


[\[Team LiB \]](#)



- [Table of Contents](#)
- [Index](#)
- [Reviews](#)
- [Reader Reviews](#)
- [Errata](#)

.NET Framework Essentials, 3rd Edition

By [Hoang Lam](#), [Thuan L. Thai](#)

Publisher: O'Reilly

Pub Date: August 2003


ISBN: 0-596-00505-9

Pages: 384

Fully updated for version 1.1 of the .NET Framework, *.NET Framework Essentials*, 3rd Edition is an objective, concise, no-nonsense overview of the Microsoft .NET Framework for developing web applications and services. Written for intermediate to advanced VB, C/C++, Java, and Delphi developers, *.NET Framework Essentials*, 3rd Edition is also useful to system architects and leaders who are assessing tools for future projects.

[\[Team LiB \]](#)

[[Team LiB](#)]



- [Table of Contents](#)
- [Index](#)
- [Reviews](#)
- [Reader Reviews](#)
- [Errata](#)

.NET Framework Essentials, 3rd Edition

By [Hoang Lam](#), [Thuan L. Thai](#)

Publisher: O'Reilly

Pub Date: August 2003

ISBN: 0-596-00505-9

Pages: 384

[Copyright](#)

[Preface](#)

[Audience](#)

[About This Book](#)

[Assumptions This Book Makes](#)

[Conventions Used in This Book](#)

[How to Contact Us](#)

[Acknowledgments](#)

[Chapter 1. .NET Overview](#)

[Section 1.1. Behind Microsoft .NET](#)

[Section 1.2. The .NET Platform](#)

[Section 1.3. .NET Framework Design Goals](#)

[Section 1.4. .NET Framework](#)

[Chapter 2. The Common Language Runtime](#)

[Section 2.1. CLR Environment](#)

[Section 2.2. CLR Executables](#)

[Section 2.3. Metadata](#)

[Section 2.4. Assemblies and Manifests](#)

[Section 2.5. Intermediate Language \(IL\)](#)

[Section 2.6. The CTS and CLS](#)

[Section 2.7. CLR Execution](#)

[Section 2.8. Summary](#)

[Chapter 3. .NET Programming](#)

[Section 3.1. Common Programming Model](#)

[Section 3.2. Core Features and Languages](#)

[Section 3.3. Language Integration](#)

[Section 3.4. Summary](#)

[Chapter 4. Working with .NET Components](#)

[Section 4.1. Deployment Options](#)

[Section 4.2. Distributed Components](#)

[Section 4.3. COM+ Services in .NET](#)

[Section 4.4. Message Queuing](#)

[Section 4.5. Summary](#)

[Chapter 5. Data and XML](#)

[Section 5.1. ADO.NET Architecture](#)

[Section 5.2. ADO.NET Benefits](#)

[Section 5.3. Content Components](#)

[Section 5.4. .NET Framework Data Providers](#)

[Section 5.5. XML in the .NET Framework](#)

[Section 5.6. Summary](#)

[Chapter 6. Web Services](#)

[Section 6.1. Web Services in Practice](#)

[Section 6.2. Web Services Framework](#)

[Section 6.3. Web Services Provider](#)

[Section 6.4. Web Services Consumers](#)

[Section 6.5. Async Web Services](#)

[Section 6.6. SOAP Header in Web Services](#)

[Section 6.7. Web Services and Security](#)

[Section 6.8. Summary](#)

[Chapter 7. ASP.NET](#)

[Section 7.1. ASP](#)

[Section 7.2. ASP.NET](#)

[Section 7.3. The System.Web.UI Namespace](#)

[Section 7.4. Web Form Syntax](#)

[Section 7.5. ASP.NET Application Development](#)

[Section 7.6. ASP.NET and Web Services](#)

[Section 7.7. Data Binding and the Use of Templates](#)

[Section 7.8. State Management and Scalability](#)

[Section 7.9. Forms Authentication in ASP.NET](#)

[Section 7.10. Summary](#)

[Chapter 8. Windows Forms](#)

[Section 8.1. Introducing Windows Forms](#)

[Section 8.2. The System.Windows.Forms Namespace](#)

[Section 8.3. Windows Forms Development](#)

[Section 8.4. Windows Forms and Web Services](#)

[Section 8.5. Summary](#)

[Chapter 9. .NET and Mobile Devices](#)

[Section 9.1. ASP.NET Mobile Controls](#)

[Section 9.2. .NET Compact Framework](#)

[Section 9.3. Summary](#)

[Appendix A. .NET Languages](#)

[Section A.1. Microsoft-Supported Languages for .NET](#)

[Section A.2. Third-Party Languages for .NET](#)

[Appendix B. Common Acronyms](#)

[Appendix C. Common Data Types](#)

[Section C.1. Usage](#)

[Appendix D. Common Utilities](#)

[Section D.1. Assembly Generation Utility \(al.exe\)](#)

[Section D.2. Assembly Registration Utility \(gacutil.exe\)](#)

[Section D.3. MSIL Assembler \(ilasm.exe\)](#)

[Section D.4. MSIL Disassembler \(ildasm.exe\)](#)

[Section D.5. C++ Compiler \(cl.exe\)](#)

[Section D.6. C# Compiler \(csc.exe\)](#)

[Section D.7. Visual Basic Compiler \(vbc.exe\)](#)

[Section D.8. PE File Format Viewer \(dumpbin.exe\)](#)

[Section D.9. Type Library Exporter \(tlbexp.exe\)](#)

[Section D.10. Type Library Importer \(tlbimp.exe\)](#)

[Section D.11. XML Schema Definition Tool \(xsd.exe\)](#)

[Section D.12. Strong Name Utility \(sn.exe\)](#)

[Section D.13. Web Service Utility \(wsdl.exe\)](#)

[Colophon](#)

[Index](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Copyright

Copyright © 2003, 2002, 2001 O'Reilly & Associates, Inc.

Printed in the United States of America.

Published by O'Reilly & Associates, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly & Associates books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (<http://safari.oreilly.com>). For more information, contact our corporate/institutional sales department: (800) 998-9938 or corporate@oreilly.com.

Nutshell Handbook, the Nutshell Handbook logo, and the O'Reilly logo are registered trademarks of O'Reilly & Associates, Inc. Microsoft, MSDN, the .NET logo, Visual Basic, Visual C++, Visual Studio, and Windows are registered trademarks of Microsoft Corporation. Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and O'Reilly & Associates, Inc. was aware of a trademark claim, the designations have been printed in caps or initial caps. The association between the image of shrimp and the topic of the .NET Framework is a trademark of O'Reilly & Associates, Inc.

While every precaution has been taken in the preparation of this book, the publisher and authors assume no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Preface

A condensed introduction to the Microsoft .NET Framework, this book aims to help programmers make the transition from traditional Windows programming to the world of .NET programming. The Microsoft .NET Framework includes the Common Language Runtime (CLR) and a set of base classes that radically simplify the development of large-scale applications and services. This book examines the CLR in detail, so that you can put its new features to good use. The book also illustrates how language integration really works and guides you through component and enterprise development using the .NET Framework. In addition, it introduces you to four key .NET technologies: data (ADO.NET) and XML, web services, Web Forms (ASP.NET), and Windows Forms.

We used the latest release of Microsoft Visual Studio .NET 2003 and the .NET Framework SDK 1.1 to prepare this manuscript and to develop all the examples and figures in this book. In .NET Framework SDK 1.1, Microsoft has fixed a number of bugs, made performance and security improvements, and added a few important changes, including:

- Support for additional ADO.NET data sources (Oracle and ODBC)
- Improved support for using C++ to build .NET Windows and Web applications
- Integration of J# into VS.NET
- Inclusion of a mobile controls package for developing ASP.NET sites for small devices, such as cell phones
- Integration of mobile application development into VS.NET

Although we have done our best to ensure that the technical content of this book is up-to-date, it is possible that some items have changed slightly from the time of writing. To stay up-to-date, regularly check <http://msdn.microsoft.com/net>, <http://www.gotdotnet.com>, and this book's O'Reilly page, <http://www.oreilly.com/catalog/dotnetfrmess3/>.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Audience

Although this book is for any person interested in learning about the Microsoft .NET Framework, it targets seasoned developers with experience in building Windows applications with Visual Studio 6 and the Visual Basic and Visual C++ languages. Java™ and C/C++ developers will also be well prepared for the material presented here. To gain the most from this book, you should have experience in object-oriented, component, enterprise, and web application development. COM programming experience is a plus.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

About This Book

Based on a short course that Thuan has delivered to numerous companies since August 2000, this book is designed so that each chapter builds on knowledge from the previous one for those unfamiliar with each technology. To give you a heads-up, here are brief summaries for the chapters and appendixes covered in this book.

[Chapter 1](#) takes a brief look at Microsoft .NET and the Microsoft .NET Platform. It then describes the .NET Framework design goals and introduces you to the components of the .NET Framework.

[Chapter 2](#) lifts the hood and peers into the CLR. This chapter surveys the rich runtime of the CLR, as well as other features.

[Chapter 3](#) introduces you to .NET programming. You'll examine a simple program that uses object-oriented and component-based concepts in five different languages: Managed C++, VB.NET, C#, J#, and IL. You'll also experience the benefits of language integration.

[Chapter 4](#) demonstrates the simplicity of component and enterprise development in .NET. Besides seeing the component-deployment features of the .NET Framework, you'll also find complete programs that take advantage of transaction, object pooling, role-base security, and message queuing—all in one chapter.

[Chapter 5](#) describes the architecture of ADO.NET and its benefits. Besides being disconnected to promote scalability, the ADO.NET dataset is also tightly integrated with XML to enhance interoperability. This chapter introduces you to the .NET data-access objects, as well as the XML namespace.

[Chapter 6](#) describes how .NET software components can be accessed through the Internet. In this chapter, we discuss the protocols that support web services, as well as how to publish and discover them. You will see how XML, used in conjunction with HTTP, breaks the proprietary nature of typical component-oriented software development and enables greater interoperability.

[Chapter 7](#) introduces you to ASP.NET, which now supports object-oriented and event-driven programming, as opposed to conventional ASP development. In this chapter, Web Forms and server controls take the center stage. In addition, we examine how to build custom server controls, perform data binding to various .NET controls, and survey state management features in ASP.NET.

[Chapter 8](#) takes conventional form-based programming a step into the future with the classes in the System.Windows.Forms namespace. Similar to Win32-based applications, Windows Forms are best used to build so-called rich or "fat" clients; however, with the new zero-effort installation procedure of .NET and the advent of web services, Windows Forms are appropriate for a host of applications.

[Chapter 9](#) shows you how to build an ASP.NET application that targets smart devices and a mobile device application that is backed by SQL Server CE.

[Appendix A](#) contains a list of links to web sites with information regarding languages that target the CLR, including some burgeoning open-source projects.

[Appendix B](#) contains a list of commonly used acronyms found in .NET literature and presentations.

[Appendix C](#) contains several lists of commonly used data types in .NET. This appendix also illustrates the use of several .NET collection classes.

[Appendix D](#) surveys the important tools that the .NET SDK provides to ease the tasks of .NET development.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Assumptions This Book Makes

This book assumes that you are a Windows or web application developer fluent in object-oriented and component-based programming. We also assume that you have some basic knowledge of XML. While COM is not a crucial prerequisite, if you have COM programming experience, you will appreciate this book and the .NET Framework all the more.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Conventions Used in This Book

We use the following font conventions in this book.

Italic is used for:

- Pathnames, filenames, and program names
- Internet addresses, such as domain names and URLs
- New terms where they are defined

Constant width is used for:

- Command lines and options that should be typed verbatim
- Direct quotes and specific method names from code examples, as well as specific values for attributes and settings within code
- XML element tags


Constant width bold is used for:

- User input in code that should be typed verbatim
- Items in code to which we'd like to draw the reader's attention

Constant width italic is used for:

- Replaceable items in code, which should be replaced with the appropriate terms

In code syntax examples, we sometimes use `[value]+` to represent one or more instances of a value and `[value]*` to mean zero or more instances of a value.

 This icon designates a tip, which contains important information about the nearby text.

 This icon designates a warning related to the nearby text.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

How to Contact Us

We have tested and verified the information in this book to the best of our abilities, but you may find that features have changed (or even that we have made mistakes!). Please let us know about any errors you find, as well as your suggestions for future editions, by writing to:

O'Reilly & Associates, Inc.
1005 Gravenstein Highway North
Sebastopol, CA 95472
(800) 998-9938 (in the United States or Canada)
(707) 829-0515 (international/local)
(707) 829-0104 (FAX)

You can also send us messages electronically. To be put on the mailing list or request a catalog, send email to:

info@oreilly.com

To ask technical questions or comment on the book, send email to:

bookquestions@oreilly.com

We have a web site for the book, where we list examples, errata, and any plans for future editions. You can access this page at:

<http://www.oreilly.com/catalog/dotnetfrmess3/>

For more information about this book and others, see the O'Reilly web site:

<http://www.oreilly.com>

For more information on .NET in general, visit the O'Reilly .NET Center at <http://dotnet.oreilly.com> and the .NET DevCenter at <http://www.oreillynet.com/dotnet/>.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Acknowledgments

The folks at O'Reilly never cease to amaze us with the support that they provide. We'd like to thank O'Reilly executive editor John Osborn for extending us the contract to write this book and for his continuous support throughout the project. Thanks to O'Reilly editor Nancy Kotary for making the first and second editions of this book such a success. We'd also like to thank O'Reilly editor Valerie Quercia who did a fine job reviewing our materials and coordinating the project for this third edition.

We'd like to give many thanks to the production and design folks at O'Reilly for making this book a reality: Emma Colby, David Futato, Matt Hutchinson, Robert Romano, and Ellie Volckhausen.

We'd like to thank Brian Jepson, who has contributed significantly to this book since the beginning of the project. Brian did an unquestionably outstanding job reading, testing, and ensuring that the technical content in every chapter lined up with the latest release. He also gave us invaluable guidance and support throughout the project. We'd like to thank Matthew MacDonald for the same reason, since he's been providing invaluable feedback and recommendations in the third edition. We'd also like to thank Dennis Angeline and Brad Merrill at Microsoft for answering technical questions on the CLR and languages.

Thuan Thai

I would like to thank my parents, siblings, and Ut Nga for their constant support during this project.

Hoang Q. Lam

I would like to thank my parents and family for their support and understanding of my being missing in-action for several months. Mom and Dad, your ongoing efforts to put your children where they are today can never be repaid.

I would like to thank my wife, Van Du, and daughter, Nina. Your smiles mean the world to me. Don't underestimate your contribution to this book.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Chapter 1. .NET Overview

The .NET Framework is a development framework that provides a new programming interface to Windows services and APIs, and integrates a number of technologies that emerged from Microsoft during the late 1990s. Microsoft announced the .NET initiative in July 2000. In April 2003, Version 1.1 of the integral .NET Framework was released. This book describes this updated version of the .NET Framework.

The .NET platform consists of four separate product groups:

Development tools and libraries

A set of languages, including C#, J#, and VB.NET; a set of development tools, including Visual Studio .NET; a comprehensive class library for building web services and web and Windows applications; as well as the Common Language Runtime (CLR). These components collectively form the largest part of the .NET Framework.

Web services

An offering of commercial web services, specifically the .NET Services initiative; for a fee, developers can use these services in building applications that require them.

Specialized servers

A set of .NET-enabled enterprise servers, including SQL Server, Exchange Server, BizTalk Server, and so on. These provide specialized functionality for relational data storage, email, and B2B commerce. Future versions of these products will increasingly support the .NET Framework.

Devices

New .NET-enabled, non-PC devices, from cell phones to game boxes.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

1.1 Behind Microsoft .NET

While the main strategy of .NET is to enable software as a service, .NET is much more than that. In addition to embracing the Web, Microsoft .NET acknowledges and responds to the following trends within the software industry:

Distributed computing

Simplifies the development of robust client/server and multi-tier (n-tier) applications. Traditional distributed technologies require high vendor-affinity and are unable to interoperate with the Web. Microsoft .NET provides remoting and web services architectures that exploit open Internet standards, including the Hypertext Transfer Protocol (HTTP), Extensible Markup Language (XML), and Simple Object Access Protocol (SOAP) and WSOL.

Componentization

Simplifies the integration of software components developed by different vendors and supports development of distributed applications. The Component Object Model (COM) has brought reality to software plug-and-play, but COM component development and deployment are too complex. Microsoft .NET provides a simpler way to build and deploy components.

Enterprise services

Allow the development of scalable enterprise applications without writing code to manage transactions, security, or pooling. Microsoft .NET continues to support COM and component services, since these services greatly reduce the development time and effort required to build large-scale applications.

Web paradigm shifts

Over the past decade, web application development has shifted from connectivity (TCP/IP), to presentation (HTML), to programmability (XML and SOAP). A key goal of Microsoft .NET is to enable the sharing of functionality across the Web among different platforms, devices, and programming languages.

Maturity of IT industry

Lessons that the software industry has learned from developing large-scale enterprise and web applications. A commercial web application must support interoperability, scalability, availability, security, and manageability. Microsoft .NET facilitates all these goals.

Although these are the sources of many ideas embodied by Microsoft .NET, what's most notable about the platform is its use of open Internet standards (HTTP, XML, and SOAP) at its core to transmit information from one machine to another across the Internet. In fact, .NET provides bidirectional mapping between XML and objects. For example, a class can be expressed as an XML Schema Definition (XSD); an object can be converted to and from an XML buffer; a method can be specified using an XML format called Web Services Description Language (WSDL); and a method call can be expressed using an XML format called SOAP.

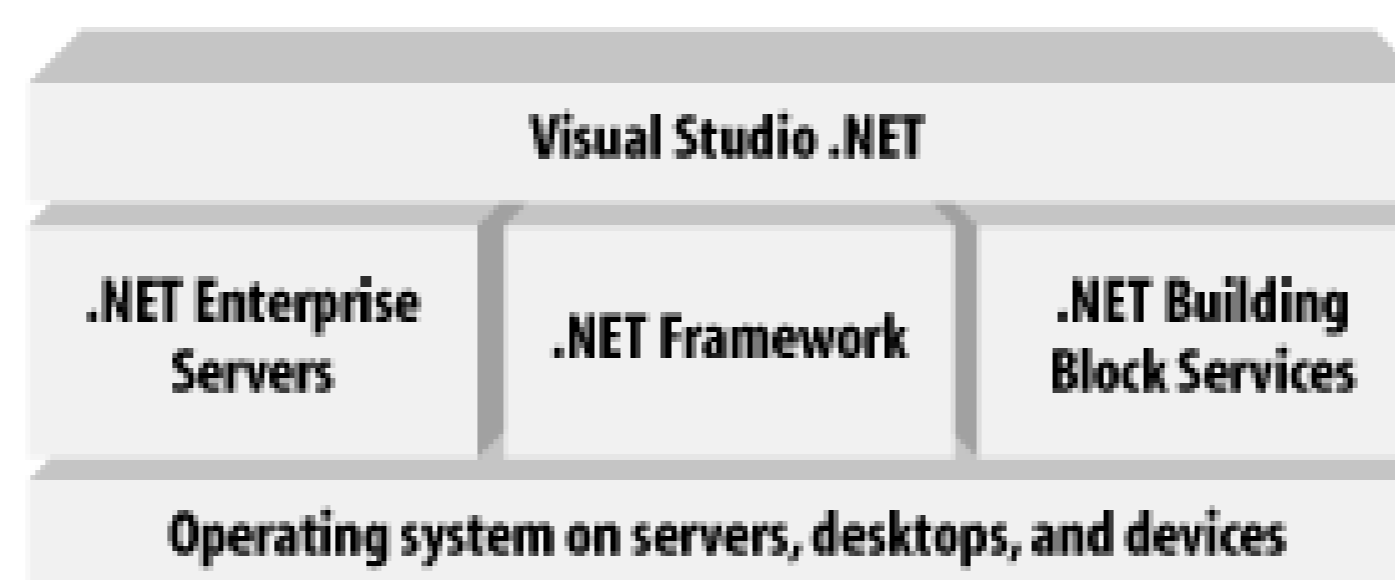
[\[Team LiB \]](#)

[\[Team LiB \]](#)

1.2 The .NET Platform

The Microsoft .NET platform consists of five main components, as shown in [Figure 1-1](#). At the lowest layer lies the operating system (OS), which can be one of a variety of Windows platforms, including Windows XP, Windows 2000, Windows Server 2003, Windows ME, and Windows CE. As part of the .NET strategy, Microsoft has promised to deliver more .NET device software to facilitate a new generation of smart devices.

Figure 1-1. The Microsoft .NET platform



On top of the operating system is a collection of specialized server products that shortens the time required to develop large-scale business systems. These server products include Application Center, BizTalk Server, Commerce Server, Exchange Server, Host Integration Server, Internet Security and Acceleration Server, and SQL Server.

Since web services are highly reusable across the Web, Microsoft provides a number of building-block services (officially called .NET Services) that applications developers can use, for a fee. Two examples of .NET Services that Microsoft offers include .NET Passport and .NET Alerts. .NET Passport allows you to use a single username and password at all web sites that support Passport authentication. .NET Alerts allow .NET Alert providers, such as a business, to alert their consumers with important or up-to-the-minute information. Microsoft plans to add newer services, such as calendar, directory, and search services. Third-party vendors are also creating new web services of their own.

The top layer of the .NET architecture is a development tool called Visual Studio .NET (VS.NET), which makes possible the rapid development of web services and other applications. A successor to Microsoft Visual Studio 6.0, VS.NET is an Integrated Development Environment (IDE) that supports four different languages and features such as cross-language debugging and the XML Schema Editor.

And at the center of .NET is the Microsoft .NET Framework—the main focus of this book. The .NET Framework is a development and runtime infrastructure that changes the development of business applications on the Windows platform. The .NET Framework includes the CLR and a common framework of classes that can be used by all .NET languages.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

1.3 .NET Framework Design Goals

The Microsoft .NET Framework embodies design goals that are both practical and ambitious. In this section, we discuss the main design goals of the Microsoft .NET Framework, including support for components, language integration, application interoperation across the Web, simplified development and deployment, improved reliability, and greater security.

1.3.1 Component Infrastructure

Prior to the introduction of COM technology, developers had no standard way to integrate binary libraries without referring to or altering their source code. With the advent of COM, programmers were able to integrate binary components into their applications, similar to the way we can plug-and-play hardware components into our desktop PCs. Although COM was great, the grungy details of COM gave developers and administrators many headaches.

Although COM permits you to integrate binary components developed using any language, it does require you to obey the COM identity, lifetime, and binary layout rules. You must also write the plumbing code that is required to create a COM component, such as `DllGetClassObject`, `CoRegisterClassObject`, and others.

Realizing that these requirements result in frequent rewrites of similar code, .NET sets out to remove them. In the .NET world, all classes are ready to be reused at the binary level. You don't have to write extra plumbing code to support componentization in the .NET Framework. You simply write a .NET class, which then becomes a part of an assembly (to be discussed in [Chapter 2](#)) that inherently supports plug-and-play.

In addition to providing a framework to make development easier, .NET removes the pain of developing COM components. Specifically, .NET removes the use of the registry for component registration and eliminates the requirements for extraneous plumbing code found in all COM components, including code to support `IUnknown`, class factories, component lifetime, registration, dynamic binding, and others.

"Component" is a nasty word because one person may use it to refer to an object and another may use it to refer to a compiled binary module. To be consistent, this book uses the term "COM component" (or simply "component") to refer to a binary module, such as a DLL or an EXE.

1.3.2 Language Integration

COM supports *language independence*, which means that you can develop a COM component in any language you want. As long as your component meets all the rules spelled out in the COM specification, it can be instantiated and used by your applications. Although this supports binary reuse, it doesn't support *language integration*. In other words, you can't reuse the code in the COM components written by someone else; you can't extend a class hosted in the COM component; you

can't catch exceptions thrown by code in the COM component; and so forth.

Microsoft .NET supports not only language independence, but also language integration. This means that you can inherit from classes, catch exceptions, and take advantage of polymorphism across different languages. The .NET Framework makes this possible with a specification called the Common Type System (CTS), which all .NET components must support. For example, everything in .NET is an object of a specific class that derives from the root class called `System.Object`. The CTS supports the general concepts of classes, interfaces, delegates (which support callbacks), reference types, and value types. The .NET base classes provide most of the base system types, such as those that support integer, string, and file manipulation. Because every language compiler must meet a minimum set of rules stipulated by the Common Language Specification (CLS) and generate code to conform to the CTS, different .NET languages can be used in the same application. We will examine the CTS and CLS in [Chapter 2](#).

1.3.3 Internet Interoperation

COM supports distributed computing through its Distributed COM (DCOM) wire protocol. A problem with DCOM is that it embeds the host TCP/IP address inside the Network Data Representation (NDR) buffer, such that it will not work through firewalls and Network Address Translation (NAT) software. In addition, the DCOM dynamic activation, protocol negotiation, and garbage collection facilities are proprietary, complex, and expensive. The solution is an open, simple, and lightweight protocol for distributed computing. The .NET Framework uses the industry-supported SOAP protocol, which is based on the widely accepted XML standards.

1.3.4 Simplified Development

If you have developed software for the Windows platforms since their inception, you have seen everything from the Windows APIs to the Microsoft Foundation Classes (MFC), the Active Template Library (ATL), the system COM interfaces, and countless other environments, such as Visual Interdev, Visual Basic, JScript, and other scripting languages. Each time you set out to develop something in a different compiler, you had to learn a new API or a class library, because there is no consistency or commonality among these different libraries or interfaces.

.NET solves this problem by providing a set of framework classes that every language uses. Such a framework removes the need for learning a new API each time you switch languages.

1.3.5 Simplified Deployment

Imagine this scenario: your Windows application, which uses three shared Dynamic Link Libraries (DLLs), works just fine for months, but stops working one day after you've installed another software package that overwrites the first DLL, does nothing to the second DLL, and adds an additional copy of the third DLL into a different directory. If you have ever encountered such a brutal-yet entirely possible-problem, you have entered DLL Hell. And if you ask a group of seasoned developers whether they have experienced DLL Hell, they will grimace at you in disgust, not because of the question you've posed, but because they have indeed experienced the pain and suffering.

To avoid DLL Hell on Windows 2000 and subsequent Windows operating systems (at least for system

DLLs), Windows 2000 stores system DLLs in a cache. If you install an application that overwrites system DLLs, Windows 2000 will overwrite the added system DLLs with the original versions from the cache.

Microsoft .NET further diminishes DLL Hell. In the .NET environment, your executable will use the shared DLL with which it was built. This is guaranteed, because a shared DLL must be registered against something similar to the Windows 2000 cache, called the Global Assembly Cache (GAC). In addition to this requirement, a shared DLL must have a unique hash value, public key, locale, and version number. Once you've met these requirements and registered your shared DLL in the GAC, its physical filename is no longer important. In other words, if you have two versions of a DLL that are both called *MyDll.dll*, both of them can live and execute on the same system without causing DLL Hell. This is possible because the executable that uses one of these DLLs is tightly bound to a specific version of the DLL during compilation.

In addition to eradicating DLL Hell, .NET also removes the need for component-related registry settings. A COM developer will tell you that half the challenge of learning COM is understanding the COM-specific registry entries for which the developer is responsible. Microsoft .NET stores all references and dependencies of .NET assemblies within a special section called *amanifest* (see [Chapter 2](#)). In addition, assemblies can be either private or shared. Private assemblies are found using logical paths or XML-based application configuration files, and public assemblies are registered in the GAC; in both cases, the system will find your dependencies at runtime. If they are missing, you get an exception telling you exactly what happened.

Finally, .NET brings back the concept of zero-impact installation and removal. This concept is the opposite of what you have to deal with in the world of COM. To set up a COM application, you have to register all your components after you have copied them over to your machine. If you fail to perform this step correctly, nothing will work and you'll end up pulling your hair out. Likewise, to uninstall the application, you should unregister your components (to remove the registry entries) prior to deleting your files. Again, if you fail to perform this step correctly, you will leave remnants in the registry that will be forever extant.

Unlike COM, but like DOS, to set up an application in .NET, you simply **xcopy** your files from one directory on a CD to another directory on your machine, and the application will run automatically.^[1] Similarly, you can just delete the directory to uninstall the application from your machine.

[1] This is true for private assemblies, but not for shared assemblies. See [Chapter 4](#) for more details.

1.3.6 Reliability

There are many programming languages and platforms in the commercial software industry, but few of them attempt to provide both a reliable language and a robust runtime or infrastructure. The most successful language that we have seen in the commercial software industry is the Java™ language and the Java Virtual Machine™, which have brought the software-development community much satisfaction. Microsoft is positioning .NET as the next big thing.

Microsoft .NET requires type safety. Unlike C++, every class in .NET is derived from the mother of all classes, Object, which supports basic features such as returning a string representation of the object indicating whether the object is equal to another, and so on. The CLR must recognize and verify types before they can be loaded and executed. This decreases the chances for rudimentary programming errors and prevents buffer overruns, which can be a security weakness.

Traditional programming languages don't provide a common error-handling mechanism. C++ and Java support exception handling, but many others leave you in the dust, forcing to invent your own error-handling facilities. Microsoft .NET supports exceptions in the CLR, providing a consistent error-handling mechanism. Put another way: exceptions work across all .NET-compatible languages.

When you program in C++, you must deallocate all heap-based objects that you have previously allocated. If you fail to do this, the allocated resources on your system will never be reclaimed even though they are no longer needed. And if this is a server application, it won't be robust because the accumulation of unused resources in memory will eventually bring down the system. Similar to Java, the .NET runtime tracks and garbage-collects all allocated objects that are no longer needed.

1.3.7 Security

When developing applications in the old days of DOS, Microsoft developers cared little about security because their applications ran on a single desktop with a single thread of execution. As soon as developers started developing client and server applications, things got a bit complicated: multiple users might then have accessed the servers, and sensitive data might be exchanged between the client and the server. The problem became even more complex in the web environment, since you could unknowingly download and execute malicious applets on your machine.

To mitigate these problems, .NET provides a number of security features. Windows NT and Windows 2000 protect resources using access-control lists and security identities, but don't provide a security infrastructure to verify access to parts of an executable's code. Unlike traditional security support in which only access to the executable is protected, .NET goes further to protect access to specific parts of the executable code—this is known as *code access security*. For example, to take advantage of declarative security checks, you can prefix your method implementations with security attributes without having to write any code. To take advantage of imperative security checks, you write the code in your method to explicitly cause a security check. .NET provides other security features to make it harder to penetrate your applications and system.

