Novel	Installation, hosting, and so on
4	Learning PHP Syntax and Variables	Novel but easy	"Hello world" for PHP
5	Learning PHP Control and Functions	Somewhat familiar	Basic constructs are similar, with syntactic differences
6	Passing Information with PHP	Mostly familiar	Specific to web scripting
7	Learning PHP String Handling	Somewhat familiar	Treatment of single-quoted and double-quoted strings essentially the same; functions are mostly novel

Chapter	Chapter Title	Verdict?	Notes
8	Learning Arrays	Novel	No exactly corresponding data type in Perl
9	Learning PHP Number Handling	Mostly familiar	Novel section on arbitrary-precision (BC) math functions.
10	PHP Gotchas	Novel	Same gotchas around unintentionally unassigned variables; other gotchas specific to PHP parsing

Appendix C

PHP for HTML Coders

This appendix contains specific advice for HTML-only jocks looking to trade up to something a little more powerful on the server side. If you already know ASP, JavaScript, or almost any real programming language, this appendix is not going to help you much.

The Good News

If you're already proficient at HTML, starting to use PHP is not a huge step. Because PHP is usually embedded in HTML, extending the functionality of static web pages with a programming language can be a very natural progression. There are plenty of reasons to believe that you can learn PHP fairly quickly, such as the following factors.

You already know HTML

Because PHP is often embedded in HTML, and because PHP generally uses HTML for display to the browser, you won't be able to see anything that your scripts are doing unless you output HTML. In fact, you can think of PHP as simply adding functionality to web pages — it can do other things, but lots of people use it mostly for form handling and dynamic page generation.

You presumably have a lot of practice debugging HTML, which is all to the good. Many errors occur within the HTML parts of scripts or during the transitions between modes, so the ability to read and write HTML with great facility is crucial.

If you're strong on the design side, as are many HTML coders, you have the ability to produce a good-looking and well-laid-out product. This skill is important for the community because a lot of early PHP developers were not exactly known for their UI skills (including ourselves, we hasten to admit). So go out there and show the world that PHP sites don't need to be ugly, clunky, or, at best, really plain — you can make us all proud.

PHP is an easy first programming language to learn

Unlike many major programming languages, PHP enables you to do useful stuff from the very beginning instead of making you play endless games of tic-tac-toe or code up incomprehensible math problems. The web browser and markup languages, however primitive and clunky they are now, point the way to the universal I/O, windowing, and multimedia solution that the world has been waiting for. PHP takes full advantage of the web's power; plus it has very little overhead and takes a loose, inclusive approach to issues such as types, variables, and syntax. All the nitpicky angst that programmers used to put into these areas you can now apply more directly to functionality.

And frankly, PHP enables you to learn just those parts that are useful to you and ignore the rest. Unlike some programming languages, which pretty much require a firm grasp of all the basic principles before you can do anything useful (try telling a C programmer that you just haven't seen a need for structs yet if you want a laugh), no one is going to give you a quiz on all the hundreds of PHP functions before entrusting you with a text editor. So if you don't need to write some huge math function right off the bat, go ahead and skip that chapter — we promise not to tell. If you ever need them, the math capabilities will still be there.

Web development is increasingly prefab anyway

Finally, the web is increasingly making development a matter of altering prefab open source code rather than hacking it all up yourself. Much of this work is about changing how the page looks rather than how it functions. Learn to be a smart script shopper, and you're more than halfway there.

The Bad News

Before we get too carried away, honesty compels us to admit that you may face a few hurdles before you become a power PHP user.

If programming were that easy, you'd already know how

PHP is a real programming language, similar to C (albeit generally web server dwelling), rather than a tag-based markup concept such as HTML or ColdFusion. This point introduces whole new levels of complexity. It simply takes time and practice to develop a bag of tricks, work out routines for solving problems, and just get better at development — and there are no shortcuts for these skills.

So here's the bottom line: Most of PHP is probably completely new to you. Unlike new PHP developers who are already proficient with ASP, JavaScript, or C, you can't expect to pick up any specific points here that are highly similar to things you already know how to do. Uh, sorry.

However, if you already know some JavaScript or have taken an "Intro to C" class in school — even if you wouldn't describe yourself as a JavaScript or C guru — you're ahead of the curve. Some of the logic is sure to come back to you as you begin to work with PHP.

Backend servers can add complexity

PHP is mostly useful in conjunction with backend servers, such as database and mail servers, which have their own syntax and implementation issues that you need to learn about. Because open source software such as PHP is commonly used in noncorporate settings, most PHP developers probably don't have the luxury of a team of database, network, and design experts doing their various things while they just worry about the middle and front tiers.

If possible, don't try to learn everything at once. The most important task is to become comfortable with the web server itself; Apache, in particular, is an extremely powerful but involved piece of software that rewards study. (This advice may not be relevant if you have IT staff to install and maintain the web server for you.) After that, you will almost certainly want to learn SQL if you don't know it already. Mail service is also a very rewarding subject. After you master those three, other new servers should be easier to learn.

Concentrate On . . .

Learning PHP quickly requires a strategy. Here are some things you might want to concentrate on doing when you're first learning.

Reading other people's code

Learning to read other people's code can be harder than it sounds. One of the best things about PHP is its loose syntax and inclusive "don't worry, be happy" design — but that can also mean that different scripts can look very different, even if they return similar results. Beginners can be boggled by stylistic issues and may find it difficult to sort out which parts of a script are functionally irreducible and which are the products of one individual's programming quirks. But regardless of difficulty, the sooner you can parse other people's PHP and the more code you can look at, the better off you'll be.

One potentially helpful exercise is to visit the mailing-list archive or a code exchange (see Appendix D) and print out multiple examples of code that solve the same issue (preferably one you're interested in). Then lay the sheets side by side, take a big ol' red pen and go through it all, circling the common parts. Give extra brownie points to any scripters who comment their code well (which doesn't necessarily mean the most voluminous comments but rather the most useful) and look for more code from those people. Also look for code that is generally well laid out and logically organized, even if it isn't extensively commented.

Working on what interests you

We firmly believe that learning is motivation. If you find tasks that you want to accomplish, you will automatically be motivated to learn what you need to know to accomplish them. Don't let anyone tell you that the right way to learn programming is through some highly structured program of math problems, games, and stock-market simulations. Wanting to put pictures of your dog on the Internet is much more important than making sure that you know what's in every byte of memory your program is using.

Thinking about programming

As we said earlier, learning PHP is inevitably going to take time, practice, and lots of example code. There is just no way around it, and there is not a whole lot more to say on the subject.

One thing that may prove helpful to new developers, particularly those of a narrative rather than mathematical bent, is the judicious use of *pseudocode*. For example, you might start out with mostly pseudocode and gradually add real PHP as we do in the steps that follow:

1. Write down the tasks that you want this page to accomplish. Being complete is more important at this stage than being cogent. Following is an example of a script mission statement:

```
This page should display a form with any old answers already filled in.
Then you can submit the form to update some of the values if you want
to. And I want the form to be password-protected, so it needs to handle
a User ID passed in from the login screen.
```

2. Break the mission statement down into steps and substeps, as in a recipe. Rearrange these if necessary. Following is an example of this step:

```
1. Get the User ID that is passed from the login screen; if none,
   don't display anything.
2. Display the HTML form.
3. Make any old values from the database appear in the form.
a. Connect to the database server.
b. Download data about this item.
c. Put the data into the HTML form's "value=X" variables.
4. Change the values and put them into the database too.
5. Pass the User ID to the next page.
```

3. Pick one of the steps and turn it into actual PHP code. Starting with a core PHP task — such as sending e-mail or returning something to the screen — is generally easier than beginning with peripheral tasks, such as connecting to a database. Any time you might want to connect to a database, use a commented variable, array, or include file for the moment. The following example illustrates this step:

```
1. Get the User ID that is passed from the login screen; if none,
   don't display anything.
```

```php
<?php
// Dummy UserID pretending to be passed from login.
// Will be superseded later.
$UserID = 1;
?>

2. Display the HTML form.

<HTML><HEAD></HEAD>
<BODY>
<FORM>
First name: <input type="text" size=30 name="FirstName"><BR>
Last name: <input type="text" size=30 name="LastName"><BR>
E-mail: <input type="text" size=30 name="Email">

3. Make any old values from the database appear in the form.

<?php
// I'm using these variables now, but later I'll get
// them from the database instead.

$FirstName = "Joyce";
$LastName = "Park";
$Email = "root@localhost"
?>
Oh, I think I need to put them in before the form is rendered.

4. Change the values and put them into the database, too.
5. Pass the User ID to the next page.
```

4. Gradually fill in more and more of the code, fixing any new issues that arise. You may want to keep some of the pseudocode, suitably edited, as comments, as shown in the following example:

```
/* Pass the User ID to the next page. The best way is to have it
   show up as a hidden input type and PHP variable in the form;
   then HTML can pass it with the rest of the POST values. */
```

Learning SQL and other protocols

Spending some time interacting with backend servers directly, via whatever interface the server provides, is generally a good idea before you add the complexity of PHP between you and the server.

You can kill two birds with one stone by using the backend server's own interface to construct the database (or whatever), even though there may be nice PHP tools for some of these tasks. For instance, even though the `phpMyAdmin` and MySQL Control Center are both very slick and handy ways to deal with the MySQL database, the newbie database administrator can learn a heck of a lot more by using MySQL's deliberately primitive command-line interface.

813

Beginning with PHP5, a lightweight, embedded database engine called SQLite is included and ready for you to use. SQLite is rudimentary and not up to the task of replacing MySQL and other database engines on a production site; however, you may find it an excellent tool for getting your feet wet before moving on to MySQL.

Making cosmetic changes to prefab PHP applications

One common way to ease into using PHP is to enlist your frontend web development skills to customize a preexisting PHP package.

■ First try just changing the colors — that's generally pretty safe. If that goes well, try customizing the buttons. The next safest thing is spacing — table widths, columns, and so on. You can also add graphics, add links, or play around with style sheets pretty much without worries.

■ If the application has include files (especially `header.inc`), the cosmetic part is often in there. Look first in headers and footers for colors, the basis of page layouts, and so on. Remember to match header changes with corresponding footer changes and vice versa.

■ Never, ever erase a line beginning with a conjunction (such as `if`, `while`, or `for`). If you are not 100 percent sure of what you're doing, comment out code blocks rather than deleting them.

Debugging is programming

Few people truly enjoy debugging; as one of our colleagues once observed, "I'd rather implement new features than eat someone else's leftovers." Debugging can turn out to be a useful learning experience, however, because you can fix things at the edges of a big project instead of jumping into writing the whole thing from scratch.

One of the most efficient ways to debug is in pairs. If you're tired or have seen a piece of code too many times, focusing on every detail can prove difficult. At this point, talking through your logic becomes very helpful — one of you briefly stating why you're doing what you're doing and the other checking each step off very deliberately. A fresh set of eyes can often find cheap mistakes such as misspelled variable names or missing brackets more quickly, too. If you get an opportunity to debug with a more experienced programmer, take it.

Avoid at First . . .

A few things in PHP are extremely unfamiliar to HTML coders and generally are not extremely necessary to writing functional PHP. Try to avoid the elements that we describe in the following sections if you can, at least at first.

Maximal PHP style

See Chapter 32 on PHP style. The maximal style is deprecated by Rasmus Lerdorf himself, and only hardcore C programmers have the slightest excuse to use it, except in very specific, brief instances. It includes too many single quotation marks, double quotation marks, forward slashes, backslashes, ASCII line breaks, and HTML line breaks for most coders. Beginners are better off if they don't waste their time worrying about stray punctuation when they could be spending the time and effort grasping larger concepts instead.

Programming large applications from scratch

Why reinvent the wheel? In Open-Source-Land, you don't need to. Becoming a good customizer and recycler of other people's code is often more efficient than trying to become the world's greatest programmer from scratch. Learn to shop for what you need.

Consider This . . .

The ideas in the following sections are completely optional but may prove helpful. You may not agree with them all, but we offer them for what they're worth.

Reading a book on C programming

Unfortunately, we can't write a complete programming tutorial. Part I of this book explains programming topics but necessarily very briefly. We've tried to comment our code samples extensively, but we can do only so much to explain these techniques in passing.

Mailing-list regulars frequently counsel new PHP developers to buy a book on C programming — but in a snotty, RTFM way that too often elicits a naturally passive-aggressive response. Nevertheless, if you separate the suggestion from the unspoken message that you must be a clueless idiot, it's good advice and something to seriously consider if you're having trouble with the programming aspect of PHP.

A clearly written, brief tutorial book is Patrick Henry Winston's *On to C* (Addison-Wesley, 1994). It's fewer than 300 pages, and a lot of the PHP-relevant material is right at the beginning. The standard reference is *The C Programming Language* by Brian W. Kernighan and Dennis M. Ritchie (Prentice Hall, 1988), which is quite definitive but more reference-oriented and, therefore, perhaps less appropriate for HTML-only coders. A friendlier introduction might be Dan Gookin's two-volume *C For Dummies* (Wiley, 1997), which has cartoons and dry humor.

Minimal PHP style

Of the range of PHP styles, the easiest for the HTML coder to work with at first is the most minimal. In other words, we suggest that you separate the HTML and PHP sections completely. Not only does this technique avoid many stylistic difficulties, but by using this method, you avoid mixing PHP and

HTML glitches on the same page, which makes diagnosing problems more than twice as hard. We discuss this topic thoroughly in Chapter 32.

Perhaps the easiest way to use the minimal style is to finish the HTML pages first, using whatever tool you're most enamored of. Take the time to debug your HTML completely and perhaps run it through a tidying utility. Then, and only then, tackle the PHP parts, secure and comfy in the knowledge that any difficulties you encounter are certainly on the PHP side rather than the HTML side.

One downside of this style is that you can't have pages pass their variables back to themselves. This is particularly relevant with forms, so if your site has a lot of forms, you may want to change your style a bit as your PHP skills improve.