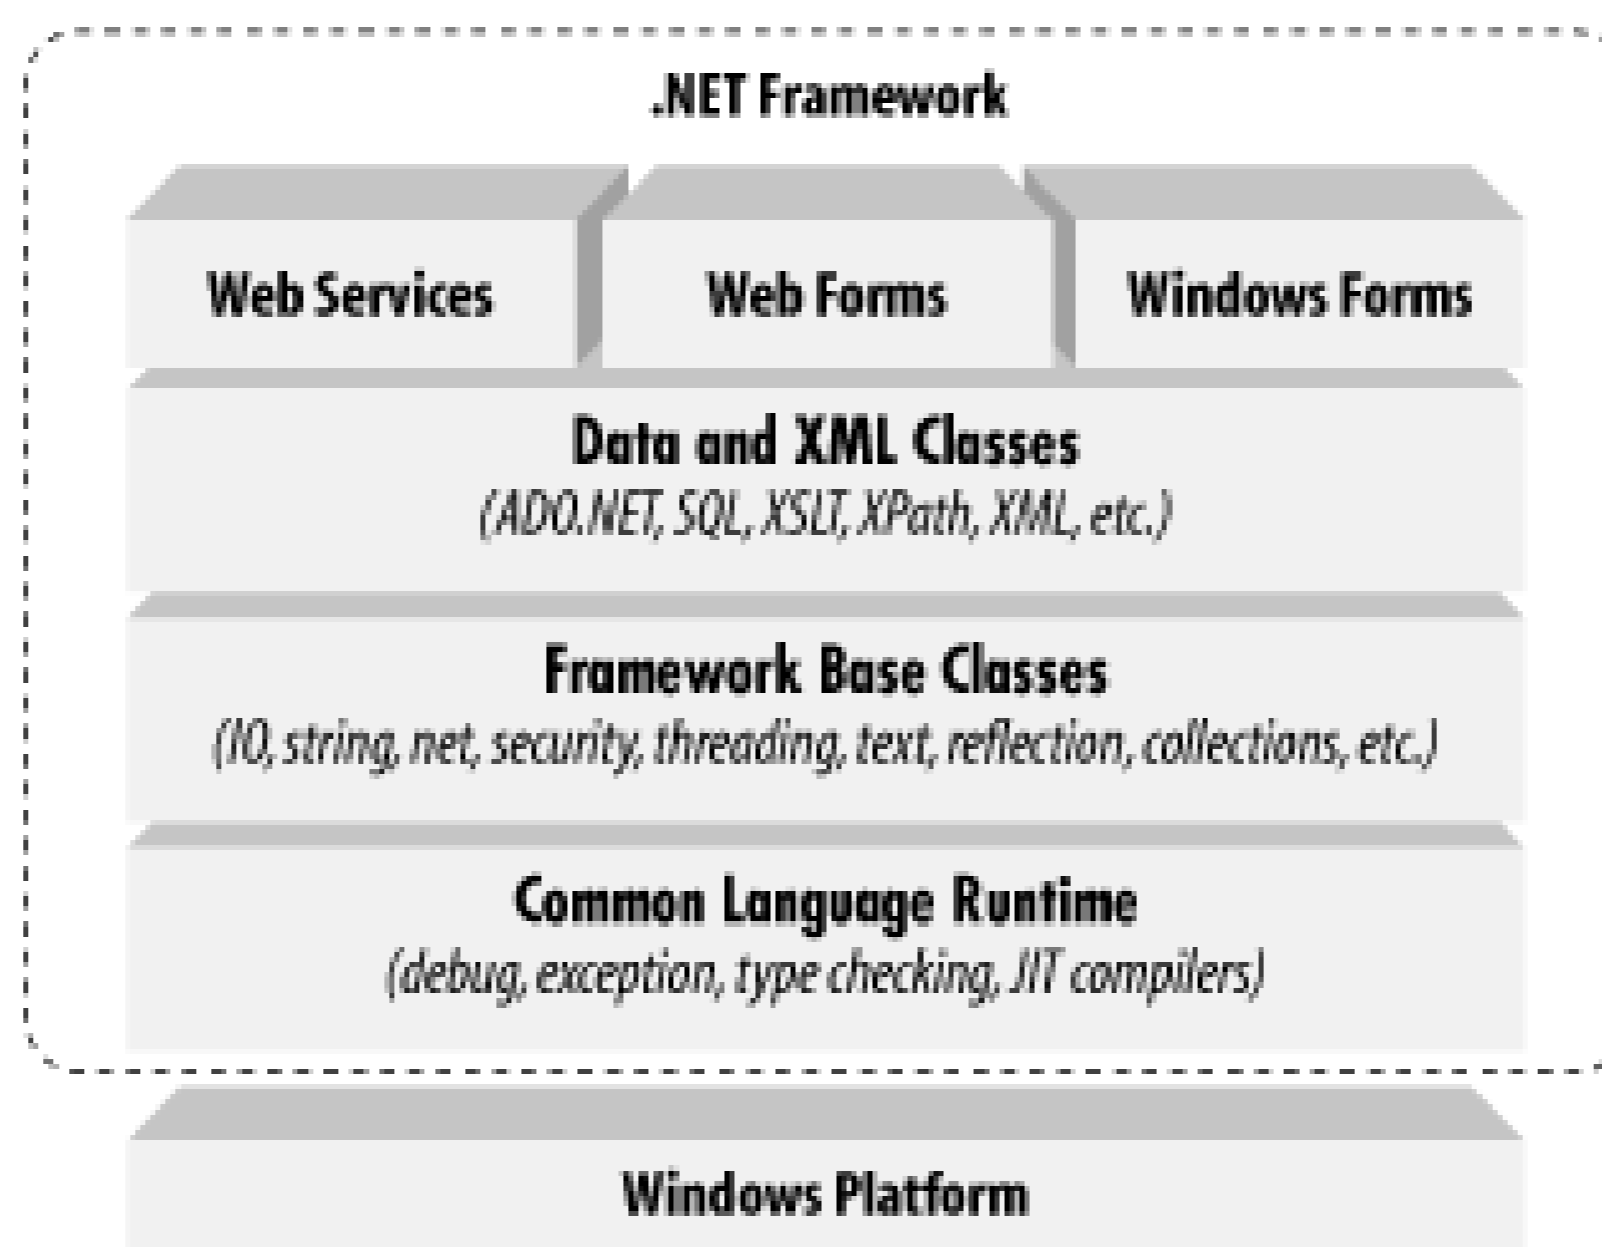
[\[Team LiB \]](#)

[[Team LiB](#)]

1.4 .NET Framework

Now that you are familiar with the major goals of the .NET Framework, let's briefly examine its architecture. As you can see in [Figure 1-2](#), the .NET Framework sits on top of the operating system, which can be a few different flavors of Windows and consists of a number of components (each of these components is discussed in greater detail starting with [Chapter 4](#), as described in the Preface). .NET is essentially a system application that runs on Windows.

Figure 1-2. The .NET Framework



The most important component of the framework is the CLR. If you are a Java programmer, think of the CLR as the .NET equivalent of the Java Virtual Machine (JVM). If you don't know Java, think of the CLR as the heart and soul of the .NET architecture. At a high level, the CLR activates objects, performs security checks on them, lays them out in memory, executes them, and garbage-collects them.

Conceptually, the CLR and the JVM are similar in that they are both runtime infrastructures that abstract the underlying platform differences. However, while the JVM officially supports only the Java language, the CLR supports any language that can be represented in its Common Intermediate Language (CIL). The JVM executes bytecode, so it can, in principle, support many languages, too. Unlike Java's bytecode, though, CIL is never interpreted. Another conceptual difference between the two infrastructures is that Java code runs on any platform with a JVM, whereas .NET code runs only on platforms that support the CLR. In April, 2003, the International Organization for Standardization and the International Electrotechnical Committee (ISO/IEC) recognized a functional subset of the CLR, known as the Common Language Interface (CLI), as an international standard. This development, initiated by Microsoft and developed by ECMA International, a European standards organization, opens the way for third parties to implement their own versions of the CLR on other platforms, such as Linux or Mac OS X. For information on third-party and open source projects working to implement the ISO/IEC CLI and C# specifications, see [Appendix A](#).

In [Figure 1-2](#), the layer on top of the CLR is a set of framework base classes. This set of classes is

similar to the set of classes found in STL, MFC, ATL, or Java. These classes support rudimentary input and output functionality, string manipulation, security management, network communications, thread management, text management, reflection functionality, collections functionality, as well as other functions.

On top of the framework base classes is a set of classes that extend the base classes to support data management and XML manipulation. These classes, called ADO.NET, support persistent data management—data that is stored on backend databases. Alongside the data classes, the .NET Framework supports a number of classes to let you manipulate XML data and perform XML searching and XML translations.

Classes in three different technologies (including web services, Web Forms, and Windows Forms) extend the framework base classes and the data and XML classes. Web services include a number of classes that support the development of lightweight distributed components, which work even in the face of firewalls and NAT software. These components support plug-and-play across the Internet, because web services employ standard HTTP and SOAP.

Web Forms, the key technology behind ASP.NET, include a number of classes that allow you to rapidly develop web Graphical User Interface (GUI) applications. If you're currently developing web applications with Visual Interdev, you can think of Web Forms as a facility that allows you to develop web GUIs using the same drag-and-drop approach as if you were developing the GUIs in Visual Basic. Simply drag-and-drop controls onto your Web Form, double-click on a control, and write the code to respond to the associated event.

Windows Forms support a set of classes that allow you to develop native Windows GUI applications. You can think of these classes collectively as a much better version of the MFC in C++ because they support easier and more powerful GUI development and provide a common, consistent interface that can be used in all languages.

In the next chapter, we examine the internals of the CLR and how it supports and executes .NET components, formally called *assemblies* in .NET.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Chapter 2. The Common Language Runtime

The most important component of the .NET Framework is the Common Language Runtime (CLR). The CLR manages and executes code written in .NET languages and is the basis of the .NET architecture, similar to the Java Virtual Machine. The CLR activates objects, performs security checks on them, lays them out in memory, executes them, and garbage-collects them.

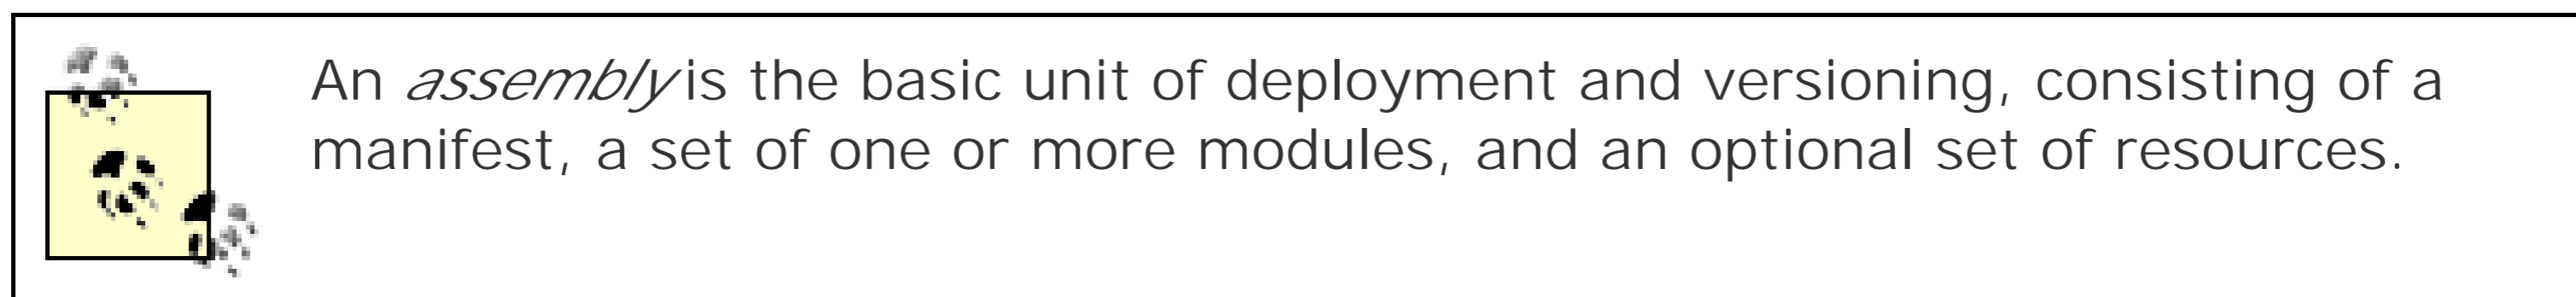
In this chapter, we describe the CLR environment, executables (with examples in several languages), metadata, assemblies, manifests, the CTS, and the CLS.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

2.1 CLR Environment

The CLR is the underlying .NET infrastructure. Its facilities cover all the goals that we spelled out in [Chapter 1](#). Unlike software libraries such as MFC or ATL, the CLR is built from a clean slate. The CLR manages the execution of code in the .NET Framework.



[Figure 2-1](#) shows the two portions of the .NET environment, with the bottom portion representing the CLR and the top portion representing the CLR executables or Portable Executable (PE) files, which are .NET assemblies or units of deployment. The CLR is the runtime engine that loads required classes, performs just-in-time compilation on needed methods, enforces security checks, and accomplishes a bunch of other runtime functionalities. The CLR executables shown in [Figure 2-1](#) are either EXE or DLL files that consist mostly of metadata and code.

Figure 2-1. The CLR environment

[\[Team LiB \]](#)

[\[Team LiB \]](#)

2.2 CLR Executables

Microsoft .NET executables are different from typical Windows executables in that they carry not only code and data, but also metadata (see [Section 2.3](#) and [Section 2.5](#) later in this chapter). In this section, we start off with the code for several .NET applications, and discuss the .NET PE format.

2.2.1 Hello, World: Managed C++

Let's start off by examining a simple *Hello, World* application written in Managed C++, a Microsoft .NET extension to the C++ language. Managed C++ includes a number of new CLR-specific keywords that permit C++ programs to take advantage of CLR features, including garbage collection. Here's the Managed C++ version of our program:

```
#using <mscorlib.dll>

using namespace System;

void main( )
{
    Console::WriteLine(L"C++ Hello, World!");
}
```

As you can see, this is a simple C++ program with an additional directive, **#using** (shown in bold). If you have worked with the Microsoft Visual C++ compiler support features for COM, you may be familiar with the **#import** directive. While **#import** reverse-engineers type information to generate wrapper classes for COM interfaces, **#using** makes all types accessible from the specified DLL, similar to a **#include** directive in C or C++. However, unlike **#include**, which imports C or C++ types, **#using** imports types for any .NET assembly, written in any .NET language.

The one and only statement within the `main()` method is self-explanatory—it means that we are invoking a static or class-level method, `WriteLine()`, on the `Console` class. The `L` that prefixes the literal string tells the C++ compiler to convert the literal into a Unicode string. You may have already guessed that the `Console` class is a type hosted by `mscorlib.dll`, and the `WriteLine()` method takes one string parameter.

One thing that you should also notice is that this code signals to the compiler that we're using the types in the `System` namespace, as indicated by the `using namespace` statement. This allows us to refer to `Console` instead of having to fully qualify this class as `System::Console`.

Given this simple program, compile it using the new C++ command-line compiler shipped with the .NET SDK:

```
cl hello.cpp /CLR /link /entry:main
```

The `/CLR` command-line option is extremely important, because it tells the C++ compiler to generate a .NET PE file instead of a normal Windows PE file.

When this statement is executed, the C++ compiler generates an executable called *hello.exe*. When you run *hello.exe*, the CLR loads, verifies, and executes it.

2.2.2 Hello, World: C#

Because .NET is serious about language integration, we'll illustrate this same program using C#, a language especially designed for .NET. Borrowing from Java and C++ syntax, C# is a simple and object-oriented language that Microsoft has used to write the bulk of the .NET base classes and tools. If you are a Java (or C++) programmer, you should have no problem understanding C# code. Here's *Hello, World* in C#:

```
using System;

public class MainApp
{
    public static void Main( )
    {
        Console.WriteLine("C# Hello, World!");
    }
}
```

C# is similar to Java in that it doesn't have the concept of a header file: class definitions and implementations are stored in the same *.cs* file. Another similarity to Java is that `Main()` is a public, static function of a particular class, as you can see from the code. This is different from C++, where `main()` itself is a global function.

The `using` keyword here functions similar to `using namespace` in the previous example, in that it signals to the C# compiler that we want to use types within the System namespace. Here's how to compile this C# program:

```
csc hello.cs
```

In this command, `csc` is the C# compiler that comes with the .NET SDK. Again, the result of executing this command is an executable called *hello.exe*, which you can execute like a normal EXE but it's managed by the CLR.

2.2.3 Hello, World: VB.NET

Here is the same program in Visual Basic .NET (VB.NET):

```
Imports System

Public Module MainApp
    Public Sub Main( )
        Console.WriteLine ("VB Hello, World!")
    End Sub
End Module
```

If you are a VB programmer, you may be in for a surprise. The syntax of the language has changed

quite a bit, but luckily these changes make the language mirror other object-oriented languages, such as C# and C++. Look carefully at this code snippet, and you will see that you can translate each line of code here into an equivalent in C#. Whereas C# uses the keywords `using` and `class`, VB.NET uses the keywords `Import` and `Module`, respectively. Here's how to compile this program:

```
vbc /t:exe /out:hello.exe hello.vb
```

Microsoft now provides a command-line compiler, *vbc*, for VB.NET. The `/t` option specifies the type of PE file to be created. In this case, since we have specified an EXE, *hello.exe* will be the output of this command.

2.2.4 Hello, World: J#

And since Microsoft has added the Visual J# compiler, which allows programmers to write Java code that targets the CLR, we'll show the same program in J# for completeness:

```
import System.*;

public class MainApp
{
    public static void main( )
    {
        Console.WriteLine("J# hello world!");
    }
}
```

If you carefully compare this simple J# program with the previously shown C# program, you'll notice that the two languages are very similar. For example, the only difference (other than the obvious literal string) is that the J# version uses the `import` directive, instead of the `using` directive. Here's how to compile this program:

```
vjc hello.jsl
```

In this command, `vjc` is the J# compiler that comes with the .NET SDK. The result of executing this command is an executable called *hello.exe*, targeting the CLR.

In all four versions of this *Hello, World* program, the `Console` class and the `WriteLine()` method have remained constant. That is, no matter which language you're using, once you know how to do something in one language, you can do it in all other languages that target the CLR. This is an extreme change from traditional Windows programming, in which if you know how to write to a file in C++, you probably won't know how to do it for VB, Java, or Cobol.

2.2.5 .NET Portable Executable File

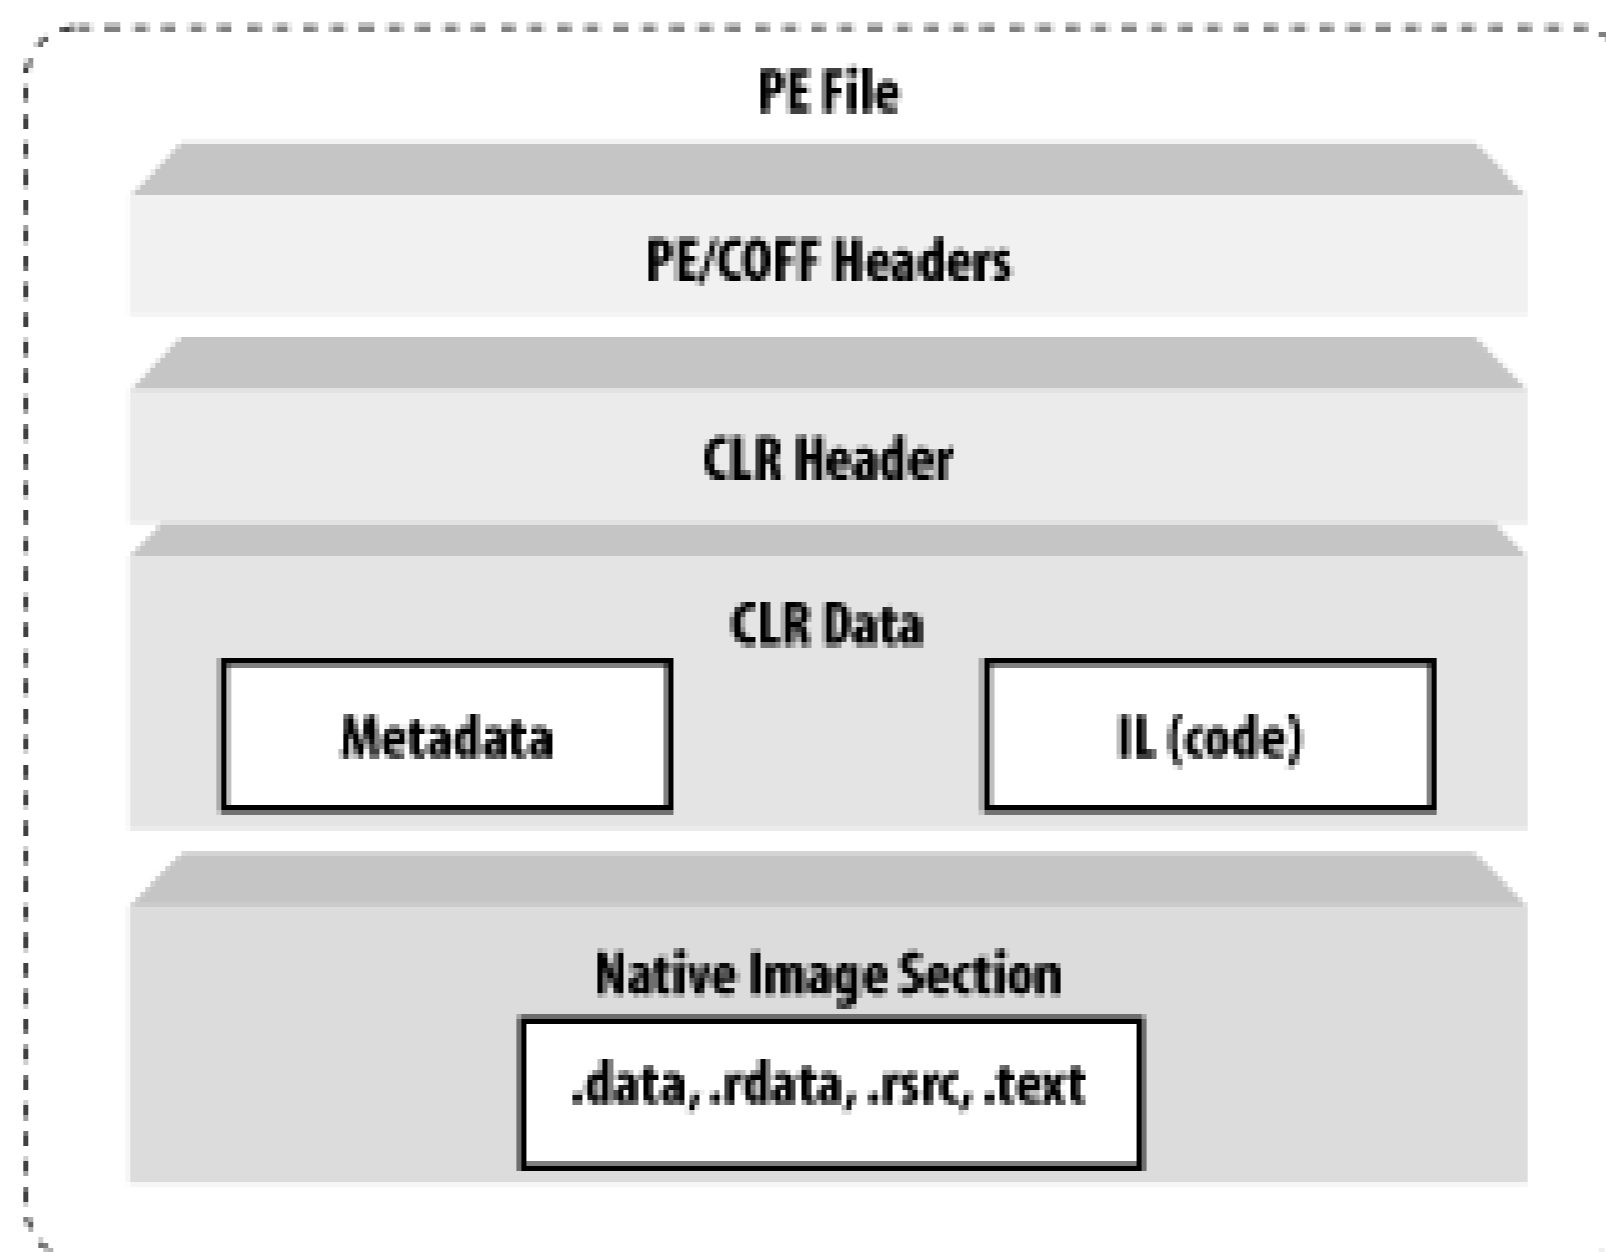
A Windows executable, EXE or DLL, must conform to a file format called the *PE* file format, which is a derivative of the Common Object File Format (COFF). Both of these formats are fully specified and publicly available. The Windows OS knows how to load and execute DLLs and EXEs because it

understands the format of a PE file. As a result, any compiler that wants to generate Windows executables must obey the PE/COFF specification.

A standard Windows PE file is divided into a number of sections, starting off with an MS-DOS header, followed by a PE header, followed by an optional header, and finally followed by a number of native image sections, including the `.text`, `.data`, `.rdata`, and `.rsrc` sections. These are the standard sections of a typical Windows executable, but Microsoft's C/C++ compiler allows you to add your own custom sections into the PE file using a compiler `#pragma` directive. For example, you can create your own data section to hold encrypted data that only you can read.

To support the CLR, Microsoft has extended the PE/COFF file format to include metadata and IL code. The CLR uses metadata to determine how to load classes and uses the IL code to turn it into native code for execution. As shown in [Figure 2-2](#), the extensions that Microsoft has added to the normal PE format include the CLR header and CLR data. The CLR header mainly stores relative virtual addresses (RVA) to locations that hold pertinent information to help the CLR manage program execution. The CLR data portion contains metadata and IL code, both of which determine how the program will be executed. Compilers that target the CLR must emit both the CLR header and data information into the generated PE file, otherwise the resulting PE file will not run under the CLR.

Figure 2-2. The format of a .NET PE file



If you want to prove to yourself that a .NET executable contains this information, use the `dumpbin.exe` utility, which dumps the content of a Windows executable in readable text^[1]

^[1] Note that you can dump the same information, in a more readable format, using the `ildasm.exe` utility, to be discussed later in this chapter.

For example, running the following command on the command prompt:

```
dumpbin.exe hello.exe /all
```

generates the following data (for brevity, we have shown only the main elements that we want to illustrate):

```
Microsoft (R) COFF/PE Dumper Version 7.10.2292
Copyright (C) Microsoft Corporation. All rights reserved.
```

Dump of file hello.exe

PE signature found

File Type: EXECUTABLE IMAGE

FILE HEADER VALUES /* 128-BYTE MS-DOS/COFF HEADER */

14C machine (x86)

. . .

OPTIONAL HEADER VALUES /* FOLLOWED BY PE AND OPTIONAL HEADERS */

10B magic # (PE32)

. . .

SECTION HEADER #1 /* CODE SECTION */

.text name

. . .

Looking at this text dump of a .NET PE file, you can see that a PE file starts off with the MS-DOS/COFF header, which all Windows programs must include. Following this header, you will find the PE header that supports Windows 32-bit programs. Immediately after the PE headers, you can find the code section for this program. The raw data (**RAW DATA #1**) of this section stores the CLR header, as follows:

RAW DATA #1

. . .

clr Header: /* CLR HEADER */

48 cb

2.00 **runtime version**

207C [214] RVA [size] of **MetaData** Directory

1 flags

6000001 **entry point token**

0 [0] RVA [size] of Resources Directory

0 [0] RVA [size] of StrongNameSignature Directory

0 [0] RVA [size] of CodeManagerTable Directory

0 [0] RVA [size] of VTableFixups Directory

0 [0] RVA [size] of ExportAddressTableJumps Directory

Section contains the following imports:

mscoree.dll

. . .

0 **_CorExeMain**

. . .

As mentioned earlier, the CLR header holds a number of pertinent details required by the runtime, including:

Runtime version

Indicates the runtime version that is required to run this program

MetaData directory

Is important because it indicates the location of the metadata needed by the CLR at runtime
Entry point token

Is even more important because, for a single file assembly, this is the token that signifies the entry point, such as `Main()`, that the CLR executes

Below the CLR Header, note that there is an imported function called `_CorExeMain`, which is implemented by `mscorlib.dll`, the core execution engine of the CLR.^[2] At the time of this writing, Windows 98, 2000, and Me have an OS loader that knows how to load standard PE files. To prevent massive changes to these operating systems and still allow .NET applications to run on them, Microsoft has updated the OS loaders for all these platforms. The updated loaders know how to check for the CLR header, and, if this header exists, it executes `_CorExeMain`, thus not only jumpstarting the CLR but also surrendering to it. You can then guess that the CLR will call `Main()`, since it can find the *entry point token* within the CLR header.^[3]

[2] We invite you to run `dumpbin.exe` and view the exports of `mscorlib.dll` at your convenience. You will also find `_CorDllMain`, `_CorExeMain`, `_CorImageUnloading`, and other interesting exports. It's interesting to note that this DLL is an in-process COM server, attesting that .NET is created using COM techniques.

[3] For brevity, we've covered only the important content of this header. If you want to learn the meanings of the rest, see this chapter's example code, which you can download from <http://www.oreilly.com/catalog/dotnetfrmess3/>.

Now that we've looked at the contents of the CLR header, let's examine the contents of the CLR data including metadata and code, which are arguably the most important elements in .NET.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

2.3 Metadata

Metadata is machine-readable information about a resource, or "data about data." Such information might include details on content, format, size, or other characteristics of a data source. In .NET, metadata includes type definitions, version information, external assembly references, and other standardized information.

In order for two systems, components, or objects to interoperate with one another, at least one must know something about the other. In COM, this "something" is an interface specification, which is implemented by a component provider and used by its consumers. The interface specification contains method prototypes with full signatures, including the type definitions for all parameters and return types.

Only C/C++ developers were able to readily modify or use Interface Definition Language (IDL) type definitions—not so for VB or other developers, and more importantly, not for tools or middleware. So Microsoft invented something other than IDL that everyone could use, called a *type library*. In COM, type libraries allow a development environment or tool to read, reverse engineer, and create wrapper classes that are most appropriate and convenient for the target developer. Type libraries also allow runtime engines, such as the VB, COM, MTS, or COM+ runtime, to inspect types at runtime and provide the necessary plumbing or intermediary support for applications to use them. For example, type libraries support dynamic invocation and allow the COM runtime to provide universal marshaling^[4] for cross-context invocations.

^[4] In COM, *universal marshaling* is a common way to marshal all data types. A universal marshaler can be used to marshal all types, so you don't have to provide your own proxy or stub code.

Type libraries are extremely rich in COM, but many developers criticize them for their lack of standardization. The .NET team invented a new mechanism for capturing type information. Instead of using the term "type library," we call such type information *metadata* in .NET.

2.3.1 Type Libraries on Steroids

Just as type libraries are C++ header files on steroids, metadata is a type library on steroids. In .NET, metadata is a common mechanism or dialect that the .NET runtime, compilers, and tools can all use. Microsoft .NET uses metadata to describe all types that are used and exposed by a particular .NET assembly. In this sense, metadata describes an assembly in detail, including descriptions of its identity (a combination of an assembly name, version, culture, and public key), the types that it references, the types that it exports, and the security requirements for execution. Much richer than a type library, metadata includes descriptions of an assembly and modules, classes, interfaces, methods, properties, fields, events, global methods, and so forth.

Metadata provides enough information for any runtime, tool, or program to find out literally everything that is needed for component integration. Let's take a look at a short list of consumers that make intelligent use of metadata in .NET, just to prove that metadata is indeed like type libraries on steroids:

CLR

The CLR uses metadata for verification, security enforcement, cross- context marshaling, memory layout, and execution. The CLR relies heavily on metadata to support these runtime features, which we will cover in a moment.

Class loader

A component of the CLR, the class loader uses metadata to find and load .NET classes. This is because metadata records detailed information for a specific class and where the class is located, whether it is in the same assembly, within or outside of a specific namespace, or in a dependent assembly somewhere on the network.

Just-in-time (JIT) compilers

JIT compilers use metadata to compile IL code. IL is an intermediate representation that contributes significantly to language-integration support, but it is not VB code or bytecode, which must be interpreted. .NET JIT compiles IL into native code prior to execution, and it does this using metadata.

Tools

Tools use metadata to support integration. For example, development tools can use metadata to generate callable wrappers that allow .NET and COM components to intermingle. Tools such as debuggers, profilers, and object browsers can use metadata to provide richer development support. One example of this is the IntelliSense features that Microsoft Visual Studio .NET supports. As soon as you have typed an object and a dot, the tool displays a list of methods and properties from which you can choose. This way, you don't have to search header files or documentation to obtain the exact method or property names and calling syntax.

Like the CLR, any application, tool, or utility that can read metadata from a .NET assembly can make use of that assembly. You can use the .NET reflection classes to inspect a .NET PE file and know everything about the data types that the assembly uses and exposes. The CLR uses the same set of reflection classes to inspect and provide runtime features, including memory management, security management, type checking, debugging, remoting, and so on.

Metadata ensures language interoperability, an essential element to .NET, since all languages must use the same types in order to generate a valid .NET PE file. The .NET runtime cannot support features such as memory management, security management, memory layout, type checking, debugging, and so on without the richness of metadata. Therefore, metadata is an extremely important part of .NET-so important that we can safely say that there would be no .NET without metadata.

2.3.2 Examining Metadata

At this point, we introduce an important .NET tool, the IL disassembler (*ildasm.exe*), which allows you to view both the metadata and IL code within a given .NET PE file. For example, if you execute *ildasm.exe* and open the *hello.exe* .NET PE file that you built earlier in this chapter, you will see something similar to [Figure 2-3](#).

Figure 2-3. The ildasm.exe tool



The *ildasm.exe* tool displays the metadata for your .NET PE file in a tree view, so that you can easily drill down from the assembly, to the classes, to the methods, and so on. To get full details on the contents of a .NET PE file, you can press Ctrl-D to dump the contents out into a text file.^[5]

[5] The *ildasm.exe* tool also supports a command-line interface. You can execute *ildasm.exe /h* to view the command-line options. As a side note, if you want to view exactly which types are defined and referenced, press Ctrl-M in the *ildasm.exe* GUI, and it will show you further details.

Here's an example of an *ildasm.exe* dump, showing only the contents that are relevant to the current discussion:

```
.assembly extern mscorlib
{
}

.assembly hello
{
}

.module hello.exe

.class private auto ansi beforefieldinit MainApp
    extends [mscorlib]System.Object
{
    .method public hidebysig static
        void Main( ) cil managed
    {
    } // End of method MainApp::Main

    .method public hidebysig specialname rtspecialname
        instance void .ctor( ) cil managed
    {
    } // End of method MainApp::.ctor
} // End of class MainApp
```

As you can see, this dump fully describes the type information and dependencies in a .NET assembly. While the first IL instruction, *.assembly extern*, tells us that this PE file references (i.e., uses) an external assembly called *mscorlib*, the second IL instruction describes our assembly, the one that is called *hello*. We will discuss the contents of the *.assembly* blocks later, as these are collectively called a *manifest*. Below the manifest, you see an instruction that tells us the module name, *hello.exe*.

Next, you see a definition of a class in IL, starting with the *.class* IL instruction. Notice this class,

MainApp, derives from System.Object, the mother of all classes in .NET. Although we didn't derive MainApp from System.Object when we wrote this class earlier in Managed C++, C#, J#, or VB.NET, the compiler automatically added this specification for us because System.Object is the implicit parent of all classes that omit the specification of a base class.

Within this class, you see two methods. While the first method, `Main()`, is a static method that we wrote earlier, the second method, `.ctor()`, is automatically generated. `Main()` serves as the main entry point for our application, and `.ctor()` is the constructor that allows anyone to instantiate `MainApp`.

As this example illustrates, given a .NET PE file, we can examine all the metadata that is embedded within a PE file. The important thing to keep in mind here is that we can do this without the need for source code or header files. If we can do this, imagine the exciting features that the CLR or a third-party tool can offer by simply making intelligent use of metadata. Of course, everyone can now see your code, unless you use different techniques (e.g., obfuscation and encryption) to protect your property rights.

2.3.3 Inspecting and Emitting Metadata

To load and inspect a .NET assembly to determine what types it supports, use a special set of classes provided by the .NET Framework base class library. Unlike API functions, these classes encapsulate a number of methods to give you an easy interface for inspecting and manipulating metadata. In .NET, these classes are collectively called the *Reflection API*, which includes classes from the System.Reflection and System.Reflection.Emit namespaces. The classes in the System.Reflection namespace allow you to inspect metadata within a .NET assembly, as shown in the following example:

```
using System;
using System.IO;

using System.Reflection;

public class Meta
{
    public static int Main( )
    {
        // First, load the assembly.
        Assembly a = Assembly.LoadFrom("hello.exe");

        // Get all the modules that the assembly supports.
        Module[] m = a.GetModules( );

        // Get all the types in the first module.
        Type[] types = m[0].GetTypes( );

        // Inspect the first type.
        Type type = types[0];
        Console.WriteLine("Type [{0}] has these methods:", type.Name);

        // Inspect the methods supported by this type.
        MethodInfo[] mInfo = type.GetMethods( );
    }
}
```

```

        foreach ( MethodInfo mi in mInfo )
        {
            Console.WriteLine(" {0}", mi);
        }

        return 0;
    }
}

```

Looking at this simple C# program, you'll notice that we first tell the compiler that we want to use the classes in the `System.Reflection` namespace because we want to inspect metadata. In `Main()`, we load the assembly by a physical name, *hello.exe*, so be sure that you have this PE file in the same directory when you run this program. Next, we ask the loaded assembly object for an array of modules that it contains. From this array of modules, we pull off the array of types supported by the module, and from this array of types, we then pull off the first type. For *hello.exe*, the first and only type happens to be `MainApp`. Once we have obtained this type or class, we loop through the list of its exposed methods. If you compile and execute this simple program, you see the following result:

Type [MainApp] has these methods:

```

    Int32 GetHashCode( )
    Boolean Equals(System.Object)
    System.String ToString( )
    Void Main( )
    System.Type GetType( )

```

Although we've written only the `Main()` function, our class actually supports four other methods, as is clearly illustrated by this output. There's no magic here, because `MainApp` inherits these method implementations from `System.Object`, which once again is the root of all classes in .NET.

As you can see, the `System.Reflection` classes allow you to inspect metadata, and they are really easy to use. If you have used type library interfaces in COM before, you know that you can do this in COM, but with much more effort. However, what you can't do with the COM type library interfaces is create a COM component at runtime—a missing feature in COM but an awesome feature in .NET. By using the `System.Reflection.Emit` classes, you can write a simple program to generate a .NET assembly dynamically at runtime. Given the existence of `System.Reflection.Emit`, anyone can write a custom .NET compiler.

2.3.4 Interoperability Support

Because it provides a common format for specifying types, metadata allows different components, tools, and runtimes to support interoperability. As demonstrated earlier, you can inspect the metadata of any .NET assembly. You can also ask an object at runtime for its type, methods, properties, events, and so on. Tools can do the same. The Microsoft .NET SDK ships four important tools that assist interoperability, including the .NET assembly registration utility (*RegAsm.exe*), the type library exporter (*tlbexp.exe*), the type library importer (*tlbimp.exe*), and the XML schema definition tool (*xsd.exe*).

You can use the .NET assembly registration utility to register a .NET assembly into the registry so COM clients can make use of it. The type library exporter is a tool that generates a type library file (*.tlb*) when you pass it a .NET assembly. Once you have generated a type library from a given .NET assembly, you can import the type library into VC++ or VB and use the .NET assembly in exactly the

same way as if you were using a COM component. Simply put, the type library exporter makes a .NET assembly look like a COM component. The following command-line invocation generates a type library, called *hello.tlb*:

```
tlbexp.exe hello.exe
```

Microsoft also ships a counterpart to *tlbexp.exe*, the type library importer; its job is to make a COM component appear as a .NET assembly. So if you are developing a .NET application and want to make use of an older COM component, use the type library importer to convert the type information found in the COM component into .NET equivalents. For example, you can generate a .NET PE assembly using the following command:

```
tlbimp.exe COMServer.tlb
```

Executing this command will generate a .NET assembly in the form of a DLL (e.g., *COMServer.dll*). You can reference this DLL like any other .NET assembly in your .NET code. When your .NET code executes at runtime, all invocations of the methods or properties within this DLL are directed to the original COM component.



Be aware that the type library importer doesn't let you reimport a type library that has been previously exported by the type library exporter. In other words, if you try to use *tlbimp.exe* on *hello.tlb*, which was generated by *tlbexp.exe*, *tlbimp.exe* will barf at you.

Another impressive tool that ships with the .NET SDK is the XML schema definition tool, which allows you to convert an XML schema into a C# class, and vice versa. This XML schema:

```
<schema xmlns="http://www.w3.org/2001/XMLSchema"
  targetNamespace="urn:book:car"
  xmlns:t="urn:book:car">
  <element name="car" type="t:CCar"/>
  <complexType name="CCar">
    <all>
      <element name="vin" type="string"/>
      <element name="make" type="string"/>
      <element name="model" type="string"/>
      <element name="year" type="int"/>
    </all>
  </complexType>
</schema>
```

represents a type called CCar. To convert this XML schema into a C# class definition, execute the following:

```
xsd.exe /c car.xsd
```

The */c* option tells the tool to generate a class from the given XSD file. If you execute this command, you get *car.cs* as the output that contains the C# code for this type.

The XML schema definition tool can also take a .NET assembly and generate an XSD file that contains representations for the public types within the .NET assembly. For example, if you execute the following, you get an XSD file as output:

```
xsd.exe somefile.exe
```

Before we leave this topic, we want to remind you to try out these tools for yourself, because they offer many impressive features that we won't cover in this introductory book.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

2.4 Assemblies and Manifests

As we just saw, types must expose their metadata to allow tools and programs to access them and benefit from their services. Metadata for types alone is not enough. To simplify software plug-and-play and configuration or installation of the component or software, we also need metadata about the component that hosts the types. Now we'll talk about .NET assemblies (deployable units) and manifests (the metadata that describes the assemblies).

2.4.1 Assemblies Versus Components

During the COM era, Microsoft documentation inconsistently used the term *component* to mean a COM class or a COM module (DLLs or EXEs), forcing readers or developers to consider the context of the term each time they encountered it. In .NET, Microsoft has addressed this confusion by introducing a new concept, *assembly*, which is a software component that supports plug-and-play, much like a hardware component. Theoretically, a .NET assembly is approximately equivalent to a compiled COM module. In practice, an assembly can contain or refer to a number of types and physical files (including bitmap files, .NET PE files, and so forth) that are needed at runtime for successful execution. In addition to hosting IL code, an assembly is a basic unit of versioning, deployment, security management, side-by-side execution, sharing, and reuse, as we discuss next.

To review: an assembly is a logical DLL or EXE, and a manifest is a detailed description (metadata) of an assembly, including its version, what other assemblies it uses, and so on.

2.4.2 Unique Identities

Type uniqueness is important in RPC, COM, and .NET. Given the vast number of GUIDs in COM (application, library, class, and interface identifiers), development and deployment can be tedious because you must use these magic numbers in your code and elsewhere all the time. In .NET, you refer to a specific type by its readable name and its namespace. Since a readable name and its namespace are not enough to be globally unique, .NET guarantees uniqueness by using unique public/private key pairs. All assemblies that are shared (called *shared assemblies*) by multiple applications must be built with a public/private key pair. Public/private key pairs are used in public-key cryptography. Since public-key cryptography uses asymmetrical encryption, an assembly creator can sign an assembly with a private key, and anyone can verify that digital signature using the assembly creator's public key. However, because no one else will have the private key, no other individual can create a similarly signed assembly.

To sign an assembly digitally, you must use a public/private key pair to build your assembly. At build time, the compiler generates a hash of the assembly files, signs the hash with the private key, and stores the resulting digital signature in a reserved section of the PE file. The public key is also stored in the assembly.

To verify the assembly's digital signature, the CLR uses the assembly's public key to decrypt the assembly's digital signature, resulting in the original, calculated hash. In addition, the CLR uses the information in the assembly's manifest to dynamically generate a hash. This hash value is then compared with the original hash value. These values must match, or we must assume that someone has tampered with the assembly.

Now that we know how to sign and verify an assembly in .NET, let's talk about how the CLR ensures that a given application loads the trusted assembly with which it was built. When you or someone else builds an application that uses a shared assembly, the application's assembly manifest will include an 8-byte hash of the shared assembly's public key. When you run your application, the CLR dynamically derives the 8-byte hash from the shared assembly's public key and compares this value with the hash value stored in your application's assembly manifest. If these values match, the CLR assumes that it has loaded the correct assembly for you.^[6]

[6] You can use the .NET Strong (a.k.a., Shared) Name (*sn.exe*) utility to generate a new key pair for a shared assembly. Before you can share your assembly, you must register it in the Global Assembly Cache, or GAC—you can do this by using the .NET Global Assembly Cache Utility (*gacutil.exe*). The GAC is simply a directory called Assembly located under the Windows (*%windir%*) directory.

2.4.3 IL Code

An assembly contains the IL code that the CLR executes at runtime (see [Section 2.5](#) later in this chapter). The IL code typically uses types defined within the same assembly, but it also may use or refer to types in other assemblies. Although nothing special is required to take advantage of the former, the assembly must define references to other assemblies to do the latter, as we will see in a moment. There is one caveat: each assembly can have at most one entry point, such as `asDllImport()`, `WinMain()`, or `Main()`. You must follow this rule because when the CLR loads an assembly, it searches for one of these entry points to start assembly execution.

2.4.4 Versioning

There are four types of assemblies in .NET:

Static assemblies

These are the .NET PE files that you create at compile time. You can create static assemblies using your favorite compiler: *csc*, *cl*, *vjc*, or *vbc*.

Dynamic assemblies

These are PE-formatted, in-memory assemblies that you dynamically create at runtime using the classes in the `System.Reflection.Emit` namespace.

Private assemblies

These are static assemblies used by a specific application.

Shared assemblies

These are static assemblies that must have a unique shared name and can be used by any application.

An application uses a private assembly by referring to the assembly using a static path or through an

XML-based application configuration file. Although the CLR doesn't enforce versioning policies-checking whether the correct version is used-for private assemblies, it ensures that an application uses the correct shared assemblies with which the application was built. Thus, an application uses a specific shared assembly by referring to the specific shared assembly, and the CLR ensures that the correct version is loaded at runtime.

In .NET, an assembly is the smallest unit to which you can associate a version number; it has the following format:

```
<major_version>.<minor_version>.<build_number>.<revision>
```

2.4.5 Deployment

Since a client application's assembly manifest (to be discussed shortly) contains information on external references-including the assembly name and version the application uses-you no longer have to use the registry to store activation and marshaling hints as in COM. Using the version and security information recorded in your application's manifest, the CLR will load the correct shared assembly for you. The CLR does lazy loading of external assemblies and will retrieve them on demand when you use their types. Because of this, you can create downloadable applications that are small, with many small external assemblies. When a particular external assembly is needed, the runtime downloads it automatically without involving registration or computer restarts.

2.4.6 Security

The concept of a user identity is common in all development and operating platforms, but the concept of a *code identity*, in which even a piece of code has an identity, is new to the commercial software industry. In .NET, an assembly itself has a code identity, which includes information such as the assembly's shared name, version number, culture, public key, and where the code came from (local, intranet, or Internet). This information is also referred to as the assembly's evidence, and it helps to identify and grant permissions to code, particularly mobile code.

To coincide with the concept of a code identity, the CLR supports the concept of *code access*. Whether code can access resources or use other code is entirely dependent on security policy, which is a set of rules that an administrator configures and the CLR enforces. The CLR inspects the assembly's evidence and uses security policy to grant the target assembly a set of permissions to be examined during its execution. The CLR checks these permissions and determines whether the assembly has access to resources or to other code. When you create an assembly, you can declaratively specify a set of permissions that the client application must have in order to use your assembly. At runtime, if the client application has code access to your assembly, it can make calls to your assembly's objects; otherwise, a security exception will ensue. You can also imperatively demand that all code on the call stack has the appropriate permissions to access a particular resource.

2.4.7 Side-by-Side Execution

We have said that an assembly is a unit of versioning and deployment, and we've talked briefly about DLL Hell, something that .NET intends to minimize. The CLR allows any versions of the same, shared DLL (shared assembly) to execute at the same time, on the same system, and even in the same

process. This concept is known as *side-by-side execution*. Microsoft .NET accomplishes side-by-side execution by using the versioning and deployment features that are innate to all shared assemblies. This concept allows you to install any versions of the same, shared assembly on the same machine, without versioning conflicts or DLL Hell. The only caveat is that your assemblies must be public or shared assemblies, meaning that you must register them against the GAC using a tool such as the .NET Global Assembly Cache Utility (*gacutil.exe*). Once you have registered different versions of the same shared assembly into the GAC, the human-readable name of the assembly no longer matters-what's important is the information provided by .NET's versioning and deployment features.

Recall that when you build an application that uses a particular shared assembly, the shared assembly's version information is attached to your application's manifest. In addition, an 8-byte hash of the shared assembly's public key is also attached to your application's manifest. Using these two pieces of information, the CLR can find the exact shared assembly that your application uses, and it will even verify that your 8-byte hash is indeed equivalent to that of the shared assembly. Given that the CLR can identify and load the exact assembly, the end of DLL Hell is in sight.

2.4.8 Sharing and Reuse

When you want to share your assembly with the rest of the world, your assembly must have a shared or strong name, and you must register it in the GAC. Likewise, if you want to use or extend a particular class that is hosted by a particular shared assembly, you don't just import that specific class, but you import the whole assembly into your application. Therefore, the whole assembly is a unit of sharing.

Assemblies turn out to be an extremely important feature in .NET because they are an essential part of the runtime. An assembly encapsulates all types that are defined within the assembly. For example, although two different assemblies, Personal and Company, can define and expose the same type, Car, Car by itself has no meaning unless you qualify it as [Personal]Car or [Company]Car. Given this, all types are scoped to their containing assembly, and for this reason, the CLR cannot make use of a specific type unless the CLR knows the type's assembly. In fact, if you don't have an assembly manifest, which describes the assembly, the CLR will not execute your program.

2.4.9 Manifests: Assembly Metadata

An assembly *manifest* is metadata that describes everything about the assembly, including its identity, a list of files belonging to the assembly, references to external assemblies, exported types, exported resources, and permission requests. In short, it describes all the details that are required for component plug-and-play. Since an assembly contains all these details, there's no need for storing this type of information in the registry, as in the COM world.

In COM, when you use a particular COM class, you give the COM library a class identifier. The COM library looks up in the registry to find the COM component that exposes that class, loads the component, tells the component to give it an instance of that class, and returns a reference to this instance. In .NET, instead of looking into the registry, the CLR peers right into the assembly manifest determines which external assembly is needed, loads the exact assembly that's required by your application, and creates an instance of the target class.

Let's examine the manifest for the *hello.exe* application that we built earlier. Recall that we used the *ildasm.exe* tool to pick up this information.

```

.assembly extern mscorlib
{
    .publickeytoken = (B7 7A 5C 56 19 34 E0 89 )
    .ver 1:0:5000:0
}

.assembly hello
{
    .hash algorithm 0x00008004
    .ver 0:0:0:0
}
.module hello.exe
// MVID: {F828835E-3705-4238-BCD7-637ACDD33B78}

```

You'll notice that this manifest starts off identifying an external orreferenced assembly, with `mscorlib` as the assembly name, which this particular application references. The keywords `.assembly extern` tell the CLR that this application doesn't implement `mscorlib`, but makes use of it instead. This external assembly is one that all .NET applications will use, so you will see this external assembly defined in the manifest of all assemblies. You'll notice that, inside this assembly definition, the compiler has inserted a special value called the *publickeytoken*, which is basic information about the publisher of `mscorlib`. The compiler generates the value for `.publickeytoken` by hashing the public key associated with the `mscorlib` assembly. Another thing to note in the `mscorlib` block is the version number of `mscorlib`.^[7]

[7] The fascinating details are explained in Partition II Metadata.doc and Partition III CIL.doc, which come with the .NET SDK. If you really want to understand metadata IL, read these documents.

Now that we've covered the first `.assembly` block, let's examine the second, which describes this particular assembly. You can tell that this is a manifest block that describes our application's assembly because there's no `extern` keyword. The identity of this assembly is made up of a readable assembly name, `hello`, its version information, `0:0:0:0`, and an optional culture, which is missing. Within this block, the first line indicates the hash algorithm that is used to hash selected contents of this assembly, the result of which will be encrypted using the private key. However, since we are not sharing this simple assembly, there's no encryption and there's no `.publickey` value.

The last thing to discuss is `.module`, which simply identifies the output filename of this assembly, `hello.exe`. You'll notice that a module is associated with a GUID, which means you get a different GUID each time you build the module. Given this, a rudimentary test for exact module equivalence is to compare the GUIDs of two modules.

Because this example is so simple, that's all we get for our manifest. In a more complicated assembly, you can get all this, including much more in-depth detail about the make up of your assembly.

2.4.10 Creating Assemblies

An assembly can be a *single-module assembly* or a *multi-module assembly*. In a single-module assembly, everything in a build is clumped into one EXE or DLL, an example of which is the `hello.exe` application that we developed earlier. This is easy to create because a compiler takes care of creating the single-module assembly for you.

If you wanted to create a multi-module assembly, one that contains many modules and resource files, you have a few choices. One option is to use the Assembly Linker (*al.exe*) that is provided by the .NET SDK. This tool takes one or more IL or resource files and spits out a file with an assembly manifest.

2.4.11 Using Assemblies

To use an assembly, first import the assembly into your code, the syntax of which is dependent upon the language that you use. For example, this is how we import an assembly in C#, as we have seen previously in the chapter:

```
using System;
```

When you build your assembly, you must tell the compiler that you are referencing an external assembly. Again, how you do this is different depending on the compiler that you use. If you use the C# compiler, here's how it's done:

```
csc /r:microsoft.dll hello.cs
```

Earlier, we showed you how to compile *hello.cs* without the */r:* option, but both techniques are equivalent. The reference to *microsoft.dll* is inherently assumed because it contains all the base framework classes.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

2.5 Intermediate Language (IL)

In software engineering, the concept of *abstraction* is extremely important. We often use abstraction to hide the complexity of system or application services, providing instead a simple interface to the consumer. As long as we can keep the interface the same, we can change the hideous internals, and different consumers can use the same interface.

In language advances, scientists introduced different incarnations of language-abstraction layers, such as *p-code* and *bytecode*. Produced by the Pascal-P compiler, p-code is an intermediate language that supports procedural programming. Generated by Java compilers, bytecode is an intermediate language that supports object-oriented programming. Bytecode is a language abstraction that allows Java code to run on different operating platforms, as long as the platforms have a Java Virtual Machine (JVM) to execute bytecode.

Microsoft calls its own language-abstraction layer the Microsoft Intermediate Language (MSIL) or IL, for short. IL is an implementation of the Common Intermediate Language (CIL), a key element of the EMCA CLI specification. Similar to bytecode, IL supports all object-oriented features, including data abstraction, inheritance, polymorphism, and useful concepts such as exceptions and events. In addition to these features, IL supports other concepts, such as properties, fields, and enumeration. Any .NET language may be converted into IL, so .NET supports multiple languages and multiple platforms, as long as the target platforms have a CLR.

Shipped with the .NET SDK, *Partition III CIL.doc* describes the important IL instructions that language compilers should use. In addition to this specification, the .NET SDK includes another important document, *Partition II Metadata.doc*. Both of these documents are intended for developers who write compilers and tools, but you should read them to further understand how IL fits into .NET. Although you can develop a valid .NET assembly using the supported IL instructions and features, you'll find IL to be very tedious because the instructions are a bit cryptic. However, should you decide to write pure IL code, you could use the IL Assembler (*ilasm.exe*) to turn your IL code into a .NET PE file.^[8]

^[8] You can test this utility using the IL disassembler to load a .NET PE file and dump out the IL to a text file. Once you've done this, use the IL Assembler to convert the text file into a .NET PE file.

Enough with the theory: let's take a look at some IL. Here's an excerpt of IL code for the *hello.exe* program that we wrote earlier:^[9]

^[9] Don't compile this IL code: it's incomplete because we've extracted unclear details to make it easier to read. If you want to see the complete IL code, use *ildasm.exe* on *hello.exe*.

```
.class private auto ansi beforefieldinit MainApp
  extends [mscorlib]System.Object
{
  .method public hidebysig static
    void Main( ) cil managed
  {
    .entrypoint
    .maxstack 1
    ldstr "C# hello world!"
```

```

        call void [mscorlib]System.Console::WriteLine(string)
        ret
    } // End of method MainApp::Main

    .method public hidebysig specialname rtspecialname
        instance void .ctor( ) cil managed
    {
        .maxstack 1
        ldarg.0
        call instance void [mscorlib]System.Object::.ctor( )
        ret
    } // End of method MainApp::.ctor

} // End of class MainApp

```

Ignoring the weird-looking syntactic details, you can see that IL is conceptually the same as any other object-oriented language. Clearly, there is a class that is called `MainApp` that derives from `System.Object`. This class supports a static method called `Main()`, which contains the code to dump out a text string to the console. Although we didn't write a constructor for this class, our C# compiler has added the default constructor for `MainApp` to support object construction.

Since a lengthy discussion of IL is beyond the scope of this book, let's just concentrate on the `Main()` method to examine its implementation briefly. First, you see the following method signature:

```

.method public hidebysig static
    void Main( ) cil managed

```

This signature declares a method that is public (meaning that it can be called by anyone) and static (meaning it's a class-level method). The name of this method is `Main()`. `Main()` contains IL code that is to be managed or executed by the CLR. The `hidebysig` attribute says that this method hides the same methods (with the same signatures) defined earlier in the class hierarchy. This is simply the default behavior of most object-oriented languages, such as C++. Having gone over the method signature, let's talk about the method body itself:

```

{
    .entrypoint
    .maxstack 1
    ldstr "C# hello world!"
    call void [mscorlib]System.Console::WriteLine(string)
    ret
} // End of method MainApp::Main

```

This method uses two directives: `.entrypoint` and `.maxstack`. The `.entrypoint` directive specifies that `Main()` is the one and only entry point for this assembly. The `.maxstack` directive specifies the maximum stack slots needed by this method; in this case, the maximum number of stack slots required by `Main()` is one. Stack information is needed for each IL method because IL instructions are stack-based, allowing language compilers to generate IL code easily.

In addition to these directives, this method uses three IL instructions. The first IL instruction, `ldstr`, loads our literal string onto the stack so that the code in the same block can use it. The next IL instruction, `call`, invokes the `WriteLine()` method, which picks up the string from the stack. The `call` IL instruction expects the method's arguments to be on the stack, with the first argument being the first object pushed on the stack, the second argument being the second object pushed onto the

stack, and so forth. In addition, when you use the `call` instruction to invoke a method, you must specify the method's signature. For example, examine the method signature of `WriteLine()`:

```
void [mscorlib]System.Console::WriteLine(string)
```

and you'll see that `WriteLine()` is a static method of the `Console` class. The `Console` class belongs to the `System` namespace, which happens to be a part of the `mscorlib` assembly. The `WriteLine()` method takes a string (an alias for `System.String`) and returns a `void`. The last thing to note in this IL snippet is that the `ret` IL instruction simply returns control to the caller.



Since .NET assemblies contain IL code, your proprietary algorithms can be seen by anyone. To protect your intellectual property, use an obfuscator, either the one that comes with Visual Studio .NET or one that is commercially available.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

2.6 The CTS and CLS

Having seen the importance of metadata and IL, let's examine the CTS and the CLS. Both the CTS and the CLS ensure language compatibility, interoperability, and integration.

2.6.1 The Common Type System (CTS)

Because .NET treats all languages as equal, a class written in C# should be equivalent to a class written in VB.NET, and an interface defined in Managed C++ should be exactly the same as one that is specified in Managed COBOL. Languages must agree on the meanings of these concepts before they can integrate with one another. In order to make language integration a reality, Microsoft has specified a common type system by which every .NET language must abide. In this section, we outline the common types that have the same conceptual semantics in every .NET language. Microsoft .NET supports a rich set of types, but we limit our discussion to the important ones, including value types, reference types, classes, interfaces, and delegates.

2.6.1.1 Value types

In general, the CLR supports two different types: value types and reference types. *Value types* represent values allocated on the stack. They cannot be null and must always contain some data. When value types are passed into a function, they are passed by value, meaning that a copy of the value is made prior to function execution. This implies that the original value won't change, no matter what happens to the copy during the function call. Since intrinsic types are small in size and don't consume much memory, the resource cost of making a copy is negligible and outweighs the performance drawbacks of object management and garbage collection. Value types include primitives, structures, and enumerations; examples are shown in the following C# code listing:

```
int i; // Primitive
struct Point { int x, y; } // Structure
enum State { Off, On } // Enumeration
```

You can also create a value type by deriving a class from `System.ValueType`. One thing to note is that a value type is *sealed*, meaning that once you have derived a class from `System.ValueType`, no one else can derive from your class.

2.6.1.2 Reference types

If a type consumes significant memory resources, then a reference type provides more benefits over a value type. *Reference types* are so called because they contain references to heap-based objects and can be null. These types are passed by reference, meaning that when you pass such an object into a function, an address of or pointer to the object is passed—not a copy of the object, as in the case of a value type. Since you are passing a reference, the caller will see whatever the called

function does to your object. The first benefit here is that a reference type can be used as an output parameter, but the second benefit is that you don't waste extra resources because a copy is not made. If your object is large (consuming lots of memory), then reference types are a better choice. In .NET, one drawback of a reference type is that it must be allocated on the managed heap, which means it requires more CPU cycles because it must be managed and garbage-collected by the CLR. In .NET, the closest concept to destruction is finalization, but unlike destructors in C++, finalization is nondeterministic. In other words, you don't know when finalization will happen because it occurs when the garbage collector executes (by default, when the system runs out of memory). Since finalization is nondeterministic, another drawback of reference types is that if reference-type objects hold on to expensive resources that will be released during finalization, system performance will degrade because the resources won't be released until these objects are garbage-collected. Reference types include classes, interfaces, arrays, and delegates, examples of which are shown in the following C# code listing:

```
class Car {} // Class
interface ISteering {} // Interface
int[] a = new int[5]; // Array
delegate void Process( ); // Delegate
```

Classes, interfaces, and delegates will be discussed shortly.

2.6.1.3 Boxing and unboxing

Microsoft .NET supports value types for performance reasons, but everything in .NET is ultimately an object. In fact, all primitive types have corresponding classes in the .NET Framework. For example, `int` is, in fact, an alias of `System.Int32`, and `System.Int32` happens to derive from `System.ValueType`, meaning that it is a value type. Value types are allocated on the stack by default, but they can always be converted into a heap-based, reference-type object; this is called *boxing*. The following code snippet shows that we can create a box and copy the value of `i` into it:

```
int i = 1; // i - a value type
object box = i; // box - a reference object
```

When you box a value, you get an object upon which you can invoke methods, properties, and events. For example, once you have converted the integer into an object, as shown in this code snippet, you can call methods that are defined in `System.Object`, including `ToString()`, `Equals()`, and so forth.

The reverse of boxing is of course *unboxing*, which means that you can convert a heap-based, reference-type object into its value-type equivalent, as the following shows:

```
int j = (int)box;
```

This example simply uses the `cast` operator to cast a heap-based object called `box` into a value-type integer.

2.6.1.4 Classes, properties, indexers

The CLR provides full support for object-oriented concepts (such as encapsulation, inheritance, and polymorphism) and class features (such as methods, fields, static members, visibility, accessibility,

nested types, and so forth). In addition, the CLR supports new features that are nonexistent in many traditional object-oriented programming languages, including properties, indexers, and events.^[10] Events are covered in [Chapter 8](#). For now let's briefly talk about properties and indexers.

[10] An event is a callback that is implemented using delegates, which is covered shortly.

A property is similar to a field (a member variable), with the exception that there is a getter and a setter method, as follows:

```
using System;

public class Car
{
    private string make;
    public string Make
    {
        get { return make; }
        set { make = value; }
    }

    public static void Main( )
    {
        Car c = new Car( );
        c.Make = "Acura"; // Use setter.
        String s = c.Make; // Use getter.
        Console.WriteLine(s);
    }
}
```

Although this is probably the first time you've seen such syntax, this example is straightforward and really needs no explanation, with the exception of the keyword `value`. This is a special keyword that represents the one and only argument to the setter method.

Syntactically similar to a property, an indexer is analogous to `operator[]` in C++, as it allows array-like access to the contents of an object. In other words, it allows you to access an object like you're accessing an array, as shown in the following example:

```
using System;

public class Car
{
    Car( )
    {
        wheels = new string[4];
    }

    private string[] wheels;
    public string this[int index]
    {
        get { return wheels[index]; }
        set { wheels[index] = value; }
    }
}
```

```

public static void Main( )
{
    Car c = new Car( );
    c[0] = "LeftWheel"; // c[0] can be an l-value or an r-value.
    Console.WriteLine(c[0]);
}
}

```



Unlike C++, but similar to Java, classes in .NET support only single-implementation inheritance.

2.6.1.5 Interfaces

Interfaces support exactly the same concept as a C++ abstract base class (ABC) with only pure virtual functions. An ABC is a class that declares one or more pure virtual functions and thus cannot be instantiated. If you know COM or Java, interfaces in .NET are conceptually equivalent to a COM or Java interface. You specify them, but you don't implement them. A class that derives from your interface must implement your interface. An interface may contain methods, properties, indexers, and events. In .NET, a class can derive from multiple interfaces.

2.6.1.6 Delegates

One of the most powerful features of C is its support for function pointers. Function pointers allow you to build software with hooks that can be implemented by someone else. In fact, function pointers allow many people to build expandable or customizable software. Microsoft .NET supports a type-safe version of function pointers, called *delegates*. Here's an example that may take a few minutes to sink in, but once you get it, you'll realize that it's really simple:

```

using System;
class TestDelegate
{
    // 1. Define callback prototype.
    delegate void MsgHandler(string strMsg);

    // 2. Define callback method.
    void OnMsg(string strMsg)
    {
        Console.WriteLine(strMsg);
    }
}

public static void Main( )
{
    TestDelegate t = new TestDelegate( );

    // 3. Wire up our callback method.
    MsgHandler f = new MsgHandler(t.OnMsg);

    // 4. Invoke the callback method indirectly.
}

```

```
        f("Hello, Delegate.");  
    }  
}
```

The first thing to do is to define a callback function prototype, and the important keyword here is `delegate`, which tells the compiler that you want an object-oriented function pointer. Under the hood, the compiler generates a nested class, `MsgHandler`, which derives from `System.MulticastDelegate`.^[11] A multicast delegate supports many receivers. Once you've defined your prototype, you must define and implement a method with a signature that matches your prototype. Then, simply wire up the callback method by passing the function to the delegate's constructor, as shown in this code listing. Finally, invoke your callback indirectly. Having gone over delegates, you should note that delegates form the foundation of events, which are discussed in [Chapter 8](#).