Use the right tools for the job

Finally, you want to consider using a PHP-enabled text editor for the PHP parts of your scripts. (See Chapter 3 for a discussion of text editors versus WYSIWYG tools.) Some people can do wonders with just Notepad or `emacs`, but a lot of frustrated beginners are certainly using those tools just because someone told them that's what the cool programmers do. As Zsa Zsa Gabor said (in a slightly different context), macho does not mean mucho. If you work more effectively with `vim` or Visual SlickEdit, by all means use those tools.

Syntax highlighting (printing different parts of the code in different colors) can help you a lot, as it will usually make clearer when you have failed to insert a closing parenthesis or double quotation mark. Some programmer's text editors will automatically line up your curly braces for you, which helps you verify that you've closed off all of a set of nested cases. We would not even consider using a text editor without line numbering, since PHP's built-in error messages always refer to lines in your code. And last but not least, do not forget the power of "View Source" in the browser — this will help you verify that the output you are producing is in fact what you intended.

CAUTION   This advice does *not* apply to WYSIWYG editors, the use of which we deprecate. Sooner or later, you need to fix up the HTML into a human-readable form, which no WYSIWYG editor can yet produce. If it's your choice to use one, fine — but you should in no way think of this tool as a substitute for understanding and writing clean HTML by hand.

Appendix D

PHP Resources

This appendix lays out some basic resources that can help you learn more about the language. We have also tried to mention specific resources and products throughout the text.

The PHP Web Site

The URL for the PHP web site (engrave it on your forehead) is `www.php.net`. Here you'll find the latest official news, the freshest downloads, the PHP bug-tracker, a growing list of PHP users' groups, and links to PHP-friendly ISPs.

Most important, you'll find the PHP manual in the Documentation section of the PHP web site. It's available in several versions for your universal reference pleasure, including the following:

- Numerous translations — many European languages, Japanese, Korean, and so on

- Several PDF, Windows Help, and HTML download versions (useful when you're traveling; HTML versions are included with the PHP download)

- Two versions for Palm OS

- A plain HTML online version

- Links to the PHP-GTK (client-side PHP) and PEAR (PHP codebase) manuals

- Information on beta releases that are not yet ready for production use

A comprehensive listing of all downloads is available at `www.php.net/download-docs.php`.

The PEAR manual, which used to be a chapter of the general PHP manual, has now been moved to its own server. You can find it at `http://pear.php.net/manual.`

But when people talk about "the PHP manual," they mean the big annotated online version for which PHP is famous. Users from around the world have added notes and comments to each page. These additions are often clarifications of points made in the main text, additional insights, and reports of PHP's behavior on various platforms.

Navigating the PHP Manual

Our experience suggests that something about the organization of the manual makes it difficult or discouraging for many users to quickly find what they're looking for. To locate information about a function or construction, you sort of need to know that it's there and what kind of thing it is (not always clear from the headings — for example, date functions versus calendar functions), and often it helps to have an idea of what the PHP team might have named it. If you can't accurately guess these three things, you're likely to spend a lot of time wandering around looking at not-quite-right function pages. We have nothing comforting to tell frustrated users, except that PHP has so many functions and configuration options and styles that writing about and organizing them all in a way that makes sense to an outsider is very difficult. Our hope is that books such as this one can help introduce you to the main categories (for example, array functions) and the most commonly used functions — and then you'll be more prepared to use the riches of the online PHP manual.

One extremely helpful thing we recommend is reading the introductory page for each section of the online PHP manual. The manual is broken up into sections or chapters, such as "Arrays" and "Regexps"; each section has a "front page," which can be accessed via the "Function Reference" part of the online manual's table of contents. Usually (although not always) this page will have a lightning-quick introduction to the subject of the manual chapter, helpful hints, and a list of the functions in that section. By reading through these introductions, you will have a much better idea of the various things that PHP can do, and greatly increase your chances of being able to put your finger on just the function you need at any given moment.

The English-language version of the online PHP manual has a cool feature that can save a lot of time for certain users, particularly those experienced in another programming language. If you type the name of a function as a top-level URI, the official PHP site automatically forwards you to either the page devoted to that function (if it exists) or a search page with a completed search for that term (if it doesn't exist or doesn't have a discrete page). So, for instance, if you type `http://www.php.net/popen`, you'll be redirected to `www.php.net/manual/en/function.popen.php`. This is the same functionality as typing a search term into the search box in the page header. If you're looking for a PHP function that is similar to one in C or Perl, this trick can prove a great timesaver over navigating the manual by hand. It's also useful in cases where you know the name of the function but just want to check on the type of variable returned or the order of arguments to pass in.

You may want to keep a couple of points in mind when using this manual:

- The canonical manual text is written in a super-terse programmer's style, and it is organized in a not particularly discursive, notebook-like format.

The online manual is not the place to ask questions! It's intended for meaningful comments and observations only. Send e-mail to one of the previous commentators who provide their addresses — many of them will be happy to help you. Or subscribe to the mailing list or post to a PHP forum, which is faster anyway. Remember, a stupid question posted to the manual errata will go down on your (semi)permanent record.

■ The comments are edited, weeded, and verified only on an episodic (not to say extremely infrequent) basis. Proceed with extreme caution — there have been numerous instances of problems actually getting worse because a user uncritically followed the advice in the manual notes. You can write to the person to make sure that the advice is appropriate for you or even to determine whether the person really knows what he or she is talking about.

■ The manual may lag behind development by a considerable time period.

The PHP Mailing Lists

The "official" PHP community meets and greets on the PHP mailing lists. With the advent of PHP4, a decision was made to split up the lists into more specific topics. These topics will continue to proliferate with the addition of new features.

Signing up

To subscribe to any of the PHP mailing lists, go to `www.php.net/mailing-lists.php`.

You should see a large form listing the various mailing lists and options for viewing them. Just choose the list that you want, specify normal or digest versions, enter your e-mail address, and click the Subscribe button. You can also unsubscribe from a list here.

The PHP mailing list manager almost instantaneously sends you an e-mail message asking you to confirm your subscription. You aren't subscribed until you reply to this e-mail.

TIP **You also find links to local (non-English) mailing lists and newsgroups at the bottom of this page. If you want to discuss PHP in Turkish or Japanese, this page is the place to start!**

Users' lists and developers' lists

Many of the user-oriented lists are new with PHP4 or later. The following are the most popular PHP users' mailing lists:

■ **php.general**: Main mailing list — very heavy traffic, 80+ e-mail messages per day

■ **php.windows**: Specific mailing list for Windows users

■ **php.install**: An installation-related mailing list, mostly for new users

■ **php.db**: The database-related issues mailing list

- **php.i18n**: Internationalization and localization mailing list
- **php.pear**: The PEAR users' list
- **php.gtk.general**: The PHP-GTK users' list
- **php.smarty.general**: The Smarty templates users' list
- **php.bugs**: Bug reporting related to PHP itself
- **php.announce**: Announces new releases — very occasional

Lists also are available for popular PHP-based projects such as Midgard and phpNuke; you subscribe to those lists through the products' own web sites. We list the URLs of some of these sites in the "Major PHP Projects" section later in this appendix.

The following four lists are mostly intended for active developers and very early adopters — people who are going to get down in the C code and battle bugs to the death:

- **php.dev**: The main PHP developers' list
- **php.gtk.dev**: The PHP-GTK developers' list
- **php.pear.dev**: The PEAR developers' list

These lists are low-to-medium volume, meaning approximately 100 to 1,000 messages a month. They are highly technical and mostly not enlightening unless you're an active team member. Various CVS lists for developers are also available, which mail out all CVS commits on a particular branch to all subscribers; special lists for documentation writers and QA team members also exist.

If you're comfortable with Internet newsgroups (which many newer users are not), you can access the PHP mailing lists through the news gateway at `news.php.net`.

This option has one great advantage: You can send messages to the mailing lists without subscribing to them. Many new users, however, should think in terms of searching the archives for answers to old questions before (or rather than) asking new questions anyway.

Most of these mailing lists, and many others on a variety of topics, are archived and searchable at `http://marc.theaimsgroup.com`.

This archive dates back to at least 1998, although the older posts are usually less complete.

TIP   **Trying a quick keyword search on the PHP site, mail archives, and perhaps some of the other major PHP web sites before contacting the mailing lists is the polite thing to do. It's actually faster for you, plus the less time the developers must spend answering the same questions over and over, the more time they have to implement new features in the language. Actually, searching the archives is no longer just polite — it's a necessity. With so many new users, so-called "RTFM" (read the effing manual) questions are not (politely) answered on the PHP lists anymore. Also, try to ask your question on a specific list if it exists — especially installation-related questions.**

Regular and digest

The main PHP user list is so high-volume that it has a twice-daily digest version. The new special-ized mailing lists also typically have digest versions. The raw and digested versions each have advantages and disadvantages.

If you've never had 100+ e-mail messages a day pouring into your mailbox, you have no idea how distracting and time-consuming this experience can be. Just reading-and-deleting can take up a couple hours, whereas actually answering them can easily become a full-time job. Under no circum-stance should you request the full user list if your primary mailbox is a web-based free e-mail ser-vice such as Yahoo! mail or Hotmail.

> **TIP** Setting up a separate mailbox for PHP mail is almost mandatory if you're subscribing to the full user mailing list, unless you've set up good mail filters. Otherwise, you quickly start to lose mail from other sources in the flood of similarly named threads.

On the other hand, the digest version makes getting into the flow more difficult. The few brave com-munity members who get the full user list seem to answer all the questions on the half-volley before you even get the digest, making participation difficult for the time-stressed community member.

For beginners, we recommend the digest version. You can always trade up later, whenever you're ready to stop lurking and participate actively.

Everyone should also consider using one of the PHP forums (see the following section on "Other PHP Web Sites") instead of or in addition to the user mailing list. These forums are great for those who dislike mailing lists. The downside is that PHP developers generally don't hang out here, so extremely abstruse infrastructure questions usually go unanswered. The upside is that they tend to be friendlier, especially to repetitive newbie questions, because the answerers can control the amount of contact they prefer and go away if they start to become annoyed.

Mailing list etiquette

Open source mailing lists can be intimidating places, and the PHP general-users list is particularly active and fast-paced. The denizens of the mailing lists are people, and learning about their different personalities and plans over time can be fun — but they can get annoyed and fed up, too. A little netiquette can take the user a long way. The following sections offer a few tips to follow.

Remember, the community does all this work for free!

Before you turn on the flamethrower, remind yourself of your last experience with commercial-software tech support. Did it solve your problem the same day? Did it cost money? How long did it take? At what point did you get to talk to the developers of the program?

People might be sick of your question

Perhaps you're trying to install PHP for the first time and can't get it working as an Apache module — but we can assure you that tens of thousands of iterations of that question have appeared on the mail-ing lists over the years. People on the mailing lists are experiencing fatigue at answering questions that

they and others have answered in a lot of other information sources — the FAQs, the online manual, the mailing-list archives, and any number of other web sites. If you ask one of these basic questions on the general mailing list, it proves that you didn't take seriously the numerous polite requests on the PHP site to search for an answer to your question before posting. You may get an irritated e-mail informing you of all the above — or you may get no responses at all. Neither response means that the PHP community is cold-hearted and unhelpful. Try to see things from the point of view of community members of longer standing, and avoid these problems by searching for the answer to your question before you post.

Give detailed descriptions

Say as much as you can about your platform, the problem, and any steps you've already tried. Don't worry about being concise; you're far better off meandering on a little than making everyone go back and forth an extra time.

Code fragments are the very most efficient way to state your problem for debugging by the community. Many people edit their raw code to make it more anonymous and/or abstract. Remember to take out any passwords!

> **TIP** **Copy and paste your code fragments; don't retype them. List participants often post perfect code, only to be frustrated that nobody can find anything wrong with it — because they corrected their errors while retyping!**

Make sure that you use a specific subject line — the more specific the better. "Subject: PHP Help" gets you ignored by most of the mailing-list regulars. You want to say something more descriptive such as, "Subject: mysql_connect arguments not being passed in 4.0.0."

PHP is international

PHP is developed and used by people literally all over the world. In fact, the active development team has only a smallish minority of native English speakers on it at any given time.

Native English speakers should feel supremely lucky that theirs is the lingua franca of the Internet in general and the PHP world specifically. They should feel awed by the linguistic dexterity of all the citizens of other nations and perhaps slightly abashed that they can't return the favor in Finnish or Urdu. In other words, cut people some slack already! Don't assume that someone is an idiot because his or her messages aren't couched in perfectly grammatical and smooth English. Instead, you might spend the time learning how to write "Thank you" in all the languages of the various PHP community members — it makes a nice sig file for your mailing list posts.

> **TIP** **If you don't know English well, you may want to write your question twice — once in English, once in your native language. This will increase the odds that someone will be able to decipher your meaning.**

There are limits

The mailing list and other resources are meant to help you, but you must prepare to make a good-faith and even strenuous effort of your own. Help does not mean that someone comes to your office

and writes your code — this is not a remake of the Disney version of Cinderella, with dancing, sewing, chore-doing mice! Please don't ask community members to go into your server and debug your scripts for you.

Every once in a while, someone gets on the mailing list and whines about how PHP doesn't have precisely the feature that he or she is looking for — to which the developers very sensibly reply, "Why don't you implement it yourself?" Or, if you're not a good C programmer yourself, you could always pay someone else to develop your feature and contribute it to the PHP community. At the very least, you can avoid doing things that may alienate others or cause developers to burn out on the whole idea of developing open-source software!

Do it yourself

Open source software may be free to use, but you should not consider it free of all responsibilities. You are technically a "free rider" until you give back — or pay forward — to the community at large. It's your task to figure out where and how to best deploy your talents, and then to do that thing as you can. We don't mean that every casual PHP user must become a C developer, but you can contribute in many other ways. Answering questions on the PHP mailing lists or web sites is always a good thing, because it lightens the load on the core developers. If you figure something out that seemed obscure in the online PHP manual, be sure to post your findings to the User-contributed notes section of the manual. Use the PHP bug-tracker according to the instructions. Simple steps like these, in aggregate, contribute to the healthy community that has made PHP so successful.

It's probably you

If you experience a failure to communicate, you need to ask yourself whether the problem could possibly lie with you. If you do find yourself in the middle of a flame war, which happens occasionally on any mailing list, people enjoy nothing more than a little public acknowledgment of what a jerk you've (unknowingly) been.