[11] If you want to see this, use `ildasm.exe` and view the metadata of the `delegate.exe` sample that we've provided.

2.6.2 The Common Language Specification (CLS)

A goal of .NET is to support language integration in such a way that programs can be written in any language, yet can interoperate with one another, taking full advantage of inheritance, polymorphism, exceptions, and other features. However, languages are not made equal because one language may support a feature that is totally different from another language. For example, Managed C++ is case-sensitive, but VB.NET is not. In order to bring everyone to the same sheet of music, Microsoft has published the Common Language Specification (CLS). The CLS specifies a series of basic rules that are required for language integration. Since Microsoft provides the CLS that spells out the minimum requirements for being a .NET language, compiler vendors can build their compilers to the specification and provide languages that target .NET. Besides compiler writers, application developers should read the CLS and use its rules to guarantee language interoperability.

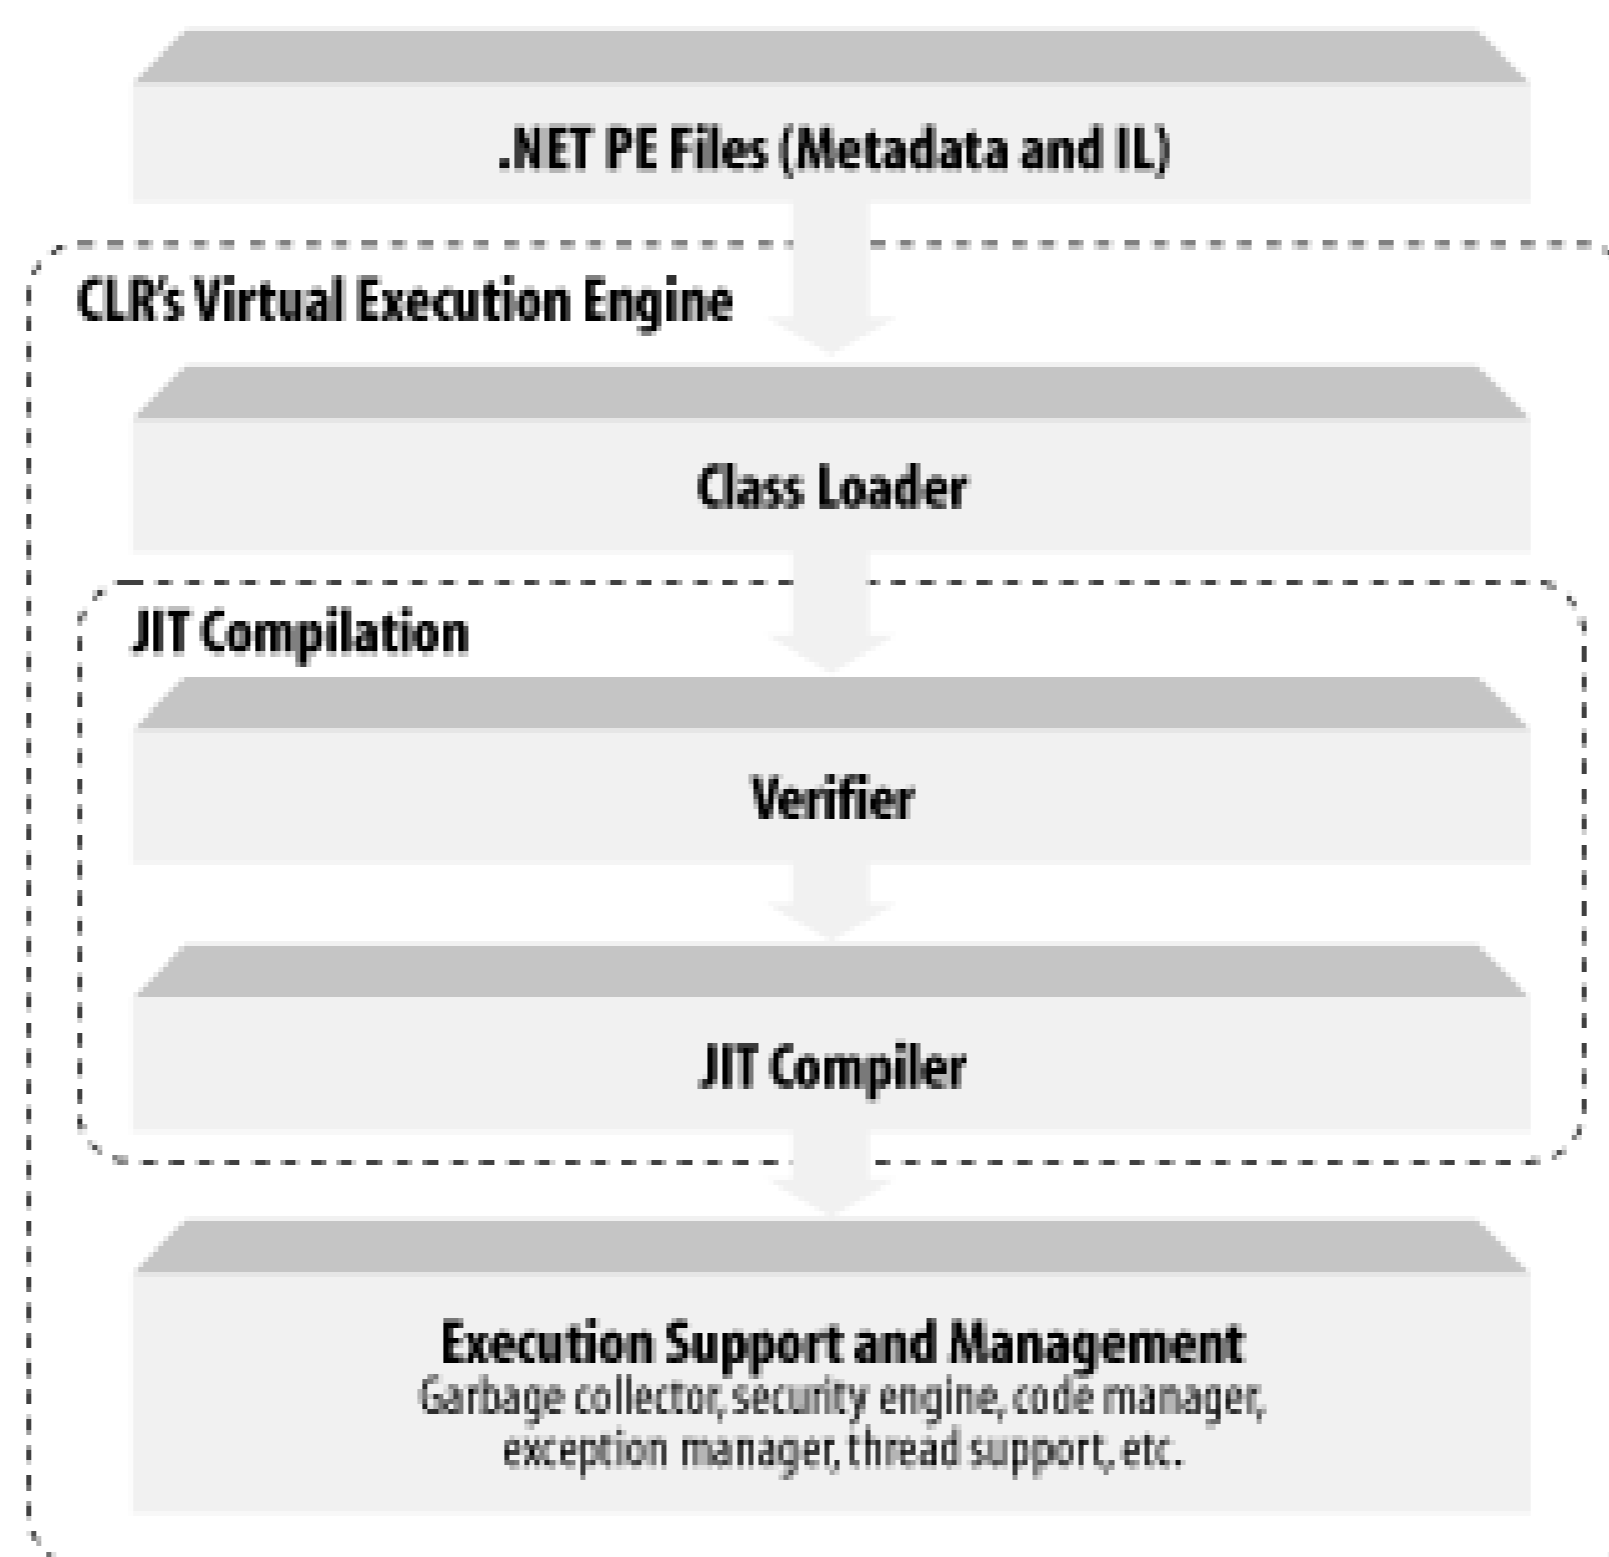
[\[Team LiB \]](#)

[[Team LiB](#)]

2.7 CLR Execution

Now that you understand the elements of a .NET executable, let's talk about the services that the CLR provides to support management and execution of .NET assemblies. There are many fascinating components in the CLR, but for brevity, we will limit our discussions to just the major components, as shown in [Figure 2-4](#).

Figure 2-4. Major CLR components: the Virtual Execution System (VES)



The major components of the CLR include the class loader, verifier, JIT compilers, and other execution support, such as code management, security management, garbage collection, exception management, debug management, marshaling management, thread management, and so on. As you can see from [Figure 2-4](#), your .NET PE files lay on top of the CLR and execute within the CLR's Virtual Execution System (VES), which hosts the major components of the runtime. Your .NET PE files will have to go through the class loader, the type verifier, the JIT compilers, and other execution support components before they will execute.

2.7.1 Class Loader

When you run a standard Windows application, the OS loader loads it before it can execute. At the time of this writing, the default loaders in the existing Windows operating systems, such as Windows 98, Windows Me, Windows 2000, and so forth, recognize only the standard Windows PE files. As a result, Microsoft has provided an updated OS loader for each of these operating systems that support the .NET runtime. The updated OS loaders know the .NET PE file format and can handle the file appropriately.

When you run a .NET application on one of these systems that have an updated OS loader, the OS loader recognizes the .NET application and thus passes control to the CLR. The CLR then finds the entry point, which is typically `Main()`, and executes it to jump-start the application. But before `Main()` can execute, the class loader must find the class that exposes `Main()` and load the class. In addition, when `Main()` instantiates an object of a specific class, the class loader also kicks in. In short, the class loader performs its magic the first time a type is referenced.

The *class loader* loads .NET classes into memory and prepares them for execution. Before it can successfully do this, it must locate the target class. To find the target class, the class loader looks in several different places, including the application configuration file (*.config*) in the current directory, the GAC, and the metadata that is part of the PE file, specifically the manifest. The information that is provided by one or more of these items is crucial to locating the correct target class. Recall that a class can be scoped to a particular namespace, a namespace can be scoped to a particular assembly, and an assembly can be scoped to a specific version. Given this, two classes, both named `Car`, are treated as different types even if the version information of their assemblies are the same.

Once the class loader has found and loaded the target class, it caches the type information for the class so that it doesn't have to load the class again for the duration of this process. By caching this information, it will later determine how much memory is needed to allocate for the newly created instance of this class. Once the target class is loaded, the class loader injects a small stub, like a function prolog, into every single method of the loaded class. This stub is used for two purposes: to denote the status of JIT compilation and to transition between managed and unmanaged code. At this point, if the loaded class references other classes, the class loader will also try to load the referenced types. However, if the referenced types have already been loaded, the class loader has to do nothing. Finally, the class loader uses the appropriate metadata to initialize the static variables and instantiate an object of the loaded class for you.

2.7.2 Verifier

Scripting and interpreted languages are very lenient on type usages, allowing you to write code without explicit variable declarations. This flexibility can introduce code that is extremely error-prone and hard to maintain, and that is often a culprit for mysterious program crashes. Unlike scripting and interpreted languages, compiled languages require types to be explicitly defined prior to their use, permitting the compiler to ensure that types are used correctly and the code will execute peacefully at runtime.

The key here is type safety, and it is a fundamental concept for code verification in .NET. Within the VES, the verifier is the component that executes at runtime to verify that the code is type safe. Note that this type verification is done at runtime and that this is a fundamental difference between .NET and other environments. By verifying type safety at runtime, the CLR can prevent the execution of code that is not type safe and ensure that the code is used as intended. In short, type safety means more reliability.

Let's talk about where the verifier fits within the CLR. After the class loader has loaded a class and before a piece of IL code can execute, the verifier kicks in for code that must be verified. The verifier is responsible for verifying that:

- The metadata is well formed, meaning the metadata must be valid.

- The IL code is type safe, meaning type signatures are used correctly.

Both of these criteria must be met before the code can be executed because JIT compilation will take place only when code and metadata have been successfully verified. In addition to checking for type safety, the verifier also performs rudimentary control-flow analysis of the code to ensure that the code is using types correctly. You should note that since the verifier is a part of the JIT compilers, it kicks in only when a method is being invoked, not when a class or assembly is loaded. You should also note that verification is an optional step because trusted code will never be verified but will be immediately directed to the JIT compiler for compilation.

2.7.3 JIT Compilers

JIT compilers play a major role in the .NET platform because all .NET PE files contain IL and metadata, not native code. The JIT compilers convert IL to native code so that it can execute on the target operating system. For each method that has been successfully verified for type safety, a JIT compiler in the CLR will compile the method and convert it into native code.

One advantage of a JIT compiler is that it can dynamically compile code that is optimized for the target machine. If you take the same .NET PE file from a one-CPU machine to a two-CPU machine, the JIT compiler on the two-CPU machine knows about the second CPU and may be able to spit out the native code that takes advantage of the second CPU. Another obvious advantage is that you can take the same .NET PE file and run it on a totally different platform, whether it be Windows, Unix, or whatever, as long as that platform has a CLR.

For optimization reasons, JIT compilation occurs only the first time a method is invoked. Recall that the class loader adds a stub to each method during class loading. At the first method invocation, the VES reads the information in this stub, which tells it that the code for the method has not been JIT-compiled. At this indication, the JIT compiler compiles the method and injects the address of the native method into this stub. During subsequent invocations to the same method, no JIT compilation is needed because each time the VES goes to read information in the stub, it sees the address of the native method. Because the JIT compiler only performs its magic the first time a method is invoked, the methods you don't need at runtime will never be JIT-compiled.

The compiled, native code lies in memory until the process shuts down and until the garbage collector clears off all references and memory associated with the process. This means that the next time you execute the process or component, the JIT compiler will again perform its magic.

If you want to avoid the cost of JIT compilation at runtime, you can use a special tool called *ngen.exe*, which compiles your IL during installation and setup time. Using *ngen*, you can JIT-compile the code once and cache it on the machine so that you can avoid JIT compilation at runtime (this process is referred to as pre-JITting). In the event that the PE file has been updated, you must PreJIT the PE file again. Otherwise, the CLR can detect the update and dynamically command the appropriate JIT compiler to compile the assembly.

2.7.4 Execution Support and Management

By now, you should see that every component in the CLR that we've covered so far uses metadata and IL in some way to successfully carry out the services that it supports. In addition to the provided metadata and generated managed code, the JIT compiler must generate managed data that the code

manager needs to locate and unwind stack frames.^[12] The *code manager* uses managed data to control the execution of code, including performing stack walks that are required for exception handling, security checks, and garbage collection. Besides the code manager, the CLR also provides a number of important execution-support and management services. A detailed discussion of these services is beyond the scope of this book, so we will briefly enumerate a few of them here:

[12] By the way, you can write a custom JIT compiler or a custom code manager for the CLR because the CLR supports the plug-and-play of these components.

Garbage collection

Unlike C++, where you must delete all heap-based objects manually, the CLR supports automatic lifetime management for all .NET objects. The garbage collector can detect when your objects are no longer being referenced and perform garbage collection to reclaim the unused memory.

Exception handling

Prior to .NET, there was no consistent method for error or exception handling, causing lots of pain in error handling and reporting. In .NET, the CLR supports a standard exception-handling mechanism that works across all languages, allowing every program to use a common error-handling mechanism. The CLR exception-handling mechanism is integrated with Windows Structured Exception Handling (SEH).

Security support

The CLR performs various security checks at runtime to make sure that the code is safe to execute and that the code is not breaching any security requirements. In addition to supporting code access security, the security engine also supports declarative and imperative security checks. Declarative security requires no special security code, but you have to specify the security requirements through attributes or administrative configuration. Imperative security requires that you write the code in your method to specifically cause security checks.

Debugging support

The CLR provides rich support for debugging and profiling. There is an API that compiler vendors can use to develop a debugger. This API contains support for controlling program execution, breakpoints, exceptions, control flow, and so forth. There is also an API for tools to support the profiling of running programs.

Interoperation support

The CLR supports interoperation between the managed (CLR) and unmanaged (no CLR) worlds. The *COM Interop* facility serves as a bridge between COM and the CLR, allowing a COM object to use a .NET object, and vice versa. The *Platform Invoke* (P/Invoke) facility allows you to call Windows API functions.

This is by no means an exhaustive list. The one thing that we want to reiterate is that like the class loader, verifier, JIT compiler, and just about everything else that deals with .NET, these execution-support and management facilities all use metadata, managed code, and managed data in some way to carry out their services.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

2.8 Summary

As you can see from this chapter, the .NET architecture strives to support language integration and componentization in every way that makes sense. Thanks to metadata, programming becomes much easier because you no longer have to worry about the registry for component deployment and other kinks (such as CoCreateInstanceEx, CLSIDs, IIDs, IUnknown, IDL, and so forth) in order to support componentization. Thanks to the CTS, CLS, metadata, and IL, you now have real language integration. Microsoft has shipped a CLR for several flavors of Windows, and has released their shared-source implementation of the CLR that will run on FreeBSD and will no doubt be portable to other Unix-like systems.^[13] Non-Microsoft implementations of the CLR have also appeared, including DotGNU Portable.NET (for more information, see http://www.southern-storm.com.au/portable_net.html) and Mono (see <http://www.go-mono.com>). .NET is thus a multilanguage and multiplatform architecture.

^[13] You can download "The Shared Source CLI 1.0 Release" from Microsoft. The CLR is submitted to ECMA as the Common Language Infrastructure (CLI).

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Chapter 3. .NET Programming

Now that you know what .NET is all about, let's talk about programming for the .NET environment. This chapter presents the common programming model that .NET provides, the core languages and features that .NET supports, and language integration-how you can take advantage of object-oriented features even across different languages that target the CLR.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

3.1 Common Programming Model

Without the .NET Framework, programmers must choose from a wealth of APIs or libraries that support system services. For example, if you want to write GUI applications on Windows, you have a slew of options from which to choose, including the Win32 API, MFC, ATL, VB, and so on. Once you've chosen the library, you have to learn how to use the structures, classes, functions, interfaces, and so forth that the library provides. Unfortunately, this knowledge doesn't transfer directly into a different environment. For instance, there's a big difference between the code to manage IO in MFC and the code to manage IO in VB.

One of the goals of the .NET Framework is to bring commonality to application development by providing a framework of common classes to developers who are using compilers that generate IL. This set of classes, known as the Base Class Library (BCL), is extremely helpful: if you know how to take advantage of IO functionality in .NET using your favorite language, you can easily port that code to another language. This is possible because the namespaces, classes, methods, and so forth are equally accessible in all languages. For example, you can output a line of text to the console the same way across all .NET languages by using the `WriteLine()` method of the `Console` object, as we have seen elsewhere in this book. This consistent framework requires less development training and enables higher programmer productivity.

Since a full discussion of the entire set of classes in the .NET BCL is beyond the scope of this book (see O'Reilly's *In a Nutshell .NET* series), we talk about the `System.Object` class and present the major namespaces in the .NET Framework, opening the doors for you to step into this world.

3.1.1 System.Object

Every type in .NET is an object, meaning that it must derive directly or indirectly from the `Object` class. If you don't specify a base class when you define a class, the compiler will inject this requirement into the IL code. The `Object` class supports a commonality that all .NET classes inherit and, thus, automatically provide to their consumers. The `Object` class exposes the public methods listed in [Table 3-1](#), which you can invoke on any given .NET object at runtime.

Table 3-1. Public methods of the `Object` class

Methods	Description
<code>Equals()</code>	Compares two objects and determines whether they are equivalent (having the same content).
<code>ReferenceEquals()</code>	Compares two object references and determines whether they are referring to the same object in memory.

Methods	Description
GetHashCode()	Gets the object's hash code. Hash codes are used as an added mechanism for determining object uniqueness at runtime. For instance, if you want your objects to be used as keys in a hashtable, you must override this function and provide a unique hash value for each instance of your class.
GetType()	Obtains the object's type at runtime. Once you have obtained the object's type, you can obtain everything about that type using the Reflection API, as explained in Chapter 2 .
ToString()	Gets a string representation of the object. Often used for debugging purposes, this method spits out the fully qualified class name by default.

Examine the following program, which illustrates the use of all these methods:

```
using System;
namespace Cpm
{
    class CPMoDel
    {
        public static void Main( )
        {
            CPMoDel c = new CPMoDel( );

            // Test for self equivalence.
            Console.WriteLine("Equivalence:\t" +
                c.Equals(c)
            );

            // Get the hash code from this object.
            Console.WriteLine("Object hash:\t" +
                c.GetHashCode( )
            );

            // Use the type to obtain method information.
            Console.WriteLine("Object method:\t" +
                c.GetType( ).GetMethods( )[1]
            );

            // Convert the object to a string.
            Console.WriteLine("String representation:\t" +
                c.ToString( )
            );
        }
    }
}
```

If you compile and run this C# program, you get the following output:

```
Equivalence:      True
Object hash:     2
Object method:   Boolean Equals(System.Object)
```

Object dump: Cpm.CPModel

The boldface line displays the second method of the CPModel class. If you look back at the program's code, you'll see that we use the GetType() method to get the type, and then we use the GetMethods() method to retrieve the array of methods supported by this type. From this array, we pull off the second method, which happens to be Equals(), a method that's implemented by System.Object.

As you can see, the System.Object class provides a mechanism for runtime type identification, equivalence, and inspection for all .NET objects.

3.1.2 Major Namespaces

[Table 3-2](#) is a short list of important namespaces and classes in the .NET Framework that provide support for almost any application that you will develop. These are the namespaces that you'll find yourself using again and again the more you develop .NET applications. For more information, consult MSDN Online or your SDK documentation, as a detailed discussion of these namespaces and classes is beyond the scope of this book.

Table 3-2. Important .NET namespaces and classes

Namespace	Description
System	Includes basic classes almost every program will use. Some simple classes that belong in this namespace are Object, Char, String, Array, and Exception. This namespace also includes more advanced classes such as GC and AppDomain.
System.IO	Provides a set of classes to support synchronous and asynchronous IO manipulation for data streams. Also provides classes that allow you to manipulate the file system, such as creating, managing, and deleting files and directories. Some of these classes are FileStream, MemoryStream, Path, and Directory.
System.Collections	Includes a set of classes that allow you to manage collections of objects. Some of these classes are ArrayList, DictionaryBase, Hashtable, Queue, and Stack.
System.Threading	Includes a set of classes that support multithreaded programming. Some of these classes are Thread, ThreadPool, Mutex, and AutoResetEvent.
System.Reflection	Includes a set of classes that support dynamic binding and type inspection. Some of these classes are Assembly, Module, and MethodInfo.
System.Security	Includes a set of classes and child namespaces that provide security support. The interesting child namespaces include Cryptography, Permissions, Policy, and Principal.

Namespace	Description
System.Net	Includes a set of classes and child namespaces that provide support for network programming. Some of these classes are IPAddress, Dns, and HttpWebRequest.
System.Data	Contains classes for ADO.NET. See Chapter 5 .
System.Web.Services	Contains classes for XML web services. See Chapter 6 .
System.Web.UI	Contains classes for ASP.NET web pages. See Chapter 7 .
System.Windows.Forms	Contains classes for Windows user interface applications. See Chapter 8 .

Keep in mind that if you know how to use any of the classes in these namespaces, you can write the code to take advantage of them in any language that targets the CLR, because the class and method names remain consistent across all .NET languages.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

3.2 Core Features and Languages

Since one of .NET's goals is to support a common paradigm for application programming, it must specify and utilize programming concepts consistently. In this section, we will examine four core Microsoft .NET languages, including Managed C++, VB.NET, C#, and J#, and several core programming concepts that all .NET languages support, including:

Namespace

Mitigates name collisions.

Interface

Specifies the methods and properties that must be implemented by objects that expose the interface.

Encapsulation

In object-oriented languages, allows a class to combine all its data and behavior.

Inheritance

Allows a class to inherit from a parent class so that it can reuse rich functionality that the parent class has implemented, thus reducing development effort and programming errors.

Polymorphism

Permits developers to specify or implement behaviors in a base class that can be overridden by a derived class. This is a very powerful feature because it allows developers to select the correct behavior based on the referenced runtime object.

Exception handling

Allows us to write easier-to-understand code because it allows us to capture all errors in a common, understandable pattern—totally opposite to that of nine levels of nested conditional blocks.

Although this is not a complete list of concepts that .NET supports, it includes all the major .NET concepts that we want to cover in this section. We will show you examples of all these features in Managed C++, VB.NET, C#, and J#. These concepts are nothing new: we're merely demonstrating how they're represented in all core Microsoft .NET languages.

Before we start, you should understand first what our examples will accomplish. First, we will create a namespace, called `Lang`, that encapsulates an interface, `ISteering`. Then we will create two classes: `Vehicle`, which is an abstract base class that implements `ISteering`, and `Car`, which is a derivative of `Vehicle`. We will support an entry point that instantiates and uses `Car` within a `try` block. We will unveil other details as we work through the examples.

3.2.1 Managed C++ Code

Managed C++ is Microsoft's implementation of the C++ programming language with some newly

added keywords and features to support .NET programming. This allows you to use C++ to develop managed objects, which are objects that run in the CLR. Using Managed C++, you can obtain the performance^[1] that is inherent in C++ programs, and at the same time, you can also take advantage of CLR features.^[2]

[1] You can easily mix managed and unmanaged code in C++ programs. The unmanaged code will perform better. See this chapter's example code, which you can download from <http://www.oreilly.com/catalog/dotnetfrmess3/>.

[2] However, if you look carefully at the features and new keywords (`_ _abstract`, `_ _box`, `_ _delegate`, `_ _gc`, `_ _nogc`, `_ _pin`, etc.) that have been added to Microsoft C++, we doubt that you'll want to use Managed C++ to write new code for the CLR, especially when you have C#.

Now let's look at an example that includes all the concepts we want to examine. As you can see in the following code listing, we start off creating a new namespace, `Lang`, which envelops everything except `main()`. With the exception of the first line and special keywords, the code listing conforms perfectly to the C++ standard:

```
#using <mscorlib.dll>
using namespace System;
```

```
namespace Lang
{
```

Next, we specify an interface, called `ISteering`. If you are a C++ programmer, you will immediately notice that there are two new keywords in the following code listing, `_ _gc` and `_ _interface`. The new keyword `_ _interface` allows you to declare an interface, which is basically equivalent to an abstract base class in C++. In other words, the two method prototypes are specified, but not implemented here. The class that implements this interface provides the implementation for these methods:

```
_ _gc _ _interface ISteering
{
    void TurnLeft( );
    void TurnRight( );
};
```

If you are a COM programmer, you know that in COM you have to manage the lifetimes of your objects and components yourself. Even worse, you also have to rely on your clients to negotiate and interoperate correctly with your COM components; otherwise, extant references will never be reclaimed. Managed C++ removes this problem by adding a new keyword, `_ _gc` that tells the CLR to garbage-collect the references to your interface when they are no longer in use. Aside from these two keywords, the previous code listing requires no other explanation for programmers who have experience with C-like languages.

Now that we have an interface, let's implement it. The following code listing is a Managed C++ class (as indicated by the `_ _gc`) that implements our `ISteering` interface. One thing to notice is that this class is an abstract base class because the `ApplyBrakes()` method is a pure virtual function, as indicated by the `=0` syntax. `Vehicle` doesn't provide the implementation for this method, but its derived class must supply the implementation:

```
_ _gc class Vehicle : public ISteering
{
    public:
```

```

void TurnLeft( )
{
    Console::WriteLine("Vehicle turns left.");
}

void TurnRight( )
{
    Console::WriteLine("Vehicle turns right.");
}

virtual void ApplyBrakes( ) = 0;
};

```

Since Vehicle is an abstract base class and can't be instantiated, we need to provide a Vehicle derivative, which we will call Car. As you can see in the following listing, everything about the class is C++, with the exception of the keyword `_ _gc`. Note that the `ApplyBrakes()` function first dumps a text message to the console and then immediately creates and throws an exception, notifying an exception handler that there has been a brake failure.

```

_ _gc class Car : public Vehicle
{
public:

void ApplyBrakes( )
{
    Console::WriteLine("Car trying to stop.");

    throw new Exception ("Brake failure!");
}
};

} // This brace ends the Lang namespace.

```

What is special here is that the Exception class is a part of the .NET Framework, specifically belonging to the System namespace. This is great because this class works exactly the same way in all languages and there's no longer a need to invent your own exception hierarchy.

Two important types that derive from Exception are SystemException (framework and Windows errors) and ApplicationException (your application's errors). If you create your own exception classes, they should derive from ApplicationException. .NET conventions also suggest throwing ApplicationException objects for application-specific errors.

Now that we have a concrete class, we can write the `main()` function to test our Car class. Notice that we have added a `try` block that encapsulates the bulk of our code so that we can handle any exceptions in the `catch` block. Looking carefully at the following code listing, you'll see that we've instantiated a new Car on the managed heap, but we've actually referred to this Car instance using a Vehicle pointer. Next, we tell the vehicle to `TurnLeft()`-there's no surprise here because we've implemented this method in Vehicle. However, in the following statement, we tell the Vehicle that we're applying the brakes, but `ApplyBrakes()` is not implemented in Vehicle. Since this is a virtual method, the correct `vptr` and `vtbl`^[3] will be used, resulting in a call to `Car::ApplyBrakes()`. Of

course `Car::ApplyBrakes()` will throw an exception, putting us into the catch block. Inside the catch block, we convert the caught exception into a string and dump it out to the console.

[3] Many C++ compilers use `vtbls` (a `vtbl` is a table of function pointers) and `vptrs` (a `vptr` is a pointer to the `vtbl`) to support dynamic binding or polymorphism.

We can do this because `Exception` is a class in the .NET Framework and all classes in the framework must derive from `System.Object`, which implements a rudimentary `ToString()` function to convert any object into a string:

```
void main( )
{
    try
    {
        Lang::Vehicle *pV = 0;    // Namespace qualifier
        pV = new Lang::Car( );    // pV refers to a car
        pV->TurnLeft( );         // Interface usage
        pV->ApplyBrakes( );       // Polymorphism in action
    }
    catch(Exception *pe)
    {
        Console::WriteLine(pe->ToString( ));
    }
}
```

Notice that you don't have to deallocate your objects on the managed heap when you've finished using them, because the garbage collector will do that for you in .NET.

Although this is a simple example, we have used Managed C++ to illustrate all major object-oriented programming concepts, including namespaces, interfaces, encapsulation, inheritance, polymorphism, and exception handling. Next, we demonstrate that you can translate this code into any other .NET language because they all support these concepts. Specifically, we'll show you this same example in VB.NET, C#, J#, and IL, just to prove that these concepts can be represented the same way in all languages that targets the CLR.

3.2.2 VB.NET Code

Microsoft has revamped VB and added full features for object-oriented programming. The new VB language, Visual Basic .NET (or VB.NET), allows you to do all that you can do with VB, albeit much more easily. If you are a VB programmer with knowledge of other object-oriented languages, such as C++ or Smalltalk, then you will love the new syntax that comes along with VB.NET. If you are a VB programmer without knowledge of other object-oriented languages, you will be surprised by the new VB.NET syntax at first, but you will realize that the new syntax simplifies your life as a programmer.^[4]

[4] To learn more about VB.NET, see O'Reilly's *VB.NET Language in a Nutshell*, Second Edition, by Steven Roman, PhD., Ron Petruscha, and Paul Lomax, or *Programming Visual Basic .NET*, Second Edition, by Jesse Liberty.

In addition to the VB-style Rapid Application Development (RAD) support, VB.NET is a modernized language that gives you full access to the .NET Framework. The VB.NET compiler generates metadata and IL code, making the language an equal citizen to that of C# or Managed C++. Unlike

VB versions prior to VB6, there will be no interpreter in VB.NET, so there should be no violent arguments about performance drawbacks of VB versus another language.

Perhaps the most potent feature is that now you can write interfaces and classes that look very similar to those written in other .NET languages. The new syntax allows you to inherit from base classes, implement interfaces, override virtual functions, create an abstract base class, and so forth. In addition, it also supports exception handling exactly as does C# and Managed C++, making error handling much easier. Finally, VB.NET ships with a command-line compiler, *vbc.exe*, introduced in [Chapter 2](#).

Let's see how to translate the previous Managed C++ program into VB.NET so that you can see the striking conceptual resemblance. First, we'll start by defining a namespace called Lang, shown here in bold:

```
Imports System
```

```
Namespace Lang
```

Next, we specify the ISteering interface, which is easy to do in VB.NET since the syntax is very straightforward, especially when you compare it with Managed C++. In the following code listing, you'll notice that instead of using opening and closing braces as in Managed C++, you start the interface definition by using the appropriate VB.NET keyword, **Interface**, and end it by prefixing the associated keyword with the word **End**. This is just normal VB-style syntax and shouldn't surprise any VB programmer:

```
Interface ISteering
```

```
    Sub TurnLeft( )
```

```
    Sub TurnRight( )
```

```
End Interface
```

With our interface specified, we can now implement it. Since our Vehicle class is an abstract base class, we must add the **MustInherit** keyword when we define it, explicitly telling the VB.NET compiler that this class cannot be instantiated. In VB.NET, the **Class** keyword allows you to define a class, and the **Implements** keyword allows you implement an interface. Another thing that you should be aware of is that `ApplyBrakes()` is not implemented in this class, and we have appropriately signaled this to the VB.NET compiler by using the **MustOverride** keyword:

```
MustInherit Class Vehicle
```

```
    Implements ISteering
```

```
    Public Sub TurnLeft( ) Implements ISteering.TurnLeft
```

```
        Console.WriteLine("Vehicle turns left.")
```

```
    End Sub
```

```
    Public Sub TurnRight( ) Implements ISteering.TurnRight
```

```
        Console.WriteLine("Vehicle turn right.")
```

```
    End Sub
```

```
    Public MustOverride Sub ApplyBrakes( )
```

```
End Class
```

As far as language differences go, you must explicitly describe the access (i.e., **public**, **private**, and

so forth) for each method separately. This is different from C++ because all members take on the previously defined access type.

Now we are ready to translate the concrete Car class. In VB.NET, you can derive from a base class by using the `Inherits` keyword, as shown in the following code. Since we have said that `ApplyBrakes()` must be overridden, we provide its implementation here. Again, notice that we're throwing an exception:

```
Class Car
    Inherits Vehicle

    Public Overrides Sub ApplyBrakes( )
        Console.WriteLine("Car trying to stop.")
        throw new Exception("Brake failure!")
    End Sub

End Class
```

```
End Namespace
```

Now that we have all the pieces in place, let's define a module with an entry point, `Main()`, that the CLR will execute. In `Main()`, you'll notice that we're handling exceptions exactly as we did in the Managed C++ example. You should also note that this code demonstrates the use of polymorphism because we first create a `Vehicle` reference that refers to a `Car` object at runtime. We tell the `Vehicle` to `ApplyBrakes()`, but since the `Vehicle` happens to be referring to a `Car`, the object that is stopping is the target `Car` object:

```
Public Module Driver

    Sub Main( )

        Try

            Dim v As Lang.Vehicle ' namespace qualifier
            v = New Lang.Car      ' v refers to a car
            v.TurnLeft( )        ' interface usage
            v.ApplyBrakes( )     ' polymorphism in action

        Catch e As Exception

            Console.WriteLine(e.ToString( ))

        End Try

    End Sub

End Module
```

This simple program demonstrates that we can take advantage of .NET object-oriented features using VB.NET. Having seen this example, you should see that VB.NET is very object oriented, with features that map directly to those of Managed C++ and other .NET languages.

3.2.3 C# Code

As you've just seen, VB.NET is a breeze compared to Managed C++, but VB.NET is not the only simple language in .NET-C# is also amazingly simple. Developed from the ground up, C# supports all the object-oriented features in .NET. It maps so closely to the Java and C++ languages that if you have experience with either of these languages, you can pick up C# and be productive with it immediately.

Microsoft has developed many tools using C#; in fact, most of the components in Visual Studio .NET and the .NET class libraries were developed using C#. Microsoft is using C# extensively, and we think that C# is here to stay.^[5]

[5] To learn more about C#, check out O'Reilly's *C# Essentials*, Second Edition, by Ben Albahari, Peter Drayton, and Brad Merrill; the forthcoming *C# in a Nutshell*, Second Edition, by Peter Drayton, Ben Albahari, and Ted Neward; and *Programming C#*, Third Edition, by Jesse Liberty.

Having said that, let's translate our previous program into C# and illustrate all the features we want to see. Again, we start by defining a namespace. As you can see, the syntax for C# maps really closely to that of Managed C++:

```
using System;
```

```
namespace Lang
{
```

Following is the ISteering interface specification in C#. Since C# was developed from scratch, we don't need to add any funny keywords like `__gc` and `__interface`, as we did in the Managed C++ version of this program:

```
interface ISteering
{
    void TurnLeft( );
    void TurnRight( );
}
```

Having defined our interface, we can now implement it in the abstract Vehicle class. Unlike Managed C++ but similar to VB.NET, C# requires that you explicitly notify the C# compiler that the Vehicle class is an abstract base class by using the `abstract` keyword. Since `ApplyBrakes()` is an abstract method-meaning that this class doesn't supply its implementation-you must make the class abstract, otherwise the C# compiler will barf at you. Put another way, you must explicitly signal to the C# compiler the features you want, including `abstract`, `public`, `private`, and so forth, each time you define a class, method, property, and so on:

```
abstract class Vehicle : ISteering
{
    public void TurnLeft( )
    {
        Console.WriteLine("Vehicle turns left.");
    }

    public void TurnRight( )
    {
        Console.WriteLine("Vehicle turn right.");
    }
}
```

```

    }

    public abstract void ApplyBrakes( );
}

```

Here's our Car class that derives from Vehicle and overrides the ApplyBrakes() method declared in Vehicle. Note that we are explicitly telling the C# compiler that we are indeed overriding a method previously specified in the inheritance chain. You must add the override modifier, or ApplyBrakes() will hide the one in the parent class. Otherwise, we are also throwing the same exception as before:

```

class Car : Vehicle
{
    public override void ApplyBrakes( )
    {
        Console.WriteLine("Car trying to stop.");
        throw new Exception("Brake failure!");
    }
}

} // This brace ends the Lang namespace.

```

Finally, here's a class that encapsulates an entry point for the CLR to invoke. If you look at this code carefully, you'll see that it maps directly to the code in both Managed C++ and VB.NET:

```

class Drive
{
    public static void Main( )
    {
        try
        {
            Lang.Vehicle v = null;    // Namespace qualifier
            v = new Lang.Car( );      // v refers to a car
            v.TurnLeft( );            // Interface usage
            v.ApplyBrakes( );         // Polymorphism in action
        }
        catch(Exception e)
        {
            Console.WriteLine(e.ToString( ));
        }
    }
}

```

There are two other interesting things to note about C#. First, unlike C++ but similar to Java, C# doesn't use header files.[\[6\]](#)

^[6] If you've never used C++, a header file is optional and usually contains class and type declarations. The implementation for these classes is usually stored in source files.

Second, the C# compiler generates XML documentation for you if you use XML comments in your code. To take advantage of this feature, start your XML comments with three slashes, as in the following examples:

```

///

```

```

/// <remarks>
///   This class is an abstract class that must be
///   overridden by derived classes.
/// </remarks>
abstract class Vehicle : ISteering
{
    /// <summary>Add juice to the vehicle.</summary>
    /// <param name="gallons">
    ///   Number of gallons added.
    /// </param>
    /// <return>Whether the tank is full.</return>
    public bool FillUp(int gallons)
    {
        return true;
    }
}

```

These are simple examples using the predefined tags that the C# compiler understands. You can also use your own XML tags in XML comments, as long as your resulting XML is well formed. Given that you have a source code file with XML comments, you can automatically generate an XML-formatted reference document by using the C# compiler's `/doc:` option, as follows:

```
csc /doc:doc.xml mylangdoc.cs
```

Although we didn't specify the types of our parameters in the XML comments shown previously, the C# compiler will detect the correct types and add the fully qualified types into the generated XML document. For example, the following generated XML listing corresponds to the XML comments for the `FillUp()` method. Notice that the C# compiler added `System.Int32` into the generated XML document:

```

<member name="M:Lang.Vehicle.FillUp(System.Int32)">
  <summary>Add juice to the vehicle.</summary>
  <param name="gallons">
    Number of gallons added.
  </param>
  <return>Whether the tank is full.</return>
</member>

```

Now that you have the generated XML document, you can write your own XSL document to translate the XML into any visual representation you prefer.

3.2.4 J# Code

Shipped with .NET Framework 1.1 (and thus with Visual Studio .NET 2003), J# is a Java language that targets the CLR. For completeness, here's the same program in J#, demonstrating that J# also supports the same object-oriented features that we've been illustrating. We simply took the preceding C# program and made a few minor changes, resulting in the J# program that we are about to examine.

Let's first look at the namespace declaration. Instead of using the keyword `namespace`, Java uses the keyword `package`, which is conceptually equivalent to the namespace concept we've been observing,

since the purpose of a package is to prevent name conflicts:

```
package Lang;
```

```
import System.Console;
```

The interface specification for `ISteering` in J# looks exactly equivalent to the one written in C#:

```
interface ISteering
{
    void TurnLeft( );
    void TurnRight( );
}
```

For the `Vehicle` class, there are two changes, which are shown in bold. First, the keyword `implements` is used to declare that a class implements one or more interfaces. Second, since Java requires thrown exceptions to be explicitly declared within the method signature, we've added this declaration in the `ApplyBrakes()` method:

```
abstract class Vehicle implements ISteering
{
    public void TurnLeft( )
    {
        Console.WriteLine("Vehicle turns left.");
    }

    public void TurnRight( )
    {
        Console.WriteLine("Vehicle turn right.");
    }

    public abstract void ApplyBrakes( ) throws Exception;
}
```

There are also two changes for the `Car` class, which are shown in bold. The `extends` keyword is used to declare that a class derives from (or extends) another class. The declaration for `ApplyBrakes()` must match it's parents signature, so we've explicitly indicated that an exception may be thrown from this method, as shown in bold:

```
// extends - used to derive from a base class.
class Car extends Vehicle
{
    public void ApplyBrakes( ) throws Exception
    {
        Console.WriteLine("Car trying to stop.");
        throw new Exception("Brake failure!");
    }
}
```

Finally, we've made one minor change in the `Drive` class: we simply changed `Main()` to `main()`, as required by J#:

```
class Drive
```

```

{
    public static void main( )
    {
        try
        {
            Lang.Vehicle v = null;    // Namespace qualifer
            v = new Lang.Car( );    // v refers to a car
            v.TurnLeft( );    // Interface usage
            v.ApplyBrakes( );    // Polymorphism in action
        }
        catch(Exception e)
        {
            Console.WriteLine(e.ToString( ));
        }
    }
}

```

Like C#, J# supports all the object-oriented concepts we've been studying. Also, J# and C# are syntactically very similar.

3.2.5 Intermediate Language (IL) Code

Since all languages compile to IL, let's examine the IL code for the program that we've been studying. As explained in [Chapter 2](#), IL is a set of stack-based instructions that supports an exhaustive list of popular object-oriented features, including the ones that we've already examined in this chapter. It is an intermediary step, gluing .NET applications to the CLR.

Let's start by looking at the namespace declaration. Notice the `.namespace` IL declaration allows us to create our Lang namespace. Similar to C#, IL uses opening and closing braces:

```

.namespace Lang
{

```

Now for the ISteering interface. In IL, any type that is to be managed by the CLR must be declared using the `.class` IL declaration. Since the CLR must manage the references to an interface, you must use the `.class` IL declaration to specify an interface in IL, as shown in the following code listing:

```

.class interface private abstract auto ansi ISteering
{
    .method public hidebysig newslot virtual abstract
        instance void TurnLeft( ) cil managed
    {
    } // End of method ISteering::TurnLeft

    .method public hidebysig newslot virtual abstract
        instance void TurnRight( ) cil managed
    {
    } // End of method ISteering::TurnRight
} // End of class ISteering

```


In addition, you must insert two special IL attributes:

`interface`

Signals that the current type definition is an interface specification.

`abstract`

Signals that there will be no method implementations in this definition and that the implementer of this interface must provide the method implementations for all methods defined in this interface.

Other attributes shown in this definition that aren't necessarily needed to specify an interface in IL include the following:

`private`

Because we haven't provided the visibility of our interface definition in C#, the generated IL code shown here adds the `private` IL attribute to this interface definition. This means that this particular interface is visible only within the current assembly and no other external assembly can see it.

`auto`

Tells the CLR to perform automatic layout of this type at runtime.

`ansi`

Tells the CLR to use ANSI string buffers to marshal data across managed and unmanaged boundaries.

Now you know how to specify an interface in IL. Before we proceed further, let's briefly look at the attributes in the `.method` declarations—at least the attributes that we haven't examined, including:

`newslot`

Tells the JIT compiler to reserve a new slot in the type's `svtbl`, which will be used by the CLR at runtime to resolve virtual-method invocations.

`instance`

Tells the CLR that this method is an instance or object-level method, as opposed to a static or class-level method.

Having specified the `ISteering` interface in IL, let's implement it in our `Vehicle` class. As you can see in the following code fragment, there's no surprise. We extend the `System.Object` class (indicated by the `extends` keyword) and implement `Lang.ISteering` (as indicated by the `implements` keyword):

```
.class private abstract auto ansi beforefieldinit Vehicle
    extends [mscorlib]System.Object
    implements Lang.ISteering
{
    .method public hidebysig newslot final virtual
        instance void TurnLeft( ) cil managed
    {
        // IL code omitted for clarity
    } // End of method Vehicle::TurnLeft
```

```

.method public hidebysig newslot final virtual
        instance void TurnRight( ) cil managed
{
    // IL code omitted for clarity
} // End of method Vehicle::TurnRight

.method public hidebysig newslot virtual abstract
        instance void ApplyBrakes( ) cil managed
{
} // End of method Vehicle::ApplyBrakes

// .ctor omitted for clarity

} // End of class Vehicle

```

Notice also that this class is an abstract class and that the `ApplyBrakes()` method is an abstract method, similar to what we've seen in the previous examples. Another thing to note is the `final` IL attribute in the `.method` declarations for both `TurnLeft()` and `TurnRight()`. This IL attribute specifies that these methods can no longer be overridden by subclasses of `Vehicle`. Having seen all these attributes, you should realize that everything in IL is explicitly declared so that all components of the CLR can take advantage of this information to manage your types at runtime.

Now let's look at the `Car` class that derives from the `Vehicle` class. You'll notice that in the `ApplyBrakes()` method implementation, the `newobj instance` IL instruction creates a new instance of the `Exception` class. Next, the `throw` IL instruction immediately raises the exception object just created:

```

.class private auto ansi beforefieldinit Car
    extends Lang.Vehicle
{
    .method public hidebysig virtual instance void
        ApplyBrakes( ) cil managed
    {
        // IL code omitted for clarity
        newobj instance void
            [mscorlib]System.Exception::.ctor(class System.String)
        throw

    } // End of method Car::ApplyBrakes

    // .ctor omitted for clarity

} // End of class Car

} // End of namespace Lang

```

Finally, let's look at our `Main()` function, which is part of the `Drive` class. We've removed most of the IL code—which you've already learned—from this function to make the following code easier to read, but we've kept the important elements that must be examined. First, the `.locals` directive identifies all the local variables for the `Main()` function. Second, you can see that IL also supports exception handling through the `.try` instruction. In both the `.try` and `catch` blocks, notice that there is a `leave.s` instruction that forces execution to jump to the IL instruction on line `IL_0024`, thus leaving

both the `.try` and `catch` blocks:

```
.class private auto ansi beforefieldinit Drive
    extends [mscorlib]System.Object
{
    .method public hidebysig static void Main( ) cil managed
    {
        .entrypoint
        // Code size          37 (0x25)
        .maxstack 1

        .locals (class Lang.Vehicle V_0 ,
                class [mscorlib]System.Exception V_1)
        .try
        {
            // IL code omitted for clarity
            leave.s    IL_0024
        } // End .try
        catch [mscorlib]System.Exception
        {
            // IL code omitted for clarity
            leave.s    IL_0024
        } // End handler
        IL_0024: ret
    } // End of method Drive::Main

    // .ctor omitted for clarity
} // End of class Drive
```

As you can see, all the major concepts that we've examined apply intrinsically to IL. Since you've seen Managed C++, VB.NET, C#, J#, and IL code that support these features, we won't attempt to further convince you that all these features work in other languages that target the CLR.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

3.3 Language Integration

In the previous section, we saw that you can take advantage of .NET object-oriented concepts in any .NET language. In this section, we show that you can take advantage of *language integration*—the ability to derive a class from a base that is specified in a totally different language; to catch exceptions thrown by code written in a different language; or to take advantage of polymorphism across different languages, and so forth.

Before we discuss the examples in this section, let's first understand what we want to accomplish (see [Figure 3-1](#)). We will first use Managed C++ to develop a Vehicle class that is an abstract base class. The Vehicle class exposes three polymorphic methods, including TurnLeft(), TurnRight(), and ApplyBrakes(). We will then use VB.NET to develop a Car class that derives from Vehicle and overrides these three virtual methods. In addition, we will use C# to develop the Plane class that derives from Vehicle and overrides these three virtual methods.

Figure 3-1. Polymorphism across languages

In the upcoming code example, we can tell a Vehicle to TurnLeft() or TurnRight(), but what turns left or right depends on the target object, whether a Car or a Plane. Unlike the examples in the previous section, the examples here illustrate that we can inherit classes and call virtual functions from ones that are defined in another language. In addition, we will demonstrate in our test program that exception handling works across different languages.

3.3.1 Vehicle Class in Managed C++

Let's use Managed C++ to develop the Vehicle class, which is an abstract base class because ApplyBrakes() is a pure virtual function. Vehicle implements the ISteering interface to support turning left and turning right. Since the ApplyBrakes() function is a pure virtual function, any concrete derivative of Vehicle must implement this method:

```
#using <mscorlib.dll>
```

```

using namespace System;

public _ _gc _ _interface ISteering
{
    void TurnLeft( );
    void TurnRight( );
};

public _ _gc class Vehicle : public ISteering
{
    public:

        virtual void TurnLeft( )
        {
            Console::WriteLine("Vehicle turns left.");
        }

        virtual void TurnRight( )
        {
            Console::WriteLine("Vehicle turns right.");
        }

        virtual void ApplyBrakes( ) = 0;
};

```

Given this abstract base class, we can create a DLL that hosts this definition. The first command here shows how we use the Managed C++ compiler to compile (as indicated by the `/c` option) the `vehicle.cpp` file, which contains the previous code. The second command shows how we use the C++ linker to create a DLL with metadata and IL code:

```

cl /CLR /c vehicle.cpp
link -dll /out:vehicle.dll vehicle.obj

```

Given just a few lines of Managed C++ code, we can build a DLL that can be used by another component. Note that there is no need to provide code for `IUnknown`, `DllGetClassObject()`, `DllCanUnloadNow()`, `DllRegisterServer()`, `DllUnregisterServer()`, and so forth. In the old days, you had to provide code for these functions and interfaces for legacy COM DLLs.

3.3.2 Car Class in VB.NET

Given this abstract Vehicle class, the Car class can derive from it and provide the implementation for the three virtual methods defined by Vehicle. In the following code, note that we've overridden and provided the implementation for `TurnLeft()`, `TurnRight()`, and `ApplyBrakes()`. The `ApplyBrakes()` method is special in that it throws an exception, which will be caught by code written in J#, as we'll see later:

```

Imports System

Public Class Car
    Inherits Vehicle

```

```

Overrides Public Sub TurnLeft( )
    Console.WriteLine("Car turns left.")
End Sub

Overrides Public Sub TurnRight( )
    Console.WriteLine("Car turns right.")
End Sub

Overrides Public Sub ApplyBrakes( )
    Console.WriteLine("Car trying to stop.")
    throw new Exception("Brake failure!")
End Sub

End Class

```

With this code, we can build a DLL using the command-line VB.NET compiler, as follows:

```
vbc /r:vehicle.dll /t:library /out:car.dll car.vb
```

Since we want the VB.NET compiler to generate a DLL, we must signal this by using the `/t:library` option. Also, since Car derives from Vehicle, the VB.NET compiler must resolve the references to Vehicle. We can tell the VB.NET compiler the location of external references using the `/r:` option. It is important to note that you don't need to have the source code for the vehicle DLL to reuse its code because all type information can be obtained from any .NET assembly. In addition, you should note that from this example, we have proven that you can derive a VB.NET class from a Managed C++ class.

3.3.3 Plane Class in C#

Now let's use C# to develop the Plane class, which derives from the Vehicle class written in Managed C++. Similar to the Car class, the Plane class implements the three virtual functions from the Vehicle class. Unlike the Car class, though, the `ApplyBrakes()` method of this class doesn't throw an exception:

```

using System;

public class Plane : Vehicle
{
    override public void TurnLeft( )
    {
        Console.WriteLine("Plane turns left.");
    }

    override public void TurnRight( )
    {
        Console.WriteLine("Plane turns right.");
    }

    override public void ApplyBrakes( )
    {
        Console.WriteLine("Air brakes being used.");
    }
}

```

```

    }
}

```

You can build a DLL from this code using the following command:

```
csc /r:vehicle.dll /t:library /out:plane.dll plane.cs
```

Notice that we have used the `/r:` option to tell the C# compiler that `Vehicle` is defined in `vehicle.dll`.

3.3.4 Test Driver in J#

Having developed `vehicle.dll`, `car.dll`, and `plane.dll`, we are now ready to demonstrate that polymorphism and exception handling work across different languages. Written in J#, the next code listing contains a `main()` method with a `Vehicle` reference and an exception handler.

Inside the `try` block, we first instantiate a `Plane` class and refer to this instance using the local `Vehicle` reference. Instead of telling the `Plane` to `TurnLeft()` or `ApplyBrakes()`, we tell the `Vehicle` to do so. Similarly, we instantiate a `Car` and refer to this instance using the local `Vehicle` reference. Again, instead of telling the `Car` to `TurnLeft()` or `ApplyBrakes()`, we tell the `Vehicle` to do so. In both cases, we tell the `Vehicle` either to `TurnLeft()` or `ApplyBrakes()`, but the actual vehicle that employs `TurnLeft()` or `ApplyBrakes()` is the `Plane` instance in the first case and the `Car` instance in the second case; that's polymorphism, and it works across languages.

You should note that the second call to `ApplyBrakes()` would cause an exception because we threw an exception from `Car's ApplyBrakes()`. Although `Car's ApplyBrakes()` was written using VB.NET, we could still catch the exception that it's throwing in J#, proving that exception handling works across languages:

```

class TestDrive
{
    public static void main( )
    {
        Vehicle v = null; // Vehicle reference

        try
        {
            Plane p = new Plane( );
            v = p;
            v.TurnLeft( );
            v.ApplyBrakes( );

            Car c = new Car( );
            v = c;
            v.TurnLeft( );
            v.ApplyBrakes( ); // Exception
        }
        catch(System.Exception e)
        {
            System.Console.WriteLine(e.ToString( ));
        }
    }
}

```

```
}  
}
```

If you want to test out these features, you can create an EXE using the following command:

```
vjc /r:vehicle.dll;car.dll;plane.dll /t:exe /out:drive.exe drive.jsl
```

Since we have used the Vehicle, Car, and Plane classes in this code, we must include references to *vehicle.dll*, *car.dll*, and *plane.dll*. And since we are building an EXE, we need to signal this to the J# compiler using the `/t:exe` option. Once you have built this EXE and executed it, you get the following output:

```
Plane turns left.  
Air brakes being used.  
Car turns left.  
Car trying to stop.  
System.Exception: Brake failure!  
   at Car.ApplyBrakes( )  
   at TestDrive.main( )
```

As expected, the plane first turns left and then uses its air brakes. Then the car turns left, tries to stop, but can't, so it throws an exception that is caught in the `main()` method.

In this simple example, we have shown that you can now take advantage of inheritance, polymorphism, and exception handling across different languages that target the CLR.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

3.4 Summary

We started this chapter by telling you that .NET provides a common programming model, which reduces the learning curve and increases productivity. Once you've learned how to do something using the classes in the .NET Framework, this knowledge will transfer to any .NET language. We then illustrated that we could write the same type of code, supporting major .NET features, in any given language that targets the CLR. Finally, we proved to you that .NET indeed supports language integration, which was never possible using Microsoft platforms and tools, prior to .NET.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Chapter 4. Working with .NET Components

Having seen the language-integration examples in the previous chapter, you now know that all .NET assemblies are essentially binary components.^[1] You can treat each .NET assembly as a component that you can plug into another component or application, without the need for source code, since all the metadata for the component is stored inside the .NET assembly. While you have to write a ton of plumbing code to build a component in COM, creating a component in .NET involves no extra work, as all .NET assemblies are components by nature.

^[1] Remember, as we explained in [Chapter 1](#), we're using the term "component" as a binary, deployable unit, not as a COM class.

In this chapter, we examine the more advanced topics, including component deployment, distributed components, and enterprise services, such as transaction management, object pooling, role-based security, and message queuing.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

4.1 Deployment Options

For a simple program like *hello.exe* that we built in [Chapter 2](#), deployment is easy: copy the assembly into a directory, and it's ready to run. When you want to uninstall it, remove the file from the directory. However, when you want to share components with other applications, you've got to do some work.

In COM, you must store activation and marshaling^[2] information in the registry for components to interoperate; as a result, any COM developer can discuss at length the pain and suffering inherent in COM and the system registry. In .NET, the system registry is no longer necessary for component integration.

[2] Distributed applications require a communication layer to assemble and disassemble application data and network streams. This layer is formally known as a marshaler in Microsoft terminology. Assembling and disassembling an application-level protocol network buffer are formally known as marshaling and unmarshaling, respectively.

In the .NET environment, components can be *private*, meaning that they are unpublished and used by known clients, or *shared*, meaning that they are published and can be used by any clients. This section discusses several options for deploying private and shared components.

4.1.1 Private Components

If you have private components that are used only by specific clients, you have two deployment options. You can store the private components and the clients that use these components in the same directory, or you can store the components in a component-specific directory that the client can access. Since these clients use the exact private components that they referenced at build time, the CLR doesn't support version checking or enforce version policies on private components.

To install your applications in either of these cases, perform a simple *xcopy* of your application files from the source installation directory to the destination directory. When you want to remove the application, remove these directories. You don't have to write code to store information into the registry, so there's no worrying about whether you've missed inserting a registry setting for correct application execution. In addition, because nothing is stored in the registry, you don't have to worry about registry residues.

4.1.1.1 One-directory deployment

For simplicity, you can place supporting assemblies in the same directory as the client application. For example, in [Chapter 3](#), we placed the *vehicle*, *car*, and *plane* assemblies in the same directory as the client application, *drive.exe*. Since both the client application and supporting assemblies are stored within the same directory, the CLR has no problem resolving this reference at runtime (i.e., find and load *plane.dll* and activate the *Plane* class). If you move any of the DLLs to a different directory (e.g. put *vehicle.dll* in *c:\temp*), you will get an exception when you execute *drive.exe*. This is because the CLR looks for the *vehicle* assembly in the following order:

1. It looks for a file called *vehicle.dll* within the same directory as *drive.exe*.
2. Assuming that the CLR hasn't found *vehicle.dll*, it looks for *vehicle.dll* in a subdirectory with the same name as the assembly name (i.e., *vehicle*).
3. Assuming that the CLR hasn't found the vehicle assembly, it looks for a file called *vehicle.exe* in the same directory as *drive.exe*.
4. Assuming that the CLR hasn't found *vehicle.exe*, it looks for *vehicle.exe* in a subdirectory with the same name as the assembly name (i.e., *vehicle*).
5. At this point, the CLR throws an exception, indicating that it has failed to find the *vehicle* assembly.

The search for other supporting assemblies, such as *car* and *plane*, follows the same order. In steps 2 and 4, the CLR looks for the supporting assemblies in specific subdirectories. Let's investigate an example of this in the following section, which we've called multiple-directory deployment.

4.1.1.2 Multiple-directory deployment

Instead of storing all assemblies in the same directory as your client application, you can also use multiple, private subdirectories to segregate your assemblies so that they are easier to find and manage. For example, we will separate the vehicle, car, and plane assemblies into their own private directories, as shown in [Figure 4-1](#). We will leave the *drive.exe* application in the top directory, *MultiDirectories*.

Figure 4-1. Multiple-directory tree of components

When you build the *vehicle* assembly, you don't have to do anything special, as it doesn't reference or use any third-party assemblies. However, when you build the car or plane assembly, you must refer to the correct vehicle component (i.e., the one in the *vehicle* directory). For example, to build the plane assembly successfully, you must explicitly refer to *vehicle.dll* using a specific or relative path, as shown in the following command (`cd` to the *plane* directory):

```
csc /r:..\vehicle\vehicle.dll /t:library /out:plane.dll plane.cs
```

You can build the car assembly the same way you build the plane assembly. To compile your client application, you must also refer to your dependencies using the correct paths (`cd` to the main directory, *MultiDirectories*, before you type this command all on one line):

```
vjc /r:vehicle\vehicle.dll;car\car.dll;plane\plane.dll  
/t:exe /out:drive.exe drive.jsl
```

Based on the previously discussed search algorithm, the CLR can find the supporting assemblies within the appropriate subdirectories.

4.1.2 Shared Components

Unlike application-private assemblies, *shared assemblies*—ones that can be used by any client application—must be published or registered in the system Global Assembly Cache (GAC). When you register your assemblies against the GAC, they act as system components, such as a system DLL that every process in the system can use. A prerequisite for GAC registration is that the component must possess originator and version information. In addition to other metadata, these two items allow multiple versions of the same component to be registered and executed on the same machine. Again, unlike COM, you don't have to store any information in the system registry for clients to use these shared assemblies.

There are three general steps to registering your shared assemblies against the GAC:

1. Use the shared named (*sn.exe*) utility to obtain a public/private key pair. This utility generates a random key pair for you and saves the key information in an output file—for example, *originator.key*.
2. Build your assembly with an assembly version number and the key information from *originator.key*.
3. Use the .NET Global Assembly Cache Utility (*gacutil.exe*) to register your assembly in the GAC. This assembly is now a shared assembly and can be used by any client.

The commands that we use in this section refer to relative paths, so if you're following along, make sure that you create the directory structure, as shown in [Figure 4-2](#). The *vehicle*, *plane*, and *car* directories hold their appropriate assemblies, and the *key* directory holds the public/private key pair that we will generate in a moment. The *car-build* directory holds a car assembly with a modified build number, and the *car-revision* directory holds a car assembly with a modified revision number.

Figure 4-2. Directory structure for examples in this section

4.1.2.1 Generating a random key pair

We will perform the first step once and reuse the key pair for all shared assemblies that we build in this section. We're doing this for brevity only because you can use different key information for each assembly, or even each version, that you build. Here's how to generate a random key pair (be sure to issue this command in the *key* directory):

```
sn -k originator.key
```

The `-k` option generates a random key pair and saves the key information into the *originator.key* file. We will use this file as input when we build our shared assemblies. Let's now examine steps 2 and 3

registering your shared assemblies against the GAC.