There are now commercial alternatives

If the whole ethos of the PHP mailing lists is driving you crazy, remember that you can now pay to play instead. Many companies are now staffed by well-known PHP developers who are willing to do everything from answering single questions to building a custom PHP extension for you. ThinkPHP (`www.thinkphp.de`), for example, is a German consultancy, associated with PHP team member Thies Arntzen, that offers support, training, and performance evaluations. In the United States, the supremely helpful Richard Lynch of PHP-mailing-list fame takes requests at `www.l-i-e.com`.

Other PHP Web Sites

Besides the *official* PHP resources that we mention in the preceding sections, some well-known community members have put up some extraordinarily helpful web sites. Some of these enjoy a special relationship with PHP, and are "quasi-official." The following sections describe some of these sites.

Core scripting engine and tools

The core of PHP is the Zend scripting engine. It is produced by an Israeli company called Zend, and besides being a free part of PHP, it can be embedded in other applications. Zend also produces various PHP tools and add-ons, such as a graphical debugger and a precompiler. You can find information on Zend products at `www.zend.com`.

Zend.com is the home of the core PHP5 scripting engine, as well as a center of PHP commercialization. Although the company sells support and custom development services to larger companies, the vast majority of PHP developers are most interested in the add-on products being developer by core developers Zeev Suraski and Andi Gutmans and their team.

For most PHP users, the most useful product is the Zend Studio IDE, now in 3.0.1 release. (See Chapters 3 and 32 for more information and screenshots.) This program is the first PHP-specific development tool available, with many well-designed features for the PHP professional. Because of Zend's unique relationship to PHP development, the company understands the language completely and can design an editing tool that is customized to the needs of hardcore PHP users.

Two other Zend products are primarily of value to companies. PHP consulting firms should find the Zend Encoder useful, as it enables them to ship their code in a platform-independent, optimized intermediate representation. Large PHP sites can get a quick return on investment by using the Zend Accelerator, which boosts performance by optimizing and caching, thereby requiring less capital investment in hardware.

The Zend site also offers unique content on a regular basis, including biographies of major figures in the PHP world, a handy weekly newsletter summarizing current issues in the development of our beloved programming language, and great articles on advanced topics in PHP development. Zend.com is one of the few web sites that consistently offers articles of interest to corporate PHP developers and architects, often showcasing the finer points of PHP functionality, such as reference counting, output buffering, and changes to the `include` functions. Topics such as these may seem abstruse at first, but advanced users usually enjoy learning about the underlying structure and logic of the programming language so that they can write the tightest, cleanest, most secure, best-architected, and most well-thought-out PHP code possible. C programmers who want to contribute to PHP can also find inspiration and information in these articles — because Zeev and Andi are driving the course of PHP core development, getting a feeling for their aesthetic and decision making is important if you want to delve into the heart of PHP.

PHP knowledgebase

PHP has a great knowledgebase, something like a FAQ-o-matic but more full-featured, called PHP Faqts (previously known as E-gineer). It is available at `http://php.faqts.com`.

PHP Faqts is an interesting and nicely executed concept: an archived knowledgebase of answered and unanswered questions from real PHP users with a decent search function. For common questions, this site is much easier to use than the mailing lists or even many forums — and for that reason, we recommend it to new PHP users.

The way that Faqts (brainchild of Australian PHP whiz Nathan Wallace) works is that community members ask questions in one of several "buckets," such as Installation and Setup, Common Problems, Database Backed Sites, and the extremely cool Not Quite PHP. Other community members come along and add multiple answers to these questions. They can also associate other questions (basically other ways of stating the same thing) with that question/answer pair. Everyone can vote anonymously on whether the question/answer was useful or not. Thus, you get an accretion of knowledge over time and some way to discern whether a particular answer is good at a glance.

Going to this site, however, is not the fastest way to get your question answered, and it remains to be seen whether Faqts can scale indefinitely — but it's a cool idea and definitely a resource to try.

Articles and tutorials

Articles and tutorials take a "teach a man to fish . . ." approach. Often they can't really walk you through all the steps involved in building your web site; instead, they attempt to guide you in thinking about what to do. Following are two sites that we recommend for such information.

- **PHPBuilder** (`www.phpbuilder.com`). Founded by Tim Perdue, the top PHP app developer responsible for Sourceforge.net and Geocrawler.com, this site has long been one of the most comprehensive and well-run PHP sites. The specialty of the site is a deep backlog of articles that focus on the correct architecture of PHP sites, with subjects such as user authentication, cross-platform development, database abstraction, and documentation and style. PHPBuilder also boasts one of the most active PHP-related web forums, with excellent response times to most questions.

- One downside to the site is that articles are not dated, so determining whether the advice is still relevant to current versions of the programming language can prove difficult. Despite this drawback, PHPBuilder is a must-visit for the more conceptual PHP programmer who wants to read well-argued position papers on the right way to code PHP.

- **Devshed** (`www.devshed.com/Server_Side/PHP` and `www.devshed.com/Server_Side/MySQL`). A big commercial site with good tutorials and a forum, Devshed covers all the scripting languages (ASP, JavaScript, Python, and so on) as well as MySQL, making it the best one-stop for those still in the shopping phase.

PHP codebases

Codebases take a "give a man a fish . . ." approach, simply offering their donated wares to all takers. The code quality can vary widely, from first scripts to elegant classes contributed by experts in a particular area. The following sections describe a few such sites that you can visit.

CAUTION Although codebases can seem attractive to new developers and do embody the power of open source, there are reasons to be wary. For one thing, no one is guaranteeing the quality or safety of this code. If you're more comfortable cutting and pasting than writing it yourself, you may not have sufficient skill at reading other people's PHP to use contributed code in an intelligent way. Proceed with caution!

- **PHP Classes Repository** (`www.phpclasses.org`): Originally a collection of classes by Manuel Lemos, this site is now a hotbed of OOP PHP. We would probably not recommend this site to beginners, both because of the heavy use of object-oriented programming and the strong leaning toward code in the "Sure you can, but is it a good idea?" category. You should also possess an understanding of the changes in the object model from PHP4 to PHP5. If you're good at intelligently reading other people's code and adapting it to your own needs, however, this site can prove instructive.

- **PX: PHP Code Exchange** (`http://px.sklar.com`): A super-plain and uninformative site design nonetheless leads to a large variety of scripts — mostly smaller ones. Look here for a standalone snippet or function in a specialized area, such as graphics or math.

> **TIP**   **A quick rule of thumb in judging contributed code: If you can't follow along pretty well just by reading the comments, take a pass and look for another code sample. It's pretty rare for a good commenter to be a bad or malicious programmer.**

Major PHP projects

These are the more ambitious standalone projects based on PHP that are becoming well known in their own right. Even organizations that are not necessarily in love with PHP are beginning to consider these projects as the best-of-breed and/or most cost-effective option in their various categories.

- **PHPMyAdmin** (`www.phpmyadmin.net`): Originated by Tobias Ratschiller, this program is a graphical frontend to MySQL that has brought database administration to the ranks of the command-line phobic. See Chapter 14 for more detailed information on how to use it.