4.1.2.2 Making the vehicle component a shared assembly

In order to add version and key information into the vehicle component (developed using Managed C++), we need to make some minor modifications to *vehicle.cpp*, as follows:

```
#using<mscorlib.dll>
using namespace System;

using namespace System::Reflection;
[assembly:AssemblyVersion("1.0.0.0")];
[assembly:AssemblyKeyFile("../key\\originator.key")];

public _ _gc _ _interface ISteering
{
    void TurnLeft( );
    void TurnRight( );
};

public _ _gc class Vehicle : public ISteering
{
    public:

        virtual void TurnLeft( )
        {
            Console::WriteLine("Vehicle turns left.");
        }

        virtual void TurnRight( )
        {
            Console::WriteLine("Vehicle turn right.");
        }

        virtual void ApplyBrakes( ) = 0;
};
```

The first boldface line indicates that we're using the `Reflection` namespace, which defines the attributes that the compiler will intercept to inject the correct information into our assembly manifest (For a discussion of attributes, see [Section 4.3.1](#) later in this chapter.) We use the `AssemblyVersion` attribute to indicate the version of this assembly, and we use the `AssemblyKeyFile` attribute to indicate the file containing the key information that the compiler should use to derive the public-key-token value, to be revealed in a moment.

Once you've done this, you can build this assembly using the following commands, which you've seen before:

```
cl /CLR /c vehicle.cpp
link -dll /out:vehicle.dll vehicle.obj
```

After you've built the assembly, you can use the .NET GAC Utility to register this assembly into the

GAC, as follows:

```
gacutil.exe /i vehicle.dll
```

Successful registration against the cache turns this component into a shared assembly. A version of this component is copied into the GAC so that even if you delete this file locally, you will still be able to run your client program.^[3]

^[3] However, don't delete the file now because we need it to build the car and plane assemblies.

4.1.2.3 Making the car component a shared assembly

In order to add version and key information into the car component, we need to make some minor modifications to *car.vb*, as follows:

```
Imports System

Imports System.Reflection
<Assembly:AssemblyVersion("1.0.0.0")>
<assembly:AssemblyKeyFile("../key\originator.key")>

Public Class Car
    Inherits Vehicle

    Overrides Public Sub TurnLeft( )
        Console.WriteLine("Car turns left.")
    End Sub

    Overrides Public Sub TurnRight( )
        Console.WriteLine("Car turns right.")
    End Sub

    Overrides Public Sub ApplyBrakes( )
        Console.WriteLine("Car trying to stop.")

        Console.WriteLine("ORIGINAL VERSION - 1.0.0.0.")

        throw new Exception("Brake failure!")
    End Sub
End Class
```

Having done this, you can now build it with the following command:

```
vbc /r:..\vehicle\vehicle.dll /t:library /out:car.dll car.vb
```

Once you've built this component, you can register it against the GAC:

```
gacutil /i car.dll
```

At this point, you can delete *car.dll* in the local directory because it has been registered in the GAC.

4.1.2.4 Making the plane component a shared assembly

In order to add version and key information into the plane component, we need to make some minor modifications to *plane.cs*, as follows:

```
using System;

using System.Reflection;
[assembly:AssemblyVersion("1.0.0.0")]
[assembly:AssemblyKeyFile("../key\\originator.key")]

public class Plane : Vehicle
{
    override public void TurnLeft( )
    {
        Console.WriteLine("Plane turns left.");
    }

    override public void TurnRight( )
    {
        Console.WriteLine("Plane turns right.");
    }

    override public void ApplyBrakes( )
    {
        Console.WriteLine("Air brakes being used.");
    }
}
```

Having done this, you can build the assembly with the following commands:

```
csc /r:..\vehicle\vehicle.dll /t:library /out:plane.dll plane.cs
gacutil /i plane.dll
```


Of course, the last line in this snippet simply registers the component into the GAC.

4.1.2.5 Viewing the GAC

Now that we've registered all our components into the GAC, let's see what the GAC looks like. Microsoft has shipped a shell extension, the Shell Cache Viewer, to make it easier for you to view the GAC. On our machines, the Shell Cache Viewer appears when we navigate to *C:\WINDOWS\Assembly*, as shown in [Figure 4-3](#).^[4]

[4] This path is entirely dependent upon the *%windir%* setting on your machine.

Figure 4-3. Our shared assemblies in the GAC



As you can see, the Shell Cache Viewer shows that all our components have the same version number because we used 1.0.0.0 as the version number when we built our components. Additionally it shows that all our components have the same public-key-token value because we used the same key file, *originator.key*.

4.1.2.6 Building and testing the drive.exe

You should copy the previous *drive.js*/source-code file into the *Shared Assemblies* directory, the root of the directory structure (shown in [Figure 4-2](#)) we are working with in this section. Having done this, you can build this component as follows (remember to type everything on one line):

```
vjc /r:vehicle\vehicle.dll;car\car.dll;plane\plane.dll
/t:exe /out:drive.exe drive.jsl
```

Once you've done this, you can execute the *drive.exe* component, which will use the *vehicle.dll*, *car.dll*, and *plane.dll* assemblies registered in the GAC. You should see the following as part of your output:

```
ORIGINAL VERSION - 1.0.0.0.
```

To uninstall these shared components (assuming that you have administrative privileges), select the appropriate assemblies and press the Delete key (but if you do this now, you must reregister these assemblies because we'll need them in the upcoming examples). When you do this, you've taken all the residues of these components out of the GAC. All that's left is to delete any files that you've copied over from your installation diskette-typically, all you really have to do is recursively remove the application directory.

4.1.2.7 Adding new versions

Unlike private assemblies, shared assemblies can take advantage of the rich versioning policies that the CLR supports. Unlike earlier OS-level infrastructures, the CLR enforces versioning policies during the loading of all shared assemblies. By default, the CLR loads the assembly with which your application was built, but by providing an application configuration file, you can command the CLR to load the specific assembly version that your application needs. Inside an application configuration file you can specify the rules or policies that the CLR should use when loading shared assemblies on which your application depends.

Let's make some code changes to our car component to demonstrate the default versioning support.

Remember that Version 1.0.0.0 of our car component's `ApplyBrakes()` method throws an exception, as follows:

```
Overrides Public Sub ApplyBrakes( )
    Console.WriteLine("Car trying to stop.")

    Console.WriteLine("ORIGINAL VERSION - 1.0.0.0.")

    throw new Exception("Brake failure!")
End Sub
```

Let's create a different *build* to remove this exception. To do this, make the following changes to the `ApplyBrakes()` method (store this source file in the *car-build* directory):

```
Overrides Public Sub ApplyBrakes( )
    Console.WriteLine("Car trying to stop.")

    Console.WriteLine("BUILD NUMBER change - 1.0.1.0.")

End Sub
```

In addition, you need to change the build number in your code as follows:

```
<Assembly:AssemblyVersion("1.0.1.0")>
```

Now build this component, and register it using the following commands:

```
vbc /r:..\vehicle\vehicle.dll
    /t:library /out:car.dll car.vb
gacutil /i car.dll
```

Notice that we've specified that this version is 1.0.1.0, meaning that it's compatible with Version 1.0.0.0. After registering this assembly with the GAC, execute your *drive.exe* application, and you will see the following statement as part of the output:

```
ORIGINAL VERSION - 1.0.0.0.
```

This is the default behavior-the CLR will load the version of the assembly with which your application was built. And just to prove this statement further, suppose that you provide Version 1.0.1.1 by making the following code changes (store this version in the *car-revision* directory):

```
Overrides Public Sub ApplyBrakes( )
    Console.WriteLine("Car trying to stop.")

    Console.WriteLine("REVISION NUMBER change - 1.0.1.1.")

End Sub
```

```
<Assembly:AssemblyVersion("1.0.1.1")>
```


This time, instead of changing the build number, you're changing the revision number, which should still be compatible with the previous two versions. If you build this assembly, register it against the GAC and execute *drive.exe* again; you will get the following statement as part of your output:

ORIGINAL VERSION - 1.0.0.0.

Again, the CLR chooses the version with which your application was built.

As shown in [Figure 4-4](#), you can use the Shell Cache Viewer to verify that all three versions exist on the system simultaneously. This implies that the support exists for side-by-side execution-which terminates DLL Hell in .NET.

Figure 4-4. Multiple versions of the same shared assembly



If you want your program to use a different, compatible version of the car assembly, you have to provide an application configuration file. The name of an application configuration file is composed of the physical executable name and ".config" appended to it. For example, since our client program is named *drive.exe*, its configuration file must be named *drive.exe.config*.

Here's a *drive.exe.config* file that allows you to tell the CLR to load Version 1.0.1.0 of the car assembly for you (instead of loading the default Version, 1.0.0.0). The two boldface attributes say that although we built our client with Version 1.0.0.0 (**oldVersion**) of the car assembly, load 1.0.1.0 (**newVersion**) for us when we run *drive.exe*.

```
<?xml version="1.0"?>
<configuration>
  <runtime>
    <assemblyBinding xmlns="urn:schemas-microsoft-com:asm.v1">
      <dependentAssembly>
        <assemblyIdentity name="car"
          publicKeyToken="D730D98B6BDE2BBA"
          culture="" />
        <bindingRedirect oldVersion="1.0.0.0"
          newVersion="1.0.1.0" />
      </dependentAssembly>
    </assemblyBinding>
  </runtime>
</configuration>
```

In this configuration file, the **name** attribute of the **assemblyIdentity** tag indicates the shared assembly's human-readable name that is stored in the GAC. Although the **name** value can be anything, you must replace the **publicKeyToken** value appropriately in order to execute *drive.exe*. The **publicKeyToken** attribute records the public-key-token value, which is an 8-byte hash of the public key used to build this component. There are several ways to get this 8-byte hash: you can copy it from the Shell Cache Viewer, you can copy it from the IL dump of your component, or you can

use the Shared Name utility to get it, as follows:

```
sn -T car.dll
```

Once you create the previously shown configuration file (stored in the same directory as the *drive.exe* executable) and execute *drive.exe*, you will see the following as part of your output:

```
BUILD NUMBER change - 1.0.1.0.
```

If you change the configuration file to `newVersion=1.0.1.1` and execute *drive.exe* again, you will see the following as part of your output:

```
REVISION NUMBER change - 1.0.1.1.
```

Having gone over all these examples, you should realize that you have full control over which dependent assembly versions the CLR should load for your applications. It doesn't matter which version was built with your application: you can choose different versions at runtime merely by changing a few attributes in the application configuration file.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

4.2 Distributed Components

A component technology should support distributed computing, allowing you to activate and invoke remote services, as well as services in another application domain.^[5] Distributed COM, or DCOM, is the wire protocol that provides support for distributed computing using COM. Although DCOM is fine for distributed computing, it is inappropriate for global cyberspace because it doesn't work well in the face of firewalls and NAT software. Some other shortcomings of DCOM are expensive lifecycle management, protocol negotiation, and binary formats.

^[5] Each Windows process requires its own memory address space, making it fairly expensive to run multiple Windows processes. An application domain is a lightweight or virtual process. All application domains of a given Windows process can use the same memory address space.

To eliminate or at least mitigate these shortcomings, .NET provides a host of different distributed support. The Remoting API in .NET allows you to use a host of channels, such as TCP and HTTP (which uses SOAP by default), for distributed computing. It even permits you to plug in your own custom channels, should you require this functionality. Best of all, since the framework is totally object-oriented, distributed computing in .NET couldn't be easier. To show you how simple it is to write a distributed application in .NET, let's look at an example using sockets, otherwise known as the *TCP channel* in .NET.

4.2.1 Distributed Hello Server

In this example, we'll write a distributed Hello application, which outputs a line of text to the console whenever a client invokes its exposed method, `SayHello()`. Since we're using the TCP channel, we'll tell the compiler that we need the definitions in the `System.Runtime.Remoting` and `System.Runtime.Remoting.Channels.Tcp` namespaces.

Note that this class, `CoHello`, derives from `MarshalByRefObject`.^[6]

^[6] If you fail to do this, your object will not have a distributed identity since the default is marshal-by-value, which means that a copy of the remote object is created on the client side.

This is the key to distributed computing in .NET because it gives this object a distributed identity, allowing the object to be referenced across application domains, or even process and machine boundaries. A marshal-by-reference object requires a proxy to be set up on the client side and a stub to be set up on the server side, but since both of these are automatically provided by the infrastructure, you don't have to do any extra work. Your job is to derive from `MarshalByRefObject` to get all the support for distributed computing:

```
using System;

using System.Runtime.Remoting;
using System.Runtime.Remoting.Channels;
using System.Runtime.Remoting.Channels.Tcp;

public class CoHello : MarshalByRefObject
```

```

{
    public static void Main( )
    {
        TcpChannel channel = new TcpChannel(4000);
        ChannelServices.RegisterChannel(channel);

        RemotingConfiguration.RegisterWellKnownServiceType (
            typeof(CoHello),           // Type name
            "HelloDotNet",           // URI
            WellKnownObjectMode.Singleton // SingleCall or Singleton
        );

        System.Console.WriteLine("Hit <enter> to exit . . . ");
        System.Console.ReadLine( );
    }

    public void SayHello( )
    {
        Console.WriteLine("Hello, Universe of .NET");
    }
}

```

The SayHello() method is public, meaning that any external client can call this method. As you can see, this method is very simple, but the interesting thing is that a remote client application (which we'll develop shortly) can call it because the Main() function uses the TcpChannel class. Look carefully at Main(), and you'll see that it instantiates a TcpChannel, passing in a port number from which the server will listen for incoming requests.^[7]

[7] Believe it or not, all you really have to do is replace TcpChannel with HttpChannel to take advantage of HTTP and SOAP as the underlying communication protocols.

Once we have created a channel object, we then register the channel to the ChannelServices, which supports channel registration and object resolution. Having done this, you must then register your object with the RemotingConfiguration so that it can be activated—you do this by calling the RegisterWellKnownServiceType() method of the RemotingConfiguration class. When you call this method, you must pass in the class name, a URI, and an object-activation mode. The URI is important because it's a key element that the client application will use to refer specifically to this registered object. The object-activation mode can be either Singleton, which means that the same object will service many calls, or SingleCall, which means an object will service at most one call.

Here's how to build this distributed application:

```
csc server.cs
```

Once you've done this, you can start the server program, which will wait endlessly until you hit the Enter key. The server is now ready to service client requests.

4.2.2 Remote Hello Client

Now that we have a server waiting, let's develop a client to invoke the remote SayHello() method. Instead of registering an object with the remoting configuration, we need to activate a remote object

So let's jump into the code now to see how this works. As you examine the following program, note these items:

- We're using types in the `System.Runtime.Remoting` and `System.Runtime.Remoting.Channels.Tcp` namespaces, since we want to use the TCP channel.
- Our `Client` class doesn't need to derive from anything because it's not a server-side object that needs to have a distributed identity.
- Since we're developing a client application, we don't need to specify a client port when we instantiate the `TcpChannel`.

Other than these items, the key thing to note is object activation, shown in the second boldface statement in the following code. To invoke remote methods, you must first activate the remote object and obtain an associated proxy on the client side. To activate the object and get a reference to the associated proxy, you call the `GetObject()` method of the `Activator` class. When you do this, you must pass along the remote class name and its fully qualified location, including the complete URI. Once you've successfully done this, you can then invoke remote methods.

```
using System;
using System.Runtime.Remoting;
using System.Runtime.Remoting.Channels;
using System.Runtime.Remoting.Channels.Tcp;

public class Client
{
    public static void Main( )
    {
        try
        {
            TcpChannel channel = new TcpChannel( );
            ChannelServices.RegisterChannel(channel);

            CoHello h = (CoHello) Activator.GetObject(
                typeof(CoHello),           // Remote type
                "tcp://127.0.0.1:4000/HelloDotNet" // Location
            );

            h.SayHello( );
        }
        catch(Exception e)
        {
            Console.WriteLine(e.ToString( ));
        }
    }
}
```

To build this client application, you must include references to the *server.exe* assembly:

```
csc /r:Server.exe Client.cs
```

If you're familiar with DCOM, you must be relieved to find that it's relatively simple to write

distributed applications in .NET.^[8]

[8] In fact, if you have a copy of *Learning DCOM* (O'Reilly) handy, compare these programs with their DCOM counterparts in [Appendix D](#), and you will see what we mean.

4.2.3 Distributed Garbage Collector

Because the .NET distributed garbage collector is different from that of DCOM, we must briefly cover this facility. Instead of using DCOM's delta pinging, which requires few network packets when compared to normal pinging (but still too many for a distributed protocol), .NET remoting uses leases to manage object lifetimes. If you've ever renewed the lease to an IP address on your Dynamic Host Configuration Protocol (DHCP) network, you've pretty much figured out this mechanism because it's based on similar concepts.

In .NET, distributed objects give out leases instead of relying on reference counting (as in COM) for lifetime management. An application domain where the remote objects reside has a special object called the *lease manager*, which manages all the leases associated with these remote objects. When a lease expires, the lease manager contacts a sponsor, telling the sponsor that the lease has expired. A *sponsor* is simply a client that has previously registered itself with the lease manager during an activation call, indicating to the lease manager that it wants to know when a lease expires. If the lease manager can contact the sponsor, the sponsor may then renew the lease. If the sponsor refuses to renew the lease or if the lease manager can't contact the sponsor after a configurable timeout period, the lease manager will void the lease and remove the object. There are two other ways in which a lease can be renewed: implicitly, via each call to the remote object, or explicitly, by calling the `Renew()` method of the `ILease` interface.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

4.3 COM+ Services in .NET

COM programming requires lots of housekeeping and infrastructure-level code to build large-scale, enterprise applications. To make it easier to develop and deploy transactional and scalable COM applications, Microsoft released Microsoft Transaction Server (MTS). MTS allows you to share resources, thereby increasing the scalability of an application. COM+ Services were the natural evolution of MTS. While MTS was just another library on top of COM, COM+ Services were subsumed into the COM library, thus combining both COM and MTS into a single runtime.

COM+ Services have been very valuable to the development shops using the COM model to build applications that take advantage of transactions, object pooling, role-based security, etc. If you develop enterprise .NET applications, the COM+ Services in .NET are a must.

In the following examples, rather than feeding you more principles, we'll show you examples for using major COM+ Services in .NET, including examples on transactional programming, object pooling, and role-based security. But before you see these examples, let's talk about the key element-attributes-that enables the use of these services in .NET.

4.3.1 Attribute-Based Programming

Attributes are the key element that helps you write less code and allows an infrastructure to automatically inject the necessary code for you at runtime. If you've used IDL (Interface Definition Language) before, you have seen the `in` or `out` attributes, as in the following example:

```
HRESULT SetAge([in] short age);
HRESULT GetAge([out] short *age);
```

IDL allows you to add these attributes so that the marshaler will know how to optimize the use of the network. Here, the `in` attribute tells the marshaler to send the contents from the client to the server, and the `out` attribute tells the marshaler to send the contents from the server to the client. In the `SetAge()` method, passing `age` from the server to the client will just waste bandwidth. Similarly, there's no need to pass `age` from the client to the server in the `GetAge()` method.

4.3.1.1 Developing custom attributes

While `in` and `out` are built-in attributes the MIDL compiler supports, .NET allows you to create your own custom attributes by deriving from the `System.Attribute` class. Here's an example of a custom attribute:

```
using System;

public enum Skill {  Guru, Senior, Junior }
[AttributeUsage(AttributeTargets.Class
                AttributeTargets.Field
                )]
```

```

        AttributeTargets.Method |
        AttributeTargets.Property |
        AttributeTargets.Constructor |
        AttributeTargets.Event)]
public class AuthorAttribute : System.Attribute
{
    public AuthorAttribute(Skill s)
    {
        level = s;
    }
    public Skill level;
}

```

The AttributeUsage attribute that we've applied to our AuthorAttribute class specifies the rules for using AuthorAttribute.^[9] Specifically, it says that `AuthorAttribute` can prefix or describe a class or any class member.

[9] You don't have to postfix your attribute class name with the word "Attribute", but this is a standard naming convention that Microsoft uses. C# lets you name your attribute class any way you like; for example, Author is a valid class name for your attribute.

4.3.1.2 Using custom attributes

Given that we have this attribute, we can write a simple class to make use of it. To apply our attribute to a class or a member, we simply make use of the attribute's available constructors. In our case, we have only one and it's `AuthorAttribute()`, which takes an author's skill level. Although you can use `AuthorAttribute()` to instantiate this attribute, .NET allows you to drop the `Attribute` suffix for convenience, as shown in the following code listing:

```

[Author(Skill.Guru)]
public class Customer
{
    [Author(Skill.Senior)]
    public void Add(string strName)
    {
    }
    [Author(Skill.Junior)]
    public void Delete(string strName)
    {
    }
}

```

You'll notice that we've applied the Author attribute to the Customer class, telling the world that a guru wrote this class definition. This code also shows that a senior programmer wrote the `Add()` method and that a junior programmer wrote the `Delete()` method.

4.3.1.3 Inspecting attributes

You won't see the full benefits of attributes until you write a simple interceptor-like program, which looks for special attributes and provides additional services appropriate for these attributes. Real

interceptors include marshaling, transaction, security, pooling, and other services in MTS and COM+.

Here's a simple interceptor-like program that uses the Reflection API to look for `AuthorAttribute` and provide additional services. You'll notice that we can ask a type, `Customer` in this case, for all of its custom attributes. In our code, we ensure that the `Customer` class has attributes and that the first attribute is `AuthorAttribute` before we output the appropriate messages to the console. In addition, we look for all members that belong to the `Customer` class and check whether they have custom attributes. If they do, we ensure that the first attribute is an `AuthorAttribute` before we output the appropriate messages to the console.

```
using System;
using System.Reflection;

public class interceptor
{
    public static void Main( )
    {
        Object[] attrs = typeof(Customer).GetCustomAttributes(false);
        if ((attrs.Length > 0) && (attrs[0] is AuthorAttribute))
        {
            Console.WriteLine("Class [{0}], written by a {1} programmer.",
                typeof(Customer).Name, ((AuthorAttribute)attrs[0]).level);
        }

        MethodInfo[] mInfo = typeof(Customer).GetMethods( );
        for ( int i=0; i < mInfo.Length; i++ )
        {
            attrs = mInfo[i].GetCustomAttributes(false);

            if ((attrs.Length > 0) && (attrs[0] is AuthorAttribute))
            {
                AuthorAttribute a = (AuthorAttribute)attrs[0];
                Console.WriteLine("Method [{0}], written by a {1} programmer.",
                    mInfo[i].Name, (a.level));
                if (a.level == Skill.Junior)
                {
                    Console.WriteLine("***Performing automatic " +
                        "review of {0}'s code***", a.level);
                }
            }
        }
    }
}
```

It is crucial to note that when this program sees a piece of code written by a junior programmer, it automatically performs a rigorous review of the code. If you compile and run this program, it will output the following to the console:

```
Class [Customer], written by a Guru programmer.
Method [Add], written by a Senior programmer.
Method [Delete], written by a Junior programmer.
***Performing automatic review of Junior's code***
```

Although our interceptor-like program doesn't intercept any object-creation and method invocations, it does show how a real interceptor can examine attributes at runtime and provide necessary services stipulated by the attributes. Again, the key here is the last boldface line, which represents a special service that the interceptor provides as a result of attribute inspection.

4.3.2 Transactions

In this section, we'll show you that it's easy to write a .NET class to take advantage of the transaction support that COM+ Services provide. All you need to supply at development time are a few attributes, and your .NET components are automatically registered against the COM+ catalog the first time they are used. Put differently, not only do you get easier programming, but you also get just-in-time and automatic registration of your COM+ application.^[10]

^[10] Automatic registration is nice during development, but don't use this feature in a production environment, because not all clients will have the administrative privilege to set up COM+ applications.

To develop a .NET class that supports transactions, here's what must happen:

1. Your class must derive from the `ServiceComponent` class to exploit COM+ Services.
2. You must describe your class with the correct Transaction attribute, such as `Transaction(TransactionOption.Required)`, meaning that instances of your class must run within a transaction.

Besides these two requirements, you can use the `ContextUtil` class (which is a part of the `System.EnterpriseServices` namespace) to obtain information about the COM+ object context. This class exposes the major functionality found in COM+, including methods such as `SetComplete()`, `SetAbort()`, and `IsCallerInRole()`, and properties such as `IsInTransaction` and `MyTransactionVote`.

In addition, while it's not necessary to specify any COM+ application installation options, you should do so because you get to specify what you want, including the name of your COM+ application, its activation setting, its versions, and so on. For example, in the following code listing, if you don't specify the `ApplicationName` attribute, .NET will use the module name as the COM+ application name displayed in the Component Services Explorer (or COM+ Explorer). For example, if the name of module is `crm.dll`, the name of your COM+ application will be `crm`. Other than this attribute, we also use the `ApplicationActivation` attribute to specify that this component will be installed as a library application, meaning that the component will be activated in the creator's process:

```
using System;
using System.Reflection;
using System.EnterpriseServices;

[assembly: ApplicationName(".NET Framework Essentials CRM")]
[assembly: ApplicationActivation(ActivationOption.Library)]

[assembly: AssemblyKeyFile("originator.key")]
[assembly: AssemblyVersion("1.0.0.0")]
```

The rest should look extremely familiar. In the `Add()` method, we simply call `SetComplete()` when we've successfully added the new customer into our databases. If something has gone wrong during the process, we will vote to abort this transaction by calling `SetAbort()`.

```

[Transaction(TransactionOption.Required)]
public class Customer : ServicedComponent
{
    public void Add(string strName)
    {
        try
        {
            Console.WriteLine("New customer: {0}", strName);

            // Add the new customer into the system
            // and make appropriate updates to
            // several databases.

            ContextUtil.SetComplete( );
        }
        catch(Exception e)
        {
            Console.WriteLine(e.ToString( ));
            ContextUtil.SetAbort( );
        }
    }
}

```

Instead of calling SetComplete() and SetAbort() yourself, you can also use the AutoComplete attribute, as in the following code, which is conceptually equivalent to the previously shown Add() method:

```

[AutoComplete]
public void Add(string strName)
{
    Console.WriteLine("New customer: {0}", strName);

    // Add the new customer into the system
    // and make appropriate updates to
    // several databases.
}

```

Here's how you build this assembly:

```
csc /t:library /out:crm.dll crm.cs
```

Since this is a shared assembly, remember to register it against the GAC by using the GAC utility:

```
gacutil /i crm.dll
```

At this point, the assembly has not been registered as a COM+ application, but we don't need to register it manually. Instead, .NET automatically registers and hosts this component for us in a COM+ application the first time we use this component. So, let's write a simple client program that uses this component at this point. As you can see in the following code, we instantiate a Customer object and add a new customer:

```
using System;
```

```

public class Client
{
    public static void Main( )
    {
        try
        {
            Customer c = new Customer( );
            c.Add("John Osborn");
        }
        catch(Exception e)
        {
            Console.WriteLine(e.ToString( ));
        }
    }
}

```

We can build this program as follows:

```
csc /r:crm.dll /t:exe /out:client.exe client.cs
```

When we run this application, COM+ Services automatically create a COM+ application called **.NET Framework Essentials CRM** to host our *crm.dll* .NET assembly, as shown in [Figure 4-5](#). In addition to adding our component to the created COM+ application, .NET also inspects our metadata for provided attributes and configures the associated services in the COM+ catalog.

Figure 4-5. The Component Services explorer

4.3.3 Object Pooling

A *pool* is technical term that refers to a group of resources, such as connections, threads, and objects. Putting a few objects into a pool allows hundreds of clients to share these few objects (you can make the same assertion for threads, connections, and other objects). Pooling is, therefore, a technique that minimizes the use of system resources, improves performance, and helps system scalability.

Missing in MTS, *object pooling* is a nice feature in COM+ that allows you to pool objects that are expensive to create. Similar to providing support for transactions, if you want to support object pooling in a .NET class, you need to derive from `ServiceComponent`, override any of the `Activate()`,

Deactivate(), and CanBePooled() methods, and specify the object-pooling requirements in an ObjectPooling attribute, as shown in the following example:^[11]

[¹¹] Mixing transactions and object pooling should be done with care. See *COM and .NET Component Services*, by Juval Löwy (O'Reilly).

```
using System;
using System.Reflection;
using System.EnterpriseServices;

[assembly: ApplicationName(".NET Framework Essentials CRM")]
[assembly: ApplicationActivation(ActivationOption.Library)]
[assembly: AssemblyKeyFile("originator.key")]
[assembly: AssemblyVersion("1.0.0.0")]

[Transaction(TransactionOption.Required)]
[ObjectPooling(MinPoolSize=1, MaxPoolSize=5)]
public class Customer : ServicedComponent
{
    public Customer( )
    {
        Console.WriteLine("Some expensive object construction.");
    }

    [AutoComplete]
    public void Add(string strName)
    {
        Console.WriteLine("Add customer: {0}", strName);
        // Add the new customer into the system
        // and make appropriate updates to
        // several databases.
    }

    override protected void Activate( )
    {
        Console.WriteLine("Activate");
        // Pooled object is being activated.
        // Perform the appropriate initialization.
    }

    override protected void Deactivate( )
    {
        Console.WriteLine("Deactivate");
        // Object is about to be returned to the pool.
        // Perform the appropriate clean up.
    }

    override protected bool CanBePooled( )
    {
        Console.WriteLine("CanBePooled");
        return true; // Return the object to the pool.
    }
}
```

```
}
```

Take advantage of the `Activate()` and `Deactivate()` methods to perform appropriate initialization and cleanup. The `CanBePooled()` method lets you tell COM+ whether your object can be pooled when this method is called. You need to provide the expensive object-creation functionality in the constructor, as shown in the constructor of this class.

Given this `Customer` class that supports both transaction and object pooling, you can write the following client-side code to test object pooling. For brevity, we will create only two objects, but you can change this number to anything you like so that you can see the effects of object pooling. Just to ensure that you have the correct configuration, delete the current `.NET Framework Essentials CRM` COM+ application from the Component Services Explorer before running the following code:

```
for (int i=0; i<2; i++)
{
    Customer c = new Customer( );
    c.Add(i.ToString( ));
}
```

Running this code produces the following results:

```
Some expensive object construction.
Activate
Add customer: 0
Deactivate
CanBePooled
Activate
Add customer: 1
Deactivate
CanBePooled
```

We've created two objects, but since we've used object pooling, only one object is really needed to support our calls, and that's why you see only one output statement that says `Some expensive object construction`. In this case, COM+ creates only one `Customer` object, but activates and deactivates it twice to support our two calls. After each call, it puts the object back into the object pool. When a new call arrives, it picks the same object from the pool to service the request.

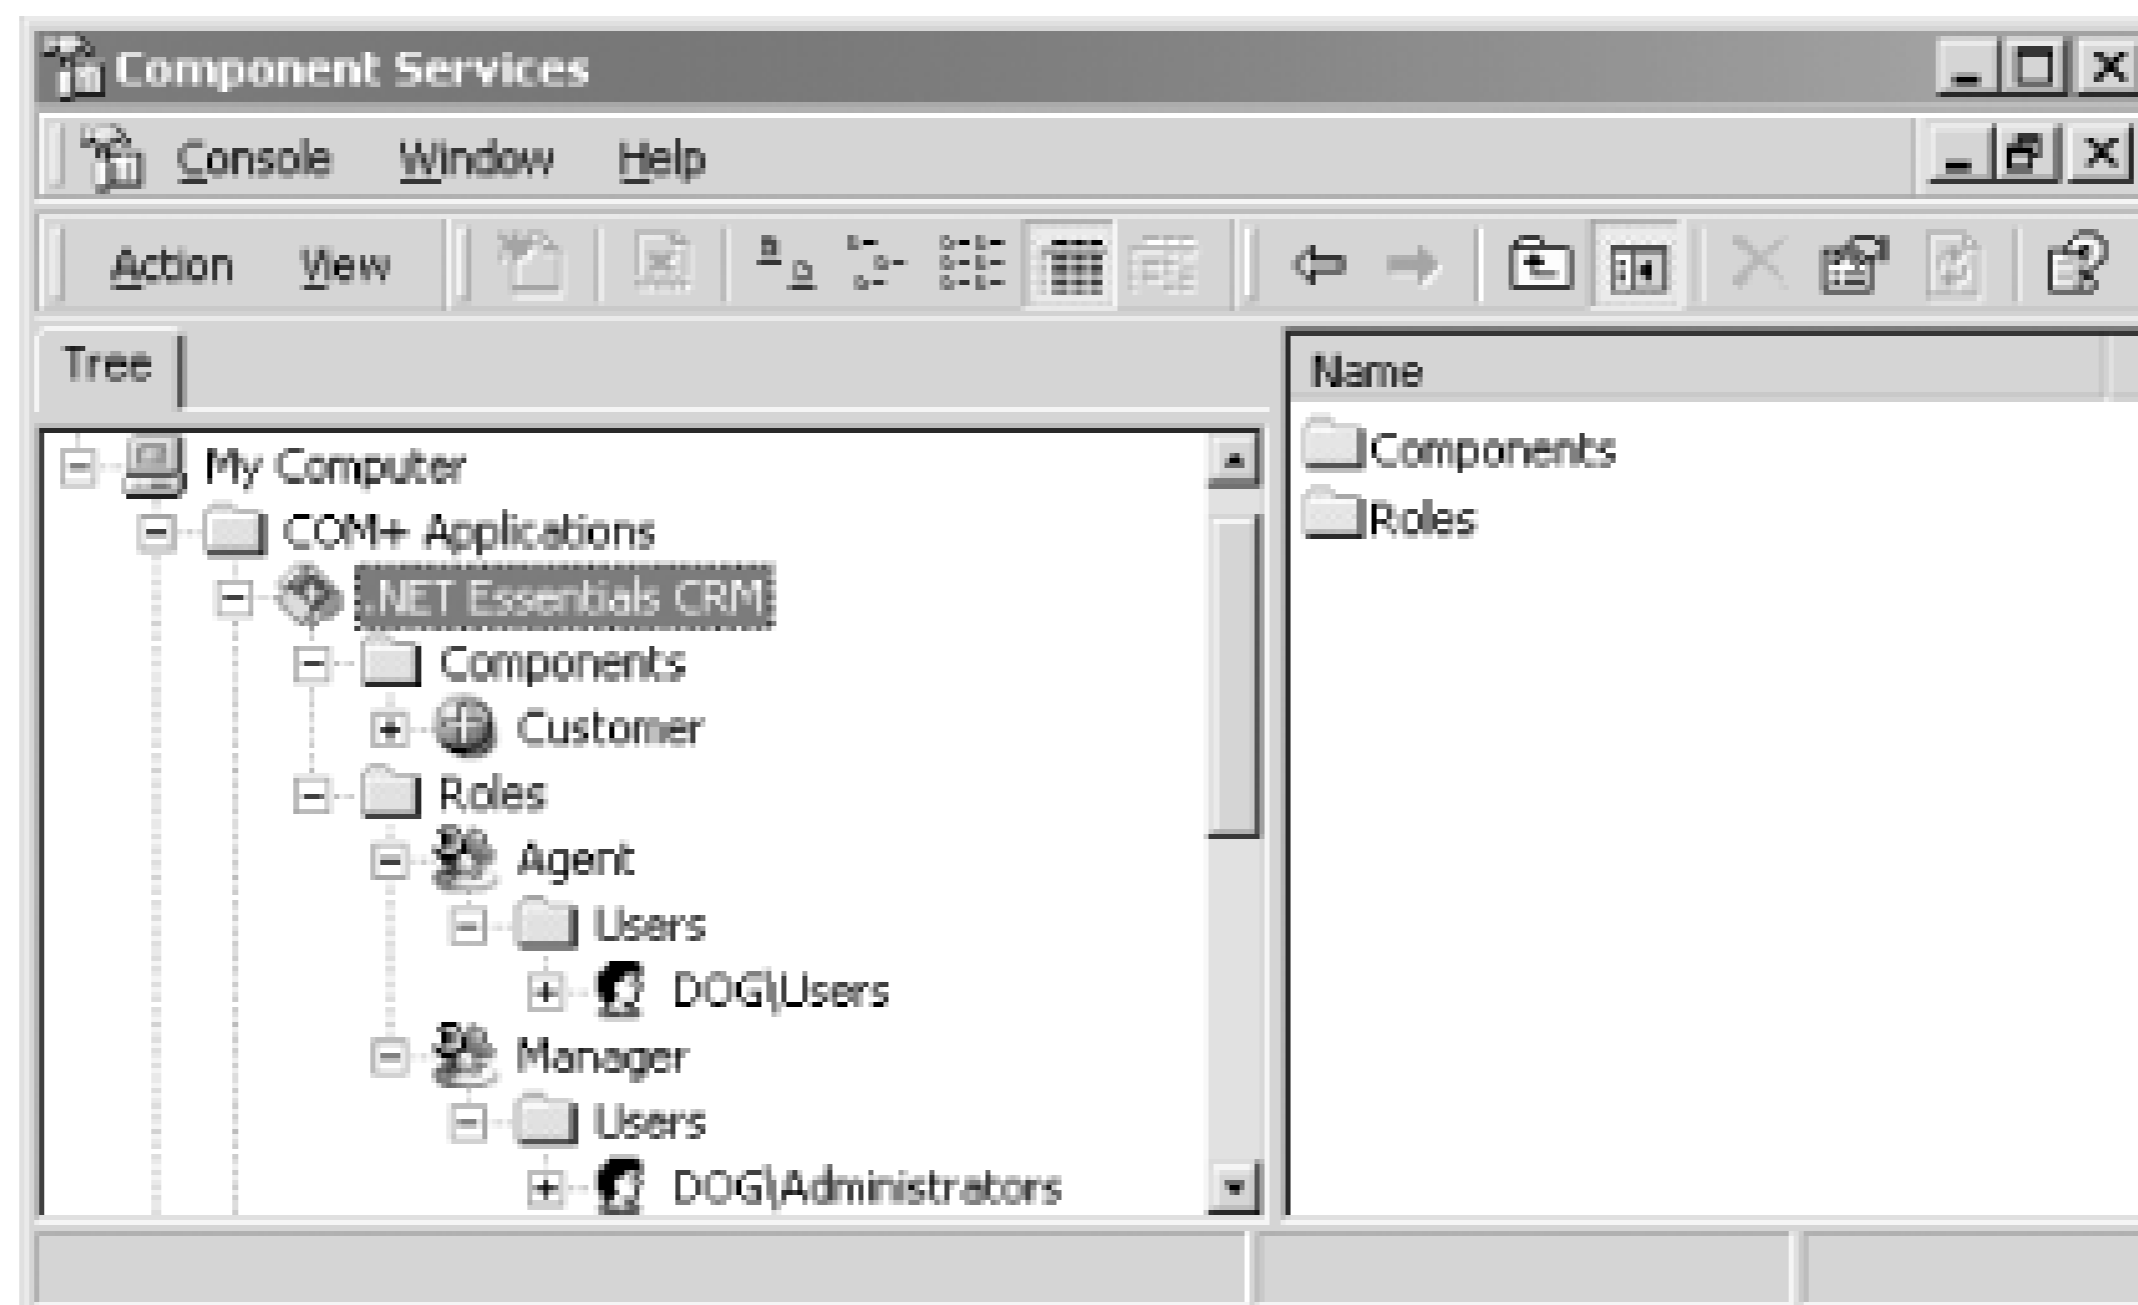
4.3.4 Role-Based Security

Role-based security in MTS and COM+ has drastically simplified the development and configuration of security for business applications. This is because it abstracts away the complicated details for dealing with access control lists (ACL) and security identifiers (SID). All .NET components that are hosted in a COM+ application can take advantage of role-based security. You can fully configure role-based security using the Component Services Explorer, but you can also manage role-based security in your code to provide fine-grain security support that's missing from the Component Services Explorer.

4.3.4.1 Configuring role-based security

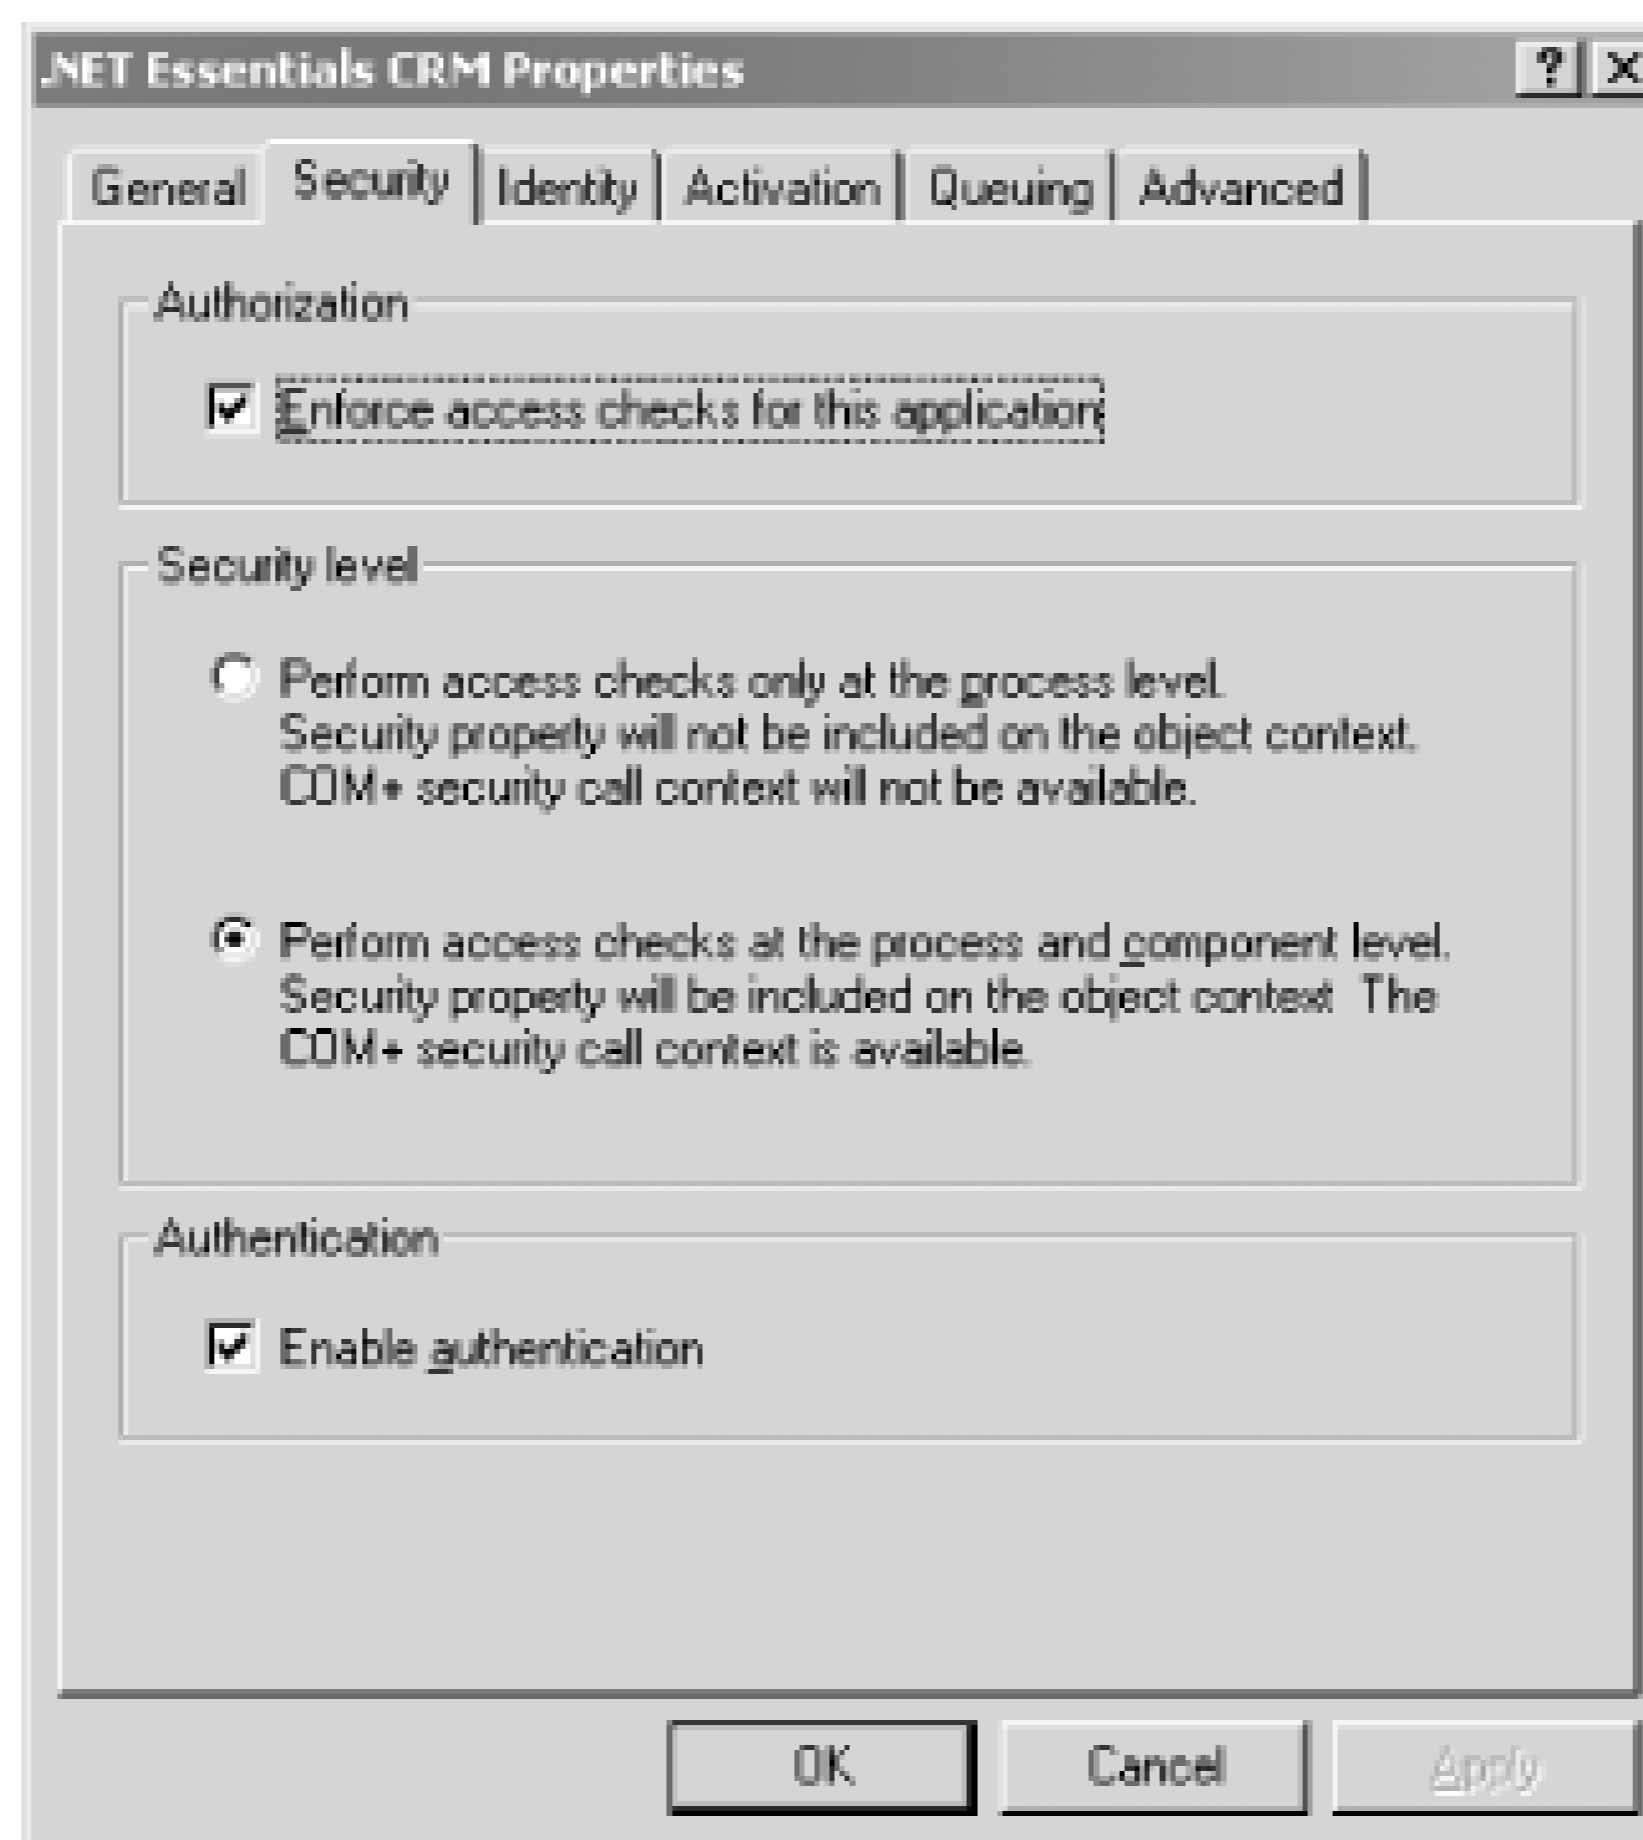
In order to demonstrate role-based security, let's add two roles to our COM+ application, **.NET Framework Essentials CRM**. The first role represents **Agent** who can use the Customer class in every way but can't delete customers. You should create this role and add to it the local **Users** group, as shown in [Figure 4-6](#). The second role represents **Manager** who can use the Customer class in every way, including deleting customers. Create this role and add to it the local **Administrators** group.

Figure 4-6. Creating roles and adding users to roles



Once you create these roles, you need to enable access checks for the **.NET Framework Essentials CRM** COM+ application. Launch the COM+ application's Properties sheet (by selecting **.NET Framework Essentials CRM** and pressing Alt-Enter), and select the Security tab. Enable access checks to your COM+ application by providing the options, as shown in [Figure 4-7](#).

Figure 4-7. Enable authorization for this COM+ application



Once you have enabled access checks at the application level, you need to enforce access checks at the class level, too. To do this, launch `Customer`'s Properties sheet, and select the Security tab. Enable access checks to this .NET class by providing the options shown in [Figure 4-8](#). Here, we're saying that no one can access the `Customer` class except for those that belong to the `Manager` or `Agent` role.

Figure 4-8. Enforce component-level access checks

Now, if you run the client application developed in the last section, everything will work because you are a user on your machine. But if you uncheck both the `Manager` [\[12\]](#) and `Agent` roles in [Figure 4-8](#) and rerun the client application, you get the following message as part of your output:

[12] Since you're a developer, you're probably an administrator on your machine, so you need to uncheck the

`Manager` role, too, in order to see an access violation in the test that we're about to illustrate.

```
System.UnauthorizedAccessException: Access is denied.
```

You're getting this exception because you've removed yourself from the roles that have access to the `Customer` class. Once you've verified this, put the configuration back to what is shown in [Figure 4-8](#) to prepare the environment for the next test that we're about to illustrate.

4.3.4.2 Programming role-based security

We've allowed anyone in the `Agent` and `Manager` roles to access our class, but let's invent a rule allowing only users under the `Manager` role to delete a customer from the system (for lack of a better example). So let's add a new method to the `Customer` class—we'll call this method `Delete()`, as shown in the following code. Anyone belonging to the `Agent` or `Manager` role can invoke this method, so we'll first output to the console the user account that invokes this method. After doing this, we'll check to ensure that this user belongs to the `Manager` role. If so, we allow the call to go through; otherwise, we throw an exception indicating that only managers can perform a deletion. Believe it or not, this is the basic premise for programming role-based security:

```
[AutoComplete]
public void Delete(string strName)
{
    try
    {
        SecurityCallContext sec;
        sec = SecurityCallContext.CurrentCall;
        string strCaller = sec.DirectCaller.AccountName;

        Console.WriteLine("Caller: {0}", strCaller);

        bool bInRole = sec.IsCallerInRole("Manager");
        if (!bInRole)
        {
            throw new Exception ("Only managers can delete customers.");
        }

        Console.WriteLine("Delete customer: {0}", strName);
        // Delete the new customer from the system
        // and make appropriate updates to
        // several databases.
    }
    catch(Exception e)
    {
        Console.WriteLine(e.ToString( ));
    }
}
```

Here's the client code that includes a call to the `Delete()` method:

```
using System;
```

```

public class Client
{
    public static void Main( )
    {
        try
        {
            Customer c = new Customer( );
            c.Add("John Osborn");
            // Success depends on the role
            // under which this method
            // is invoked.
            c.Delete("Jane Smith");
        }
        catch(Exception e)
        {
            Console.WriteLine(e.ToString( ));
        }
    }
}

```

Once you've built this program, you can test it using an account that belongs to the local `Users` group, since we added this group to the `Agent` role earlier. On Windows 2000 or XP, you can use the following command to launch a command window using a specific account:

```
runas /user:DEVTOUR\student cmd
```

Of course, you should replace `DEVTOUR` and `student` with your own machine name and user account, respectively. After running this command, you will need to type in the correct password, and a new command window will appear. Execute the client under this user account, and you'll see the following output:

```

Add customer: John Osborn
Caller: DEVTOUR\student
System.Exception: Only managers can delete customers .
    at Customer.Delete(String strName)

```

You'll notice that the `Add()` operation went through successfully, but the `Delete()` operation failed, because we executed the client application under an account that's missing from the `Manager` role.

To remedy this, we need to use a user account that belongs to the `Manager` role-any account that belongs to the `Administrators` group will do. So, start another command window using a command similar to the following:

```
runas /user:DEVTOUR\instructor cmd
```

Execute the client application again, and you'll get the following output:

```

Add customer: John Osborn
Caller: DEVTOUR\instructor
Delete customer: Jane Smith

```

As you can see, since we've executed the client application using an account that belongs to the `Manager` role, the `Delete()` operation went through without problems.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

4.4 Message Queuing

In addition to providing support for COM+ Services, .NET also supports message queuing. If you've used Microsoft Message Queuing (MSMQ) services before, you'll note that the basic programming model is the same but the classes in the System.Messaging namespace make it extremely easy to develop message-queuing applications. The System.Messaging namespace provides support for basic functionality, such as connecting to a queue, opening a queue, sending messages to a queue, receiving messages from a queue, and peeking for messages on the queue. To demonstrate how easy it is to use the classes in System.Messaging, let's build two simple applications: one to enqueue messages onto a private queue on the local computer and another to dequeue these messages from the same queue.^[13]

[13] To execute these programs, you must have MessageQueuing installed on your system. You can verify this by launching the ComputerManagement console, as shown in [Figure 4-9](#).

4.4.1 Enqueue

Here's a simple program that enqueues a Customer object onto a private queue on the local computer. Notice first that we need to include the System.Messaging namespace because it contains the classes that we want to use:

```
using System;  
using System.Messaging;
```

While the following Customer structure is very simple, it can be as complex as you want because it will be serialized into an XML-formatted buffer by default before it's placed into the queue:

```
public struct Customer  
{  
    public string Last;  
    public string First;  
}
```

Our program first checks whether a private queue on the local computer exists. If this queue is missing, the program will create it. Next, we instantiate a MessageQueue class, passing in the target queue name. Once we have this MessageQueue object, we invoke its Send() method, passing in the Customer object, as shown in the following code. This will put our customer object into our private queue:

```
public class Enqueue  
{  
    public static void Main( )  
    {  
        try  
        {  
            string path = ".\\PRIVATE$\\NFE_queue";        }  
    }  
}
```

```

    if(!MessageQueue.Exists(path))
    {
        // Create our private queue.
        MessageQueue.Create(path);
    }

    // Initialize the queue.
    MessageQueue q = new MessageQueue(path);

    // Create our object.
    Customer c = new Customer( );
    c.Last = "Osborn";
    c.First = "John";

    // Send it to the queue.
    q.Send(c);
}
catch(Exception e)
{
    Console.WriteLine(e.ToString( ));
}
}
}

```

Use the following command to build this program:

```
csc /t:exe /out:enqueue.exe enqueue.cs
```

Execute this program, examine the Computer Management console, and you will see your message in the private queue called `nfe_queue`, as shown in [Figure 4-9](#).

Figure 4-9. Our private queue, `ne_queue`, with a message

4.4.2 Dequeue

Now that there's a message in our private message queue, let's write a program to dequeue and examine the message. After ensuring that the private queue we want exists, we initialize it by instantiating a `MessageQueue` class, passing in the path to our private queue. Next, we tell the

MessageQueue object that the type of object we want to dequeue is Customer. To actually dequeue the object, we need to invoke the Receive() method, passing in a timeout in terms of a TimeSpan object, whose parameters stand for hours, minutes, and seconds, respectively. Finally, we cast the body of the received Message object into a Customer object and output its contents:

```
using System;
using System.Messaging;
using System.Runtime.Serialization;

public struct Customer
{
    public string Last;
    public string First;
}

public class Dequeue
{
    public static void Main( )
    {
        try
        {
            string strQueuePath = ".\\PRIVATE$\\NFE_queue";

            // Ensure that the queue exists.
            if (!MessageQueue.Exists(strQueuePath))
            {
                throw new Exception(strQueuePath + " doesn't exist!");
            }

            // Initialize the queue.
            MessageQueue q = new MessageQueue(strQueuePath);

            // Specify the types we want to get back.
            string[] types = {"Customer, dequeue"};
            ((XmlMessageFormatter)q.Formatter).TargetTypeNames = types;

            // Receive the message (5-second timeout).
            Message m = q.Receive(new TimeSpan(0,0,5));

            // Convert the body into the type we want.
            Customer c = (Customer) m.Body;

            Console.WriteLine("Customer: {0}, {1}", c.Last, c.First);
        }
        catch(Exception e)
        {
            Console.WriteLine(e.ToString( ));
        }
    }
}
```

Compile and execute this program, look at the [Computer Management](#) console, press F5 to refresh the

screen, and you will realize that the previous message is no longer there.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

4.5 Summary

In this chapter, we've touched on many aspects of component-oriented programming, including deployment strategies, distributed computing, and enterprise services such as transaction management, object pooling, role-based security, and message queuing. We have to give due credit to Microsoft for making componentization easier in the .NET Framework. Case in point: without .NET, it would be impossible for us to show the complete code for all of these programs in a single chapter of a book.^[14]

^[14] For additional information on programming .NET component-based applications, see O'Reilly's *Programming .NET Components*, by Juval Löwy.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Chapter 5. Data and XML

Almost everything we do in the software industry relates to data in some way. At some point, all software developers must deal with data, perhaps using a database, text file, spreadsheet, or some other method of data storage. There are many different methods and technologies for using, manipulating, and managing data, and newer methods are continually introduced to enhance existing ones. These methods range from function-based APIs to object-based frameworks and proprietary libraries.

Several years ago, it was common for a simple VB desktop application to access a private Microsoft Access database stored on the local hard disk, but this is no longer a typical scenario. Today's applications take advantage of distributed-component technologies to exploit scalability and interoperability, thus widening the reach of the application to the enterprise. Although ActiveX Data Objects (ADO) served a typical VB application well a few years ago, it fails to meet the increasing demands for better scalability, performance, and interoperability across multiple platforms that web-based applications require.

Here's where ADO.NET comes in. ADO.NET provides huge benefits that allow us to build even better enterprise applications. In this chapter, you will learn the benefits of ADO.NET, the ADO.NET architecture, the main classes in ADO.NET and how they work, and the integration of ADO.NET and XML.

[\[Team LiB \]](#)

[[Team LiB](#)]

5.1 ADO.NET Architecture

Microsoft ADO.NET's object model encompasses two distinct groups of classes: *content components* and *data-provider components*.^[1] The content components include the DataSet class and other supporting classes, such as DataTable, DataRow, DataColumn, and DataRelation. These classes contain the actual content of a data exchange. The data-provider components assist in data retrieval and updates. Developers can use the connection, command, and data reader objects to directly manipulate data. In more typical scenarios, developers use the DataAdapter class as the conduit to move data between the data store and the content components. The data can be actual rows from a database or any other form of data, such as an XML file or an Excel spreadsheet.

[1] In previous releases of this book, we referred to the Data-Provider components as managed-provider components.

[Figure 5-1](#) shows the high-level architecture of ADO.NET. ADO developers should have no problems understanding connection and command objects. We offer a brief overview then go into more detail in the rest of this chapter.

Figure 5-1. High-level architecture of ADO.NET

A data reader is a new object providing fast, forward-only, and read-only access to data. This structure is similar to an ADO Recordset, which has server-side, forward-only, and read-only cursor types.

The DataSet class is analogous to a lightweight cache of a subset of the database from the data store for disconnected operations. It also allows reading and writing of data and schema in XML, and it is tightly integrated with XmlDocument, as you will see later.

The DataAdapter class serves as a higher-level abstraction of the connection and command classes. It enables you to load content from a data store into a DataSet and reconcile DataSet changes back

to the data store.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

5.2 ADO.NET Benefits

ADO.NET brings with it a number of benefits, which fall into the following categories:

Interoperability

The ability to communicate across heterogeneous environments.

Scalability

The ability to serve a growing number of clients without degrading system performance.

Productivity

The ability to quickly develop robust data access applications using ADO.NET's rich and extensible component object model.

Performance

An improvement over previous ADO versions due to the disconnected data model.

5.2.1 Interoperability

All communication involves data exchange, whether the communication between distributed components is through a request/response methodology or a message-based facility. Current distributed systems assume that the components involved in the communication are using the same protocol and data format. This assumption is too restrictive for a client base to expand across an enterprise or for multiple companies. Data-access layers should impose no such restrictions.

In current Microsoft Windows Distributed InterNet Applications (DNA) Architecture, application components pass data back and forth as ADO disconnected recordsets. The data-providing components, as well as the data-consuming components, are required to use the Component Object Model (COM). The *payload*, the actual content we are passing around, is packaged in a data format called Network Data Representation (NDR). These NDR packages are streamed between components.

There are two issues with current Windows DNA systems. The first is the requirement that both ends of the communication pipe have the COM library. The second issue is that it is difficult to set up and manage these communications across firewalls. If your middle-tier components are COM/DCOM-based and you are using them within your intranet, you are in good shape. To put it another way: if all your components use Microsoft technology, you're fine. With the advent of electronic commerce (e-commerce), however, enterprise applications must interoperate with more than just Microsoft-development shops. ADO must improve for cross-platform components to seamlessly share data, breaking away from the limitations of COM/DCOM.

ADO.NET addresses the common data-exchange limitation by using XML as its payload data format. Since XML is text-based and simple to parse, it's a good choice for a common, platform-independent, and transportable data format. Furthermore, because XML is nothing more than structured text, employing XML as the data format on top of the HTTP network protocol minimizes firewall-related problems. With ADO and its XML format, the clients do not have to know COM to de-serialize the

packaged data. All they need is an XML parser, which is readily available in many flavors on many different platforms. The data producers and consumers need only adhere to the XML schema to exchange data among themselves.

5.2.2 Scalability

In a client/server model, it is typical for a client to acquire and hold onto a connection to the server until all requests are fulfilled. While this solution works fine in small- to medium-scale applications, it is not scalable across a large enterprise. As soon as the number of clients reaches a certain threshold, the server becomes the bottleneck, as database connections eat up network and CPU resources. ADO.NET moves away from the client/server model by promoting the use of disconnected datasets. When a client requests some data, the data is retrieved, it's transferred to the client, and-as soon as possible-the connection is torn down. Since the connection between the client and the data source is short-lived, this technique allows more clients to request information from the server, thus solving the problem of limited connections.

You might think that setting up and tearing down connections is not a good idea since the cost of establishing a connection is usually high. This is a concern only in the absence of connection pooling. ADO.NET automatically keeps connections to a data source in a pool, so when an application thinks it is tearing down a connection, it's actually returning it to the resource pool. This allows connections to be reused, avoiding the cost of reconstructing new connections from scratch.

Working with data in this disconnected fashion is not new to ADO programmers. The disconnected recordset was introduced in early versions of ADO. However, in ADO, it is up to the developer to implement this feature, whereas in ADO.NET, data is disconnected by nature.

ADO.NET has enhanced its predecessor by growing out of the client/server model and into the distributed components model. By using disconnected datasets as the paradigm for data exchange, ADO.NET is much more scalable than its predecessors.

5.2.3 Productivity

ADO.NET's rich data access classes allow developers to boost their productivity. Current ADO developers should have no problems getting up to speed with the object model, because ADO.NET is a natural evolution of ADO. The core functionality remains the same. We still have the connection object, representing the pipeline through which commands are executed^[2] With ADO.NET, the functionality is factored and distributed to each object in the model-much better than in previous versions of ADO. For example, the connection object is responsible only for connecting to and disconnecting from the data source. In ADO.NET, we can no longer execute a query directly through the connection object. Although some developers might miss this ability, it is a step in the right direction for cohesion of component development.

^[2] Along with the familiar connection and command objects, ADO.NET introduces a number of new objects, such as DataSet and DataAdapter. All of these objects are discussed earlier in this chapter.

ADO.NET also boosts developers' productivity through extensibility. Because ADO.NET framework classes are managed code, developers can inherit and extend these classes to their custom needs. If you prefer not to do this low-level legwork, you can use the Visual Studio .NET data-design environment to generate these classes for you.