- **PHP-Nuke** (`http://phpnuke.org`): This site offers a newslog-style content management system that enables multiple users on an intranet or the web to post stories and comments on an ongoing basis. You can find lots of add-on packages written by enthusiastic users as well.

- **PHPSlash** (`http://phpslash.sourceforge.net`): This site also offers a newslog-style content management system. It was originally a rewrite of Slashcode (the Perl codebase behind Slashdot) in PHP, although PHPSlash development has now diverged somewhat.

- **Midgard** (`www.midgard-project.org`): This site offers a highly customizable content management system, similar to Vignette Story Server. Midgard doesn't simply enable users or editors to post short pieces in a constant format on a web page; you can also use the program to manage workflow on all kinds of content-rich sites.

- **phpBB** (`www.phpbb.com`): This is an object-oriented, template-based bulletin-board system offering threaded or flat view, skins, avatars, and other attractive display features.

- **Phorum** (`www.phorum.org`): Phorum is a lighter-weight bulletin-board system with no graphics. Unlike most other PHP bulletin boards, Phorum displays an outline of the thread, plus the current message, plus a form to reply to that message on a single page.

- **SquirrelMail** (`www.squirrelmail.org`): IMAP web mail client.

- **Serendipity** (`www.s9y.org`): Full-featured blogware, comparable to (commercial and Perl-based) Movable Type and Blogger. One of the authors of this book is a team member.

- **PHPWiki** (`http://phpwiki.sourceforge.net`): Popular Wiki system.

- **PHPGroupware** (`www.phpgroupware.org`): This large integrated suite of PHP programs offers you web-based group-scheduling and interaction tools, including web mail, a calendar, to-do lists, chat, forums, and more.

- **Sourceforge.net** (`http://sourceforge.net/projects/alexandria-dev`): This is a web-based engineering management toolkit that includes a task tracker, a bug tracker, a CVS frontend, forums, a documentation manager, and news releases. The codebase went closed source in 2001.

Appendix E

PEAR

The PHP Extension and Application Repository (PEAR) is a broad
effort with many component parts, collectively aimed at expanding
the usefulness and reliability of the PHP language. PEAR has been
removed from PHP6 but was previously (and likely will continue to be) an
important part of many developed applications for years to come. For this
reason, while we didn't include coverage of PEAR in the main section of
the book we're including it as an appendix in the hope that you'll get some
value out of the material.

As of late, there has been discussion about how to get PEAR back into the
main PHP source, so it may be back between the time we write this and
the time you read this! That is all the more reason to include some cover-
age of PEAR.

Since PEAR isn't included in PHP6, some of the installation steps may no
longer be applicable to a PHP6 system. You can find the latest installation
instructions at `http://pear.php.net`.

With PEAR, developers should be able to write more capable software more
quickly and with greater reliability.

The most useful and best-known element of PEAR, its package manage-
ment system, attempts to reduce the frequency with which PHP developers
reinvent the wheel. Its main part is an online database of code modules,
accessible to anyone via an automated process, that give the PHP language
special capabilities. PEAR modules, for example, enable PHP program-
mers to access LDAP directories and open files in the Ogg Vorbis format
without writing utility classes for those jobs. Programmers using the PEAR

packages can focus on the functionality of their creations, rather than wasting time struggling with nuts-and-bolts problems.

The PEAR initiative also includes a set of rules about how code is to be written — a style guide, if you like. The PEAR coding style rules are meant to govern modules contributed to the PEAR archive but in fact apply well to all PHP work. You could do worse than to apply the PEAR coding style rules to all your PHP programs.

PEAR has a sister project, the PHP Extension Community Library (PECL, pronounced *pickle*). PECL modules are extensions to PHP itself, rather than just PHP modules that can be imported into PHP programs as needed. Together, PEAR and PECL make PHP much more capable and enable many more people to participate in the development of the language.

What Is PEAR?

There are many common tasks in PHP that require or strongly benefit from libraries of functions. There are many web sites where PHP community members offer code they've written, but how do you know the code is good, will be maintained and extended, and doesn't have any odd quirks or even malicious features? The PEAR project offers a large and growing library of known-good, well-maintained, well-documented PHP code that has passed many quality inspections — all free for the taking.

The PEAR project began in 1999, shortly after PHP itself came into being. It's a community-driven initiative dedicated to generating open source code that improves PHP. PEAR packages are built on top of the standard PHP functions, and are often written in an object-oriented style (for example, classes). You include these modules from your own PHP script with an `include()` or `require()` statement, as you would any other PHP function library or class.

For the most part, PEAR is to PHP as the Comprehensive Perl Archive Network (CPAN) is to Perl. It has many parts, but the best known and most frequently used is a library of open source PHP code modules that may be accessed in an automated way. The PEAR module management system makes it easy for you to keep a server's PHP installation up to date and outfitted with the elements it needs to do its job (for example, with the PEAR DB classes for standardized database access and the PEAR LDAP classes for accessing a corporate directory). You can run the package manager as an automated routine that checks for updated versions of your installed packages every week, if you like.

Other parts of the PEAR project include:

- A set of coding standards that specifically applies to PHP modules distributed by PEAR
- The PHP Foundation Classes (PFC), which are a few especially worthy PEAR classes distributed with the main PHP package
- Various code archives and mailing lists for the people doing PEAR module development work

The PHP Extension Community Library (PECL) is a collection of PHP extensions (written in C as all PHP extensions are) which are relatively rarely used and therefore do not need to be part of the core PHP distribution (which was threatening to become too large and unwieldy). PECL used to be part of PEAR, but has been split off for separate management. PECL and PEAR share the same automated distribution tool, though, and so remain related projects. The key difference between PEAR modules and PECL modules: PEAR modules are written in PHP and may be included in PHP programs as required. PECL modules are written in C, and may be incorporated into the PHP engine itself by the normal process of recompiling.

The PEAR Package System

The PEAR package system is an archive of compressed files (tar files compressed with gzip), each of which contains a series of PHP files and a manifest file in XML format. Each archive, when incorporated into a PHP installation on a server (by means of the automated package-management system that's discussed later in this chapter), adds to the overall collection of functions and classes a developer can invoke in his or her code. Widely used packages handle database abstraction, the interpretation of various file formats, the implementation of industry-specific algorithms, and all kinds of convenience functions. The universe of PEAR packages is large and expanding, and because the packages are of such high quality, you should make use of them in your own code if you can.

The PEAR homepage is `http://pear.php.net`.

A sampling of PEAR packages

Here's a much-abridged list of PEAR packages. The package name generally describes its function:

- `Auth` — User authentication
- `Benchmark` — Performance calibration
- `DB` — Database connectivity
- `Calendar` — Calendar objects and functions
- `Archive_Tar` — Interaction with tar files
- `Archive_Zip` — Interaction with Zip files
- `HTTP` — Manipulation of the HTTP protocol
- `Image_Barcode` — Barcode generation
- `I18N` — Internationalization tools
- `Log` — Logging
- `Mail` — Interaction with POP, IMAP, and SMTP
- `Oggvorbis` — Interpretation of the Ogg Vorbis open-source audio file format

- `Tree` — Tree structures for organizing objects
- `SOAP` — Implementation of the SOAP protocol

Aside from enabling PHP server administrators to incrementally adjust the capabilities of their systems, the PEAR package system is a way of dividing the labor involved in expanding the capabilities of PHP. Each of the many packages in the system — there are more than 250 as of this writing — has a separate development team behind it, complete with a project lead and several other contributors. Individual packages have version numbers and (usually) their own supporting documentation. Packages may depend on other packages (meaning that the depended-upon package must be installed); managing these dependencies is one function of the PEAR package management tool.

How the PEAR database works

The PEAR database serves two purposes: it is by design accessible to human readers as well as to the PEAR package-management client. You can use an ordinary web browser to navigate around the HTML documents at the PEAR site (`http://pear.php.net`), or you can use the package management client to interface with it via a Web service interface.

Either way, the PEAR repository is organized as a tree, with related packages grouped into hierarchies (though hierarchical relationships do not necessarily indicate dependency relationships among packages). The PEAR community manages what goes into the tree, determining when development on a particular package has progressed far enough to warrant a new release into the publicly available repository.

The Package Manager

If, like most people, you're planning to use the PEAR repository as a resource rather than as an entity to which to contribute, your main interaction with it will be through the PEAR Package Manager. The package manager is a command-line program that interacts with the online repository and allows you to download, install, and uninstall PEAR packages according to your requirements. This remainder of this section shows you how to get and use the PEAR Package Manager.

> **TIP** **Remember that if you just want to use PEAR's** `DB`, `Net_Socket`, `Net_SMTP`, `Mail`, `XML_ Parser`, **or** `phpUnit` **modules, you do not need to install the PEAR Package Manager or any packages! These packages, which are collectively referred to as the PEAR Foundation Classes, are bundled with PHP.**

Installing the PEAR Package Manager on Linux

Assuming that PEAR is no longer included in PHP, or if you're running an older version of PHP under Linux (really old, older than version 4.3), you'll need to install the PEAR Package Manager by means of a two-part command. The command looks like this:

```
$ lynx -source http://pear.php.net/go-pear | php
```

That command opens up the specified URL (which you can examine yourself through an ordinary web browser) with Lynx, a text-only HTTP client (certain Linux distributions have similarly functional programs with different names, such as `links` under Red Hat Linux). The URL contains text that defines a PHP program. The command line pipes that text to the PHP engine, thus allowing it to be interpreted.

Updating the Package Manager

Later, you may want to go through the `go-pear` procedure again to update your system and make sure it's aware of the latest contents of the PEAR repository. You may want to do this every few months if you use PEAR packages very frequently or don't reinstall PHP very often. However, most people will find that getting a new version of the PEAR Package Manager every time you install a new version of PHP is frequent enough. The basic procedure is to go to `http://go-pear.org` and save the file there — there is only one — as `go-pear.php` in a directory that's accessible to your PHP compiler.

Figure E-1 shows the `go-pear` web site.

FIGURE E-1

The go-pear web site

After saving `go-pear.php`, go to the command line and run this command:

```
php go-pear.php
```

You should see output similar to the code above. With that done, you're again ready to make use of the PEAR Package Manager.

Using the Manager

The PEAR Package Manager has a command-line interface that is common to all versions of PHP. The instructions in this section apply equally to all Unix variants (including Linux) and to Microsoft Windows.

The key executable of the PEAR Package Manager is `pear`. It resides in your PHP home directory, alongside the PHP interpreter itself.

Automatic package installation

Once you have PEAR installed and updated, you can install any package you've downloaded. The generic syntax for doing an automatic installation of a package is this:

```
pear install <package>
```

In that syntax, `<package>` is the name of a PEAR package. All available packages are listed at `http://pear.php.net/packages.php`. You can also run:

```
pear remote-list
```

to see what's available.

Here's an example of installing the PEAR DB package via the automatic PEAR Package Manager method:

```
C:\PHP>pear install DB
downloading DB-1.5.0RC2.tgz ...
Starting to download DB-1.5.0RC2.tgz (68,128 bytes)
.................done: 68,128 bytes
install ok: DB 1.5.0RC2
```

Automatic package removal

Uninstalling a package is just as easy as adding one. The generic syntax looks like this:

```
pear uninstall <package>
```

To uninstall the DB package, then, we'd do this:

```
C:\PHP>pear uninstall DB
Uninstall ok: DB
```

If you're not sure what packages are installed locally, run the following command to find out:

```
pear list
```

Semiautomatic package installation

If, for some reason, you downloaded a PEAR package in the form of a `.tgz` file, you can later use the PEAR Package Manager to install it, even if there's no connection to the Internet available. You just point the `pear` command at the local file, as follows:

```
pear install HTML_BBCodeParser-1.0.tgz
```

Using PEAR packages in your scripts

Once you've installed the PEAR modules you wish to use, you should make sure the location is included in the `include_path` variable of your `php.ini` file. This location can be tricky — it will probably be `/usr/local/lib/php` on Unix servers and whatever you specified during the `go-pear` procedure on a Windows server. Once you've done that, you can include these libraries from any PHP script with a normal include directive:

```php
<?php
include_once('Mail.php')

// Your code which uses PEAR Mail functions here

?>
```

PHP Foundation Classes (PFC)

The PEAR Foundation Classes (PFC) are a subset of the PEAR module repository. The modules that are part of the PFC are written to an especially high standard of quality, have been extensively tested, and are considered very stable and reliable. The PFC are distributed with PHP itself, so you do not have to download or install them separately. As of PHP5, the members of the PFC are these packages: `DB`, `Net_Socket`, `Net_SMTP`, `Mail`, `XML_Parser`, and `phpUnit`.

In writing modules for the PFC, programmers must aim for broad compatibility. They should avoid using any resource that's particular to a specific operating system, and try to take input and give output in the most generic possible form (for example, in plain text rather than as SOAP-formatted messages). Programmers also need to keep in mind possible future developments in PHP itself — information that can be gleaned from mailing lists and other community resources — and write their software so it is unlikely to break when new releases appear.

PHP Extension Code Library (PECL)

The PHP Extension Community Library (PECL) is conceptually very similar to PEAR, and in fact they share the PEAR Package Manager infrastructure (that is, PECL modules can be accessed and installed via the PEAR Package Manager). The main difference is that PECL is concerned with extensions to PHP itself, in the form of C modules that attach to the PHP engine. As C programs, extensions typically execute faster and more efficiently than the modules contained in the PEAR repository.

PECL used to be called the PEAR Extension Code Library and was spun off from PEAR in October 2003. The new PECL homepage is `http://pecl.php.net`.

The PEAR Coding Style

Newspapers (as well as publishers of books!) spend a lot of time and effort establishing style rules that govern how their writers use language. Are people identified by their last names (as in *The Washington Post*) or by their honorifics and last names (as in *The Economist* and *The New York Times*)? It's a matter of style.

The same sorts of questions arise among programmers, except that the issues at stake are usually matters of formatting rather than syntax. Where do brackets go, and how is code laid out on a page? It's important to have standard (if arbitrary) answers to these questions, because a standard style can be a real aid to error-spotting and maintainability.

PEAR defines its style rules online at `http://pear.php.net/manual/en/standards.php`. This section calls attention to some of the most important ones.

Indenting, whitespace, and line length

Code is much easier to read if you use indentation to indicate the relationship among lines of code that are tied together in a common functional block, as well as whitespace to logically group elements. The following code is hard to read, though it will run perfectly fine.

```
switch ($flag) {
case 1:
doWork();
break;
case 2:
doOtherWork();
break;
default:
doNothing();
break;
}
```

On the other hand, this code:

```
switch ($flag) {

case 1:
    doWork();
    break;

case 2:
    doOtherWork();
    break;

default:
    doNothing();
    break;
}
```

is both functional and more easily understood. Spotting syntax errors is hard enough; don't make the job harder by clumping your code together sloppily.

One of the big religious arguments in programming is the number of spaces to indent each new code block — some people insist that two saves space, others swear by four, and some outliers actually employ eight-space indents (the horror!). Over time and in groups, four has come to be a standard compromise position, adopted by many open source projects — including PEAR. If you want your code to be accepted into PEAR, it must use four-space indents.

Because different editors on different platforms interpret tab characters differently, it's recommended that you use groups of four space characters in all places you would, under other circumstances, use a tab character.

Formatting control structures

Control structures — such as `if`, `if/else`, `if/elseif`, and `switch` statements — can be confusing if not properly formatted. PEAR has recommended styles for all of these language constructs.

if Statements

A simple two-test `if` statement should be formatted like this:

```
if ((condition1) && (condition2)) {
    doSomething();
}
```

Note that the opening bracket appears on the same line as the conditions (so-called Kernighan and Ritchie, or K&R, braces), and that there are brackets even though there is only one line of code in the conditional block. That way, the fact that it's a block is obvious, and there's no need to remember to add them when further lines of code are added in the future. Also note that there should be a space between a conditional statement and the expression being tested.

if/else Statements

An `if/else` statement builds on the basic `if` format:

```
if ((condition1) && (condition2)) {
    doSomething();
} else {
    doSomethingElse();
}
```

The `else` appears on the same line as the closing bracket that terminates the `if` block.

if/elseif Statements

An `if/elseif` statement looks just like an `if/else` statement in terms of formatting:

```
if ((condition1) && (condition2)) {
    doSomething();
} elseif {
    doSomethingElse();
}
```

switch Statements

`Switch` statements rely on whitespace and indentation to make code blocks obvious:

```
switch ($flag) {

case 1:
    doWork();
    break;

case 2:
    doOtherWork();
    break;

default:
    doNothing();
    break;
}
```

Formatting functions and function calls

Much of PHP is concerned with defining functions, then making calls to them, and obviously code libraries like PEAR will be almost all functions. Properly formatting your functions can make it more obvious what's going on and can therefore make debugging and maintenance easier.

The PEAR style rules mandate that functions be defined with both their beginning and ending braces flush with the left margin, like this:

```
function myFunction()
```

```
    {

        // Function code goes here.

    }
```

This makes function definitions (which use braces) stand out from conditional blocks (which also use braces). Furthermore, the standards require that code within the function be indented. Everything is indented at least four spaces; some segments may be indented further:

```
function myFunction()
{
    doSomething();
    if ($is) {
        doSomethingMore();
    }
}
```

If your function takes arguments, be sure to order them so that arguments with default values go at the end of the list, like this:

```
function myFunction($a, $b, $c='Default')
{
    doSomething();
    if ($is) {
        doSomethingMore();
    }
}
```

Also note that there should be no spaces between the name of the function and the parentheses containing arguments. Again, this helps visually distinguish functions (which use parentheses) from expressions (which also use parentheses).

It is important that functions return something. The return value will either be a value that resulted from the function's processing or a Boolean value (true or false) to indicate success or failure.

Summary

In this appendix, you learned about PEAR, a previously included framework for PHP community-based projects that extend the capabilities of the language. PEAR exists to facilitate the ongoing development and widespread distribution of handy toolkits.

At the center of PEAR is its repository, an online database that contains the accumulated body of PEAR packages. This repository has an HTML interface as well as an XML_RPC (Web services) interface, which means that you can browse it manually or interact with it via a specialized command-line program: The PEAR Package Manager. The PEAR Package Manager allows you to

quickly see what's in the PEAR repository, download what you want, and install some or all of what you download. Particularly important PEAR packages are part of the PHP Foundation Classes (PFC).

Another element of the PEAR community is a definition of a coding standard, which specifies how functions should be defined, comments placed, and brackets structured in various parts of PHP programs. It's meant to ease readability and make life easier for documentation writers.

PEAR shares its automated package-distribution scheme with PECL, which manages PHP extensions written in the C language.

PEAR represents an invaluable resource to PHP programmers of all levels. Make sure that the PEAR Package Manager is installed on your PHP server, and make full use of its resources. When you're ready, join the development effort and contribute to the growth of PHP.

Index

A

abs function, 155, 159
absolute links from relative links, 384–385
absolute value, 155, 159
abstract classes
    declaring, 328
    designated, 352
abstract data structures, 363
abstraction, procedural, 60
access
    members, 319, 325–327
    source code, 474–475
access.conf file, 489
access.log file, 514–515
accessor functions, 354
Acos function, 462
ACTION attribute, 100–101
Action directive, 475, 490
actual parameters, 87–89
add function, 329
add_1 function, 451
add_new_country function, 250
add_point_to_path function, 705
addHiddenVariable function, 347
addInputButton function, 347
addInputForm function, 347, 352
addition, 154
    arbitrary-precision function, 466
    assignment operator, 156
    increment operator, 155–156
additionalAttributes function, 351
AddModule setting, 491
addslashes function, 127–128, 300
AddType directive, 475, 490
affectedRows function, 573
aggregating query results, 289
algorithm efficiency, 534
alignment in format strings, 129
allowed elements in XML documents, 652
alpha value for transparency, 699
ALTER privilege, 209
ALTER TABLE statement
    indexes, 286–287

primary keys, 284–285
    structure, 204
alternate control syntaxes, 79–80
American Standard Code for Information Interchange (ASCII) encoding, 662
ampersands (&)
    GET method, 101
    logical operator, 61
    passing arrays by reference, 145
anchor tags for links, 381–382
anchoring characters in regular expressions, 379
anchortext for links, 381–382
and operator, 37, 61
angle_given_sides function, 774–775
answer handling in Certainty Quiz, 753
answer_string function, 543
anti-style sheets, 545
Apache modules, 485
Apache Server, 7
    building, 23
    in CentOS installations, 21
    configuration files, 489–491
    extensions setup for, 24
    log files, 514–515
    RPM versions, 20
    on Windows, 25–26
APC Accelerator, 495
appendChild function, 659
application platform considerations, 4
    cost, 4
    cross-platform compatibility, 7
    ease of use, 5
    extensions, 8
    feature development, 8
    HTML-embeddedness, 5–7
    proprietary standards, 8–9
    stability, 7–8
    user communities, 9–10
apt-get command, 190
arbitrary-precision math functions, 465–466
    configuration option, 488
    converting code to, 467–469
    example, 466–467
area of intersecting circles, 773–776

C

H

J

855

869