Visual Studio .NET is a great Rapid Application Development (RAD) tool for developing applications with ADO.NET. You can have the Component Designer generate ADO.NET typed DataSets. These typed DataSets are extended types, modeled for your data. You don't have to reference database fields by their names or indices but instead by their names as properties of typed objects. For example, instead of `oRow(customerName)`, you can use `oRow.customerName`. The generated code is much more readable, when compared to previous Microsoft code generators. In addition, these generated classes are type-safe, thus reducing the chances for errors and allowing compilers and the CLR to verify type usage.

In short, ADO.NET improves developers' productivity through its rich and extensible framework classes. These features are complemented by the rich toolsets for ADO.NET in Visual Studio .NET, which enable rapid application development.

5.2.4 Performance

Because ADO.NET is mainly about disconnected datasets, the system benefits from improved performance and scalability. The database server is no longer a bottleneck when the number of connection requests goes up. Data Providers in ADO.NET also enable implicit connection pooling, which reduces the time required to open a connection.

Previous marshaling of recordsets required type conversion to make sure that the data types were all COM-based. Since the disconnected dataset is in XML format, there is no need for this type conversion during transport, as opposed to dealing with data in Network Data Representation format.

With the ADO.NET architecture, the data providers can be optimized for better performance by talking directly to the database server. .NET Framework Data Provider for SQL Server is an example of this as we can see later in this chapter.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

5.3 Content Components

Content components encapsulate data. In previous ADO versions, the Recordset object represented such a component. The data contained by the recordset component is in the form of a table, consisting of columns and rows. In ADO.NET, the data encapsulated by the DataSet component is in the form of a relational database, consisting of tables and relationships. This is a major improvement in data-access technology. In this section, we provide a high-level survey of the core classes that make up the content components, including DataSet, DataTable, DataColumn, DataRow, DataView, and DataRelation.^[3]

^[3] The complete list of all classes can be found in the Microsoft .NET SDK.

5.3.1 DataSet

If you are familiar with ADO, you know that data is typically transferred between components in *recordsets*. The recordset contains data in a tabular form. Whether the recordset includes information from one or many tables in the database, the data is still returned in the form of rows and columns as if they were from a single table. ADO.NET allows for more than just a recordset to be shared between application components. This is one of the most important features of ADO.NET: we will be transferring a DataSet instead of a recordset.

The DataSet can be viewed as an in-memory view of the database. It can contain multiple DataTable and DataRelation objects.^[4] With previous versions of ADO, the closest you could get to this functionality was to exchange data with a chain of Recordset objects. When the client application receives this chained recordset, it can get to each of the recordsets through `NextRecordset()`; however, there is no way to describe the relationship between each of the recordsets in the chain. With ADO.NET, developers can navigate and manipulate the collection of tables and their relationships.

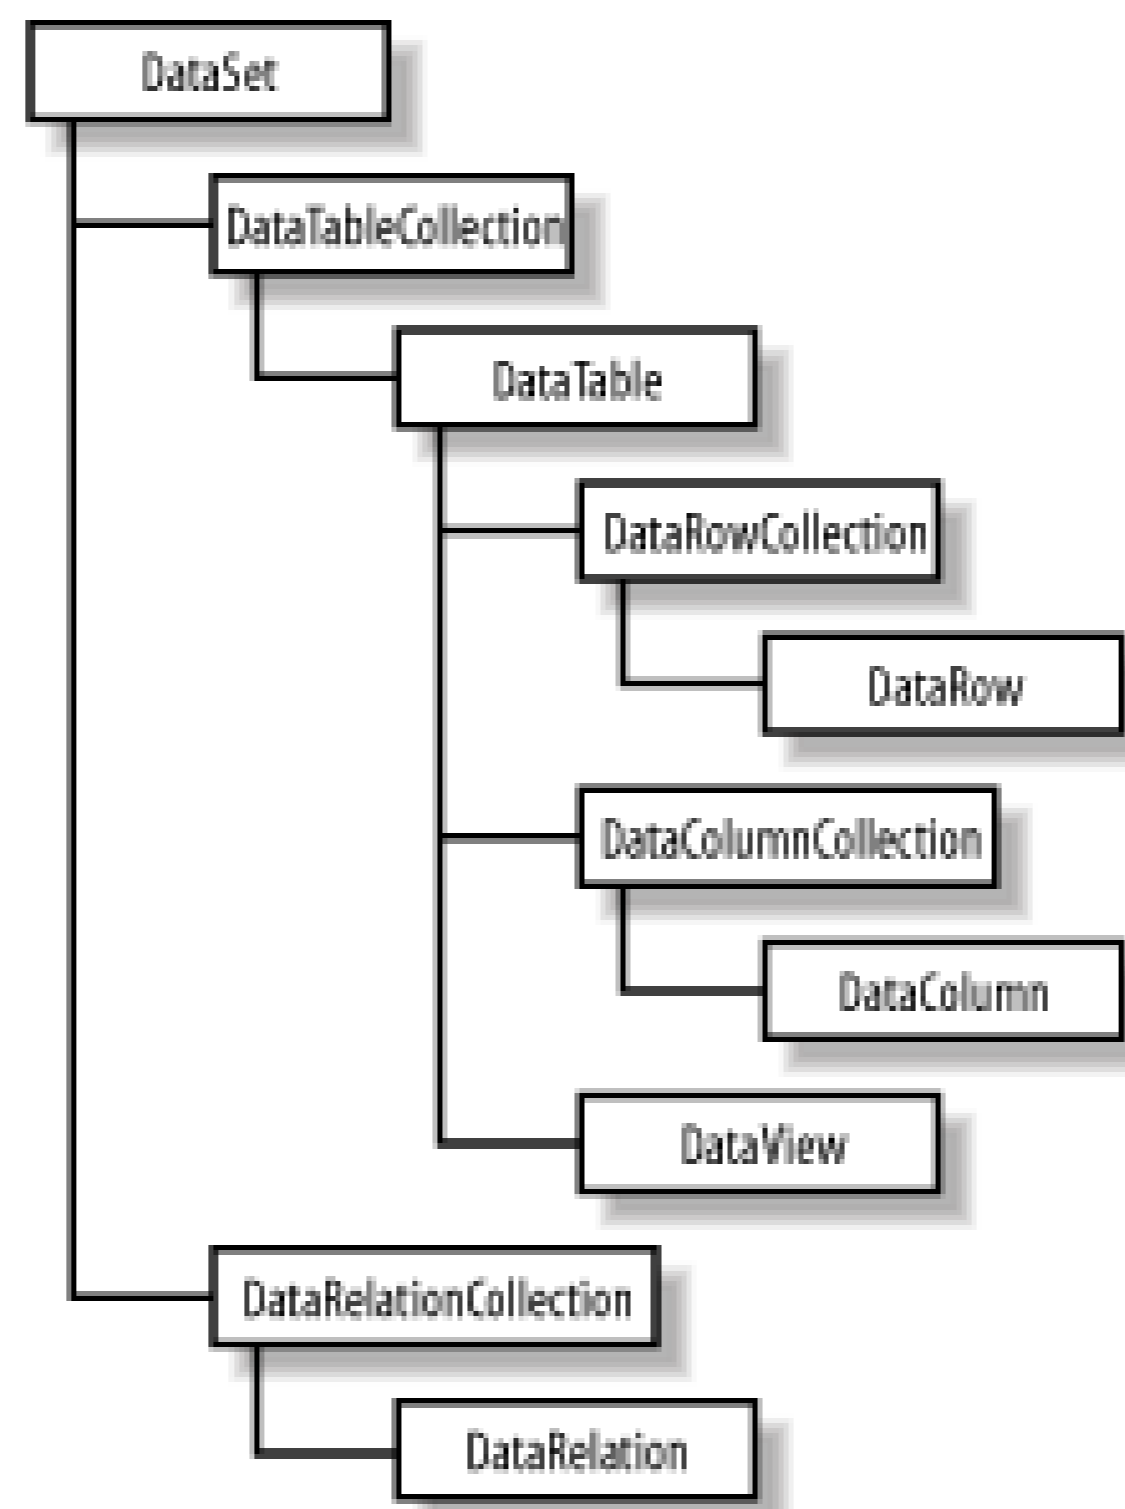
^[4] Furthermore, the DataSet can be persisted into disconnected XML datafiles so that your application can continue to work offline. More information on this topic will be presented in later sections.

As mentioned earlier, ADO.NET involves disconnected datasets because it is geared toward a distributed architecture. Since a DataSet is disconnected, it must provide a way to track changes to itself. The DataSet object provides a number of methods so that all data manipulation done to the DataSet can be easily reconciled with the actual database (or other data source) at a later time. They include: `HasChanges()`, `HasErrors`, `GetChanges()`, `AcceptChanges()`, and `RejectChanges()`. You can employ these methods to check for changes that have happened to the DataSet, obtain the modifications in the form of a changed DataSet, inspect the changes for errors, and then accept or reject the changes. If you want to communicate the changes to the data store back end (which is usually the case), you would ask the DataSet for an update.

The DataSet is intended to benefit enterprise web applications, which are disconnected by nature. You don't know that the data at the back end has changed until you have updated records you were editing or performed any other tasks that required data reconciliation with the database.

As depicted in [Figure 5-2](#), a DataSet contains two important collections. The first is the Tables (of type DataTableCollection), which holds a collection for all the tables belonging to a given DataSet. The second collection contains all the relationships between the tables, and it is appropriately named the Relations (of type DataRelationCollection).

Figure 5-2. Important ADO.NET data objects, including DataSet



5.3.1.1 Creating a DataSet: An example in C#

All the tables and relations inside the DataSet are exposed through its Tables and Relations properties, respectively. Normally, you obtain tables from some data sources such as SQL Server or other databases; however, we would like to show the nuts and bolts of the DataSet here first. The following block of C# code demonstrates how to create a DataSet dynamically that consists of two tables, Orders and OrderDetails, and a relationship between the two tables:

```

using System;
using System.Data;

// Class and method declarations omitted for brevity . . .

// Construct the DataSet object.
DataSet m_ds = new DataSet("DynamicDS");

// Add a new table named "Order" to m_ds's collection tables.
m_ds.Tables.Add ("Order");

// Add new columns to table "Order".
m_ds.Tables["Order"].Columns.Add("OrderID",
                                Type.GetType("System.Int32"));
m_ds.Tables["Order"].Columns.Add("CustomerFirstName",
                                Type.GetType("System.String"));
  
```

```

m_ds.Tables["Order"].Columns.Add("CustomerLastName",
                                Type.GetType("System.String"));
m_ds.Tables["Order"].Columns.Add("Date",
                                Type.GetType("System.DateTime"));

// Register the column "OrderID" as the primary key of table "Order".
DataColumn[] keys = new DataColumn[1];
keys[0] = m_ds.Tables["Order"].Columns["OrderID"];
m_ds.Tables["Order"].PrimaryKey = keys;

// Add a new table named "OrderDetail" to m_ds's collection of tables.
m_ds.Tables.Add("OrderDetail");

// Add new columns to table "OrderDetail".
m_ds.Tables["OrderDetail"].Columns.Add("fk_OrderID",
                                       Type.GetType("System.Int32"));
m_ds.Tables["OrderDetail"].Columns.Add("ProductCode",
                                       Type.GetType("System.String"));
m_ds.Tables["OrderDetail"].Columns.Add("Quantity",
                                       Type.GetType("System.Int32"));
m_ds.Tables["OrderDetail"].Columns.Add("Price",
                                       Type.GetType("System.Currency"));

// Get the DataColumn objects from two DataTable objects in a DataSet.
DataColumn parentCol = m_ds.Tables["Order"].Columns["OrderID"];
DataColumn childCol = m_ds.Tables["OrderDetail"].Columns["fk_OrderID"];

// Create and add the relation to the DataSet.
m_ds.Relations.Add(new DataRelation("Order_OrderDetail",
                                   parentCol,
                                   childCol));
m_ds.Relations["Order_OrderDetail"].Nested = true;

```

Let's highlight some important points in this block of code. After instantiating the DataSet object with the new operator, we add some tables with the Add method of the Tables object. We go through a similar process to add columns to each Table's Columns collection. Each of the added tables or columns can later be referenced by name. In order to assign the primary key for the Order table, we have to create the DataColumn array to hold one or more fields representing a key or a composite key. In this case, we have only a single key field, `OrderID`. We set the PrimaryKey property of the table to this array of key columns. For the relationship between the two tables, we first create the DataRelation called `Order_OrderDetail` with the two linking columns from the two tables, and then we add this DataRelation to the collection of relations of the DataSet. The last statement indicates that we want to represent the relationship between the Order and OrderDetail table as a nested structure. This makes dealing with these entities easier in XML.

The following block of C# code shows how to insert data into each of the two tables:

```

DataRow newRow;
newRow = m_ds.Tables["Order"].NewRow( );
newRow["OrderID"] = 101;
newRow["CustomerFirstName"] = "John";
newRow["CustomerLastName"] = "Doe";

```

```

newRow["Date"] = new DateTime(2001, 5, 1);
m_ds.Tables["Order"].Rows.Add(newRow);
newRow = m_ds.Tables["Order"].NewRow();
newRow["OrderID"] = 102;
newRow["CustomerFirstName"] = "Jane";
newRow["CustomerLastName"] = "Doe";
newRow["Date"] = new DateTime(2001, 4, 29);
m_ds.Tables["Order"].Rows.Add(newRow);

newRow = m_ds.Tables["OrderDetail"].NewRow();
newRow["fk_OrderID"] = 101;
newRow["ProductCode"] = "Item-100";
newRow["Quantity"] = 7;
newRow["Price"] = "59.95";
m_ds.Tables["OrderDetail"].Rows.Add(newRow);

newRow = m_ds.Tables["OrderDetail"].NewRow();
newRow["fk_OrderID"] = 101;
newRow["ProductCode"] = "Item-200";
newRow["Quantity"] = 1;
newRow["Price"] = "9.25";
m_ds.Tables["OrderDetail"].Rows.Add(newRow);

newRow = m_ds.Tables["OrderDetail"].NewRow();
newRow["fk_OrderID"] = 102;
newRow["ProductCode"] = "Item-200";
newRow["Quantity"] = 3;
newRow["Price"] = "9.25";
m_ds.Tables["OrderDetail"].Rows.Add(newRow);

```

Tables and Relations are important properties of DataSet. Not only do they describe the structure of the in-memory database, but the DataTables inside the collection also hold the content of the DataSet.

5.3.1.2 XML and tables sets

Now that you have a DataSet filled with tables and relationships, let's see how this DataSet helps in interoperability. XML is the answer. The DataSet has a number of methods that integrate DataSet tightly with XML, thus making it universally interoperable. These methods are `WriteXml()`, `WriteXmlSchema()`, `ReadXml()`, and `ReadXmlSchema()`.

`WriteXmlSchema()` dumps only the schema of the tables, including all tables and relationships between tables. `WriteXml()` can dump both the schema and table data as an XML encoded string. Both `WriteXmlSchema()` and `WriteXml()` accept a Stream, TextWriter, XmlWriter, or String representing a filename. `WriteXml()` accepts an `XmlWriteMode` as the second argument so you can optionally write the schema in addition to the data.

By default, `WriteXml()` writes only the data. To also write the schema, you will have to pass `XmlWriteMode.WriteSchema` as the second parameter to the call. You can also retrieve only the data portion of the XML by using the `XmlWriteMode.IgnoreSchema` property explicitly. Another mode that you can set is `XmlWriteMode.DiffGram`. In this DiffGram mode, the DataSet will be dumped out as

both the original data and changed data. More on this topic when we get to the `GetChanges()` method of the `DataSet`.

The `DataSet` object also provides methods to reconstruct itself from an XML document. Use `ReadXmlData()` for reading XML data documents, and `ReadXmlSchema()` for reading XML schema documents.

The following code creates an XML document from the previously created dataset:

```
// Dump the previously shown DataSet to
// the console (and to an XML file).
m_ds.WriteXml(Console.Out, XmlWriteMode.WriteSchema);
m_ds.WriteXml("DS_Orders.xml", XmlWriteMode.WriteSchema);

// Constructing a new DataSet object
DataSet ds2 = new DataSet("RestoredDS");
ds2.ReadXml("DS_Orders.xml");
```

Let's examine the resulting XML file and its representation of the dataset:

```
<?xml version="1.0" standalone="yes"?>
<DynamicDS>
  <xs:schema id="DynamicDS"
    xmlns=""
    xmlns:xs="http://www.w3.org/2001/XMLSchema"
    xmlns:msdata="urn:schemas-microsoft-com:xml-msdata">
    <xs:element name="DynamicDS" msdata:IsDataSet="true">

      <xs:complexType>
        <xs:choice maxOccurs="unbounded">

          <xs:element name="Order">
            <xs:complexType>
              <xs:sequence>
                <xs:element name="OrderID"
                  type="xs:int" />
                <xs:element name="CustomerFirstName"
                  type="xs:string" minOccurs="0" />
                <xs:element name="CustomerLastName"
                  type="xs:string" minOccurs="0" />
                <xs:element name="Date"
                  type="xs:dateTime" minOccurs="0" />

                <xs:element name="OrderDetail"
                  minOccurs="0" maxOccurs="unbounded">

                  <xs:complexType>
                    <xs:sequence>
                      <xs:element name="fk_OrderID"
                        type="xs:int" minOccurs="0" />
                      <xs:element name="ProductCode"
                        type="xs:string" minOccurs="0" />
                    </xs:sequence>
                  </xs:complexType>
                </xs:element>
              </xs:sequence>
            </xs:complexType>
          </xs:element>
        </xs:choice>
      </xs:complexType>
    </xs:element>
  </xs:schema>
</DynamicDS>
```

```

        <xs:element name="Quantity"
            type="xs:int" minOccurs="0" />
        <xs:element name="Price"
            msdata:DataType="System.Currency,
            mscorlib, Version=n.n.nnnn.n,
            Culture=neutral,
            PublicKeyToken=nnnnnnnnnnnnnnnnnn"
            type="xs:string" minOccurs="0" />
    </xs:sequence>
</xs:complexType>

</xs:element>

</xs:sequence>
</xs:complexType>
</xs:element>

</xs:choice>
</xs:complexType>

<xs:unique name="Constraint1"
    msdata:PrimaryKey="true">
    <xs:selector xpath="//Order" />
    <xs:field xpath="OrderID" />
</xs:unique>

<xs:keyref name="Order_OrderDetail"
    refer="Constraint1"
    msdata:IsNested="true">
    <xs:selector xpath="//OrderDetail" />
    <xs:field xpath="fk_OrderID" />
</xs:keyref>

</xs:element>
</xs:schema>

< . . . Data Portion . . . >

</DynamicDS>

```

The root element is named `DynamicDS` because that is the name of the dataset we created earlier. The `xsd:schema` tag contains all table and relationship definitions in this `DynamicDS` dataset. Because we've indicated that the relationship should be nested, the schema shows the `xsd:element` `OrderDetail` nested within the `xsd:element` `Order`. All columns are also represented as `xsd:elements`.

After the table definitions, the document holds definitions for various key types. The `xsd:unique` element is used with `msdata:PrimaryKey` for keys, as shown in the `xsd:unique` named `Constraint1`. The `msdata:PrimaryKey` attribute makes this a primary key, which has the added effect of enforcing uniqueness (every `OrderID` in the `Order` table must be unique).

The `xsd:keyref` element is used for foreign keys, as shown in the `Order_OrderDetail` key that

refers to the `Constraint1` key. This links the `OrderDetail` and `Order` tables where `OrderDetail.fk_OrderID = Order.OrderID`.

Let's now look at the data portion of the XML file:

```
<Order>
  <OrderID>101</OrderID>
  <CustomerFirstName>John</CustomerFirstName>
  <CustomerLastName>Doe</CustomerLastName>
  <Date>2001-05-01T00:00:00.0000000-04:00</Date>
  <OrderDetail>
    <fk_OrderID>101</fk_OrderID>
    <ProductCode>Item-100</ProductCode>
    <Quantity>7</Quantity>
    <Price>59.95</Price>
  </OrderDetail>
  <OrderDetail>
    <fk_OrderID>101</fk_OrderID>
    <ProductCode>Item-200</ProductCode>
    <Quantity>1</Quantity>
    <Price>9.25</Price>
  </OrderDetail>
</Order>
<Order>
  <OrderID>102</OrderID>
  <CustomerFirstName>Jane</CustomerFirstName>
  <CustomerLastName>Doe</CustomerLastName>
  <Date>2001-04-29T00:00:00.0000000-04:00</Date>
  <OrderDetail>
    <fk_OrderID>102</fk_OrderID>
    <ProductCode>Item-200</ProductCode>
    <Quantity>3</Quantity>
    <Price>9.25</Price>
  </OrderDetail>
</Order>
```

This part of the XML document is fairly self-explanatory. For each row of data in the `Order` table, we end up with one record of type `Order`. This is the same for the `OrderDetail` table. The `OrderDetail` that relates to a particular `Order` is nested inside the `Order` element.

Because the dataset is inherently disconnected from its source, changes to the data inside the dataset have to be tracked by the dataset itself. This is done through the following methods: `HasChanges()`, `GetChanges()`, and `Merge()`. The application can check the changes to the dataset and then ask the `DataAdapter` object to reconcile the changes with the data source through the `DataAdapter.Update()` method.

The following block of code demonstrates how to track and manage changes to a `DataSet`:

```
m_ds.AcceptChanges( );
/* Make a change to the data set. */
m_ds.Tables["OrderDetail"].Rows[0]["Quantity"] = 12;

if(m_ds.HasChanges( )){
```

```

/* Get a copy of the data set containing the changes. */
DataSet changeDS = m_ds.GetChanges( );

/* Dump the changed rows. */
changeDS.WriteXml("ChangedDS.xml" , XmlWriteMode.DiffGram);

/* Commit all changes. */
m_ds.AcceptChanges( );

}

```

Because we create this DataSet dynamically, we want to tell the DataSet to accept all changes made up to this point by first issuing an AcceptChange() call. Knowing that the DataSet should start tracking the changes again, we then change the quantity of one of the OrderDetail rows. Next, we ask the dataset for all the changes and dump it into a new dataset called `changeDS`. This dataset results in the following XML dump when using DiffGram mode. Notice that because OrderDetail is a child of Order, the change also includes the parent row:

```

<?xml version="1.0" standalone="yes"?>
<diffgr:diffgram xmlns:msdata="urn:schemas-microsoft-com:xml-msdata"
    xmlns:diffgr="urn:schemas-microsoft-com:xml-diffgram-v1">
  <DynamicDS>
    <Order diffgr:id="Order1" msdata:rowOrder="0">
      <OrderID>101</OrderID>
      <CustomerFirstName>John</CustomerFirstName>
      <CustomerLastName>Doe</CustomerLastName>
      <Date>2001-05-01T00:00:00.0000000-04:00</Date>
      <OrderDetail diffgr:id="OrderDetail1"
        msdata:rowOrder="0" diffgr:hasChanges="modified">
        <fk_OrderID>101</fk_OrderID>
        <ProductCode>Item-100</ProductCode>
        <Quantity>12</Quantity>
        <Price>59.95</Price>
      </OrderDetail>
    </Order>
  </DynamicDS>

  <diffgr:before>
    <OrderDetail diffgr:id="OrderDetail1" msdata:rowOrder="0">
      <fk_OrderID>101</fk_OrderID>
      <ProductCode>Item-100</ProductCode>
      <Quantity>7</Quantity>
      <Price>59.95</Price>
    </OrderDetail>
  </diffgr:before>

</diffgr:diffgram>

```

We would like to emphasize that the DataSet object is the most important construct in ADO.NET. Because DataSet does not tie to an underlying representation, such as SQL Server or Microsoft Access, it is extremely portable. Its data format is self-described in its schema, and its data is in pure

XML. A DataSet is self-contained regardless of how it was created, whether by reading data from a SQL Server, from Microsoft Access, from an external XML file, or even by being dynamically generated as we have seen in an earlier example. This portable XML-based entity-without a doubt-should be the new standard for data exchange.

Enough said about DataSet. Let's drill down from DataSet to DataTable.

5.3.2 DataTable

DataTable represents a table of data and, thus, contains a collection of DataColumn objects as a Columns property and a collection of DataRow objects as a Rows property. The Columns property provides the structure of the table, while the Rows property provides access to actual row data. Fields in the table are represented as DataColumn objects, and table records are represented as DataRow objects. Here is some sample code that dumps the name of each column as a row of headers, followed by each row of data:

```
/* Walk the DataTable and display all column headers
 * along with all data rows.
 */
DataTable myTable = m_ds.Tables["OrderDetail"];

/* Display all column names. */
foreach(DataColumn c in myTable.Columns) {
    Console.WriteLine(c.ColumnName + "\t");
}
Console.WriteLine(""); // Newline

/* Process each row. */
foreach(DataRow r in myTable.Rows) {

    /* Display each column. */
    foreach(DataColumn c in myTable.Columns) {
        Console.WriteLine(r[c] + "\t");
    }
    Console.WriteLine(""); // Newline
}
}
```

Here is the output of that code:

fk_OrderID	ProductCode	Quantity	Price
101	Item-100	12	59.95
101	Item-200	1	9.25
102	Item-200	3	9.25

Typically, a DataTable has one or more fields serving as a primary key. This functionality is exposed as the PrimaryKey property. Because the primary key might contain more than one field, this property is an array of DataColumn objects. We revisit this excerpt of code here to put things in context. Note that in this example, the primary key consists of only one field; hence, the array of size one.

```
// Register the column "OrderID" as the primary key of table "Order".
DataColumn[] keys = new DataColumn[1];
keys[0] = m_ds.Tables["Order"].Columns["OrderID"];
m_ds.Tables["Order"].PrimaryKey = keys;
```

5.3.2.1 Relations and constraints

Relations define how tables in a database relate to each other. The DataSet globally stores the collection of relations between tables in the Relations property; however, each of the tables participating in the relation also has to know about the relationship. ChildRelations and ParentRelations, two properties of the DataTable object, take care of this. ChildRelations enumerates all relations that this table participates in as a master table. ParentRelations, on the other hand, lists the relations in which this table acts as a slave table. We provide more information on the topic of relations when we explain the DataRelation object in an upcoming section of this chapter.

While we are on the topic of tables and relationships, it is important to understand how to set up constraint enforcements. There are two types of constraints that we can set up and enforce, UniqueConstraint and ForeignKeyConstraint. UniqueConstraint enforces the uniqueness of a field value for a table. ForeignKeyConstraint enforces rules on table relationships. For ForeignKeyConstraint, we can set up UpdateRule and DeleteRule to dictate how the application should behave upon performing update or delete on a row of data in the parent table.

[Table 5-1](#) shows the constraint settings and behavior of ForeignKeyConstraint rules.

Table 5-1. Constraint types and behaviors

Setting	Behavior
None	Nothing.
Cascade	Dependent rows (identified by foreign key) are deleted/ updated when parent row is deleted/updated.
SetDefault	Foreign keys in dependent rows are set to the default value when parent row is deleted.
SetNull	Foreign keys in dependent rows are set to null value when parent row is deleted.

Constraints are activated only when the EnforceConstraint property of the DataSet object is set to `true`.

The following block of code shows how we have altered the foreign key constraint between the `Order` and `OrderDetail` tables to allow cascading deletion:

```
m_ds.Relations["Order_OrderDetail"].ChildKeyConstraint.DeleteRule =
    Rule.Cascade;

m_ds.WriteXml("DS_BeforeCascadeDelete.xml");
m_ds.Tables["Order"].Rows[0].Delete();
m_ds.WriteXml("DS_AfterCascadeDelete.xml");
```

As the result of running this code, the DataSet is left with only one order (order 102), which contains one line item.

5.3.3 DataView

The DataView object is similar to a *view* in conventional database programming. We can create different customized views of a DataTable, each with different sorting and filtering criteria. Through these different views, we can traverse, search, and edit individual records. This ADO.NET concept is the closest to the old ADO recordset. In ADO.NET, DataView serves another important role—data binding to Windows Forms and Web Forms. We show the usage of DataView when we discuss data binding on Windows Forms and Web Forms in [Chapter 7](#) and [Chapter 8](#).

5.3.4 DataRelation

A DataSet object as a collection of DataTable objects alone is not useful enough. A collection of DataTable objects returned by a server component provides little improvement upon the chained recordset in previous versions of ADO. In order for your client application to make the most of the returned tables, you also need to return the relations between these DataTables. This is where the DataRelation object comes into play.

With DataRelation, you can define relationships between the DataTable objects. Client components can inspect an individual table or navigate the hierarchy of tables through these relationships. For example, you can find a particular row in a parent table and then traverse all dependent rows in a child table.

The DataRelation contains the parent table name, the child table name, the parent table column (primary key), and the child table column (foreign key).

Because it has multiple DataTables and DataRelations within the DataSet, ADO.NET allows for a much more flexible environment where consumers of the data can choose to use the data in whichever way they wish.

One example might be the need to display all information about a particular parent table and all of its dependent rows in a child table. You have ten rows in the parent table. Each of the rows in the parent table has ten dependent rows in the child table. Let's consider two approaches to getting this data to the data consumer. First, we will just use a join in the query string:

```
Select
    Order.CustomerFirstName, Order.CustomerLastName, Order.OrderDate,
    OrderDetail.ProductCode, OrderDetail.Quantity, OrderDetail.Price
from
    Order, OrderDetail
    where Order.OrderID = OrderDetail.fk_OrderID
```

The result set contains 100 rows, in which each group of ten rows contains duplicate information about the parent row.

A second approach is to retrieve the list of rows from the parent table first, which would be ten rows:

```
Select
    Order.OrderID,
    Order.CustomerFirstName, Order.CustomerLastName, Order.OrderDate
from
    Order
```

Then for each of the ten rows in the parent table, you would retrieve the dependent rows from the child table:

```
Select
    OrderDetail.ProductCode, OrderDetail.Quantity, OrderDetail.Price
from
    OrderDetail where fk_OrderID = thisOrderID
```

This second approach is less of a resource hog since there is no redundant data; however, you end up making 11 round-trips (one time for the parent table, and 10 times for each parent of the child table).

It's better to get the parent table, the child table, and the relation between them using one round-trip, without all the redundant data. This is one of the biggest benefits that DataSet brings. The following block of code demonstrates the power of having tables and relationships:

```
/*
 * Given an order id, display a single order.
 */
public static void DisplaySingleOrder(DataSet m_ds, int iOrderID) {

    Decimal runningTotal = 0;
    Decimal lineTotal = 0;
    Decimal dPrice = 0;
    int iQty = 0;

    DataTable oTable = m_ds.Tables["Order"];

    // Find an order from the Order table.
    DataRow oRow = oTable.Rows.Find(iOrderID);

    /* Navigate to the OrderDetail table
     * through the Order_Details relationship.
     */
    DataRow[] arrRows = oRow.GetChildRows("Order_OrderDetail");

    /* Display the order information. */
    Console.WriteLine ("Order: {0}", iOrderID);
    Console.WriteLine ("Name: {0} {1}",
        oRow["CustomerFirstName"].ToString( ),
        oRow["CustomerLastName"].ToString( ));
    Console.WriteLine ("Date: {0}", oRow["Date"].ToString( ));
    Console.WriteLine("-----");

    /*
```

```

    * Display and calculate line total for each item.
    */
    for(int i = 0; i < arrRows.Length; i++) {

        foreach(DataColumn myColumn in m_ds.Tables["OrderDetail"].Columns)
        {
            Console.Write(arrRows[i][myColumn] + " ");

        }

        iQty = System.Int32.Parse(arrRows[i]["Quantity"].ToString( ));
        dPrice = System.Decimal.Parse(arrRows[i]["Price"].ToString( ));

        lineTotal = iQty * dPrice;
        Console.WriteLine("{0}", lineTotal);

        /* Keep a running total. */
        runningTotal += lineTotal;
    }

    /* Display the total of the order. */
    Console.WriteLine("Total: {0}", runningTotal);
}

```

DisplaySingleOrder finds a single row in the Order table with a given order ID. Once this row is found we ask the row for an array of dependent rows from the OrderDetail table according to the `Order_OrderDetail` relationship. With the returned array of DataRows, we then proceed to display all fields in the row. We also calculate the `lineTotal` value based on the quantity ordered and the price of the item, as well as keeping a `runningTotal` for the whole order. The following shows the output from the DisplaySingleOrder function:

```

Order: 101
Name: John Doe
Date: 5/1/2001 12:00:00 AM
-----
101 Item-100 12 59.95 719.4
101 Item-200 1 9.25 9.25
Total: 728.65

```

[\[Team LiB \]](#)

[\[Team LiB \]](#)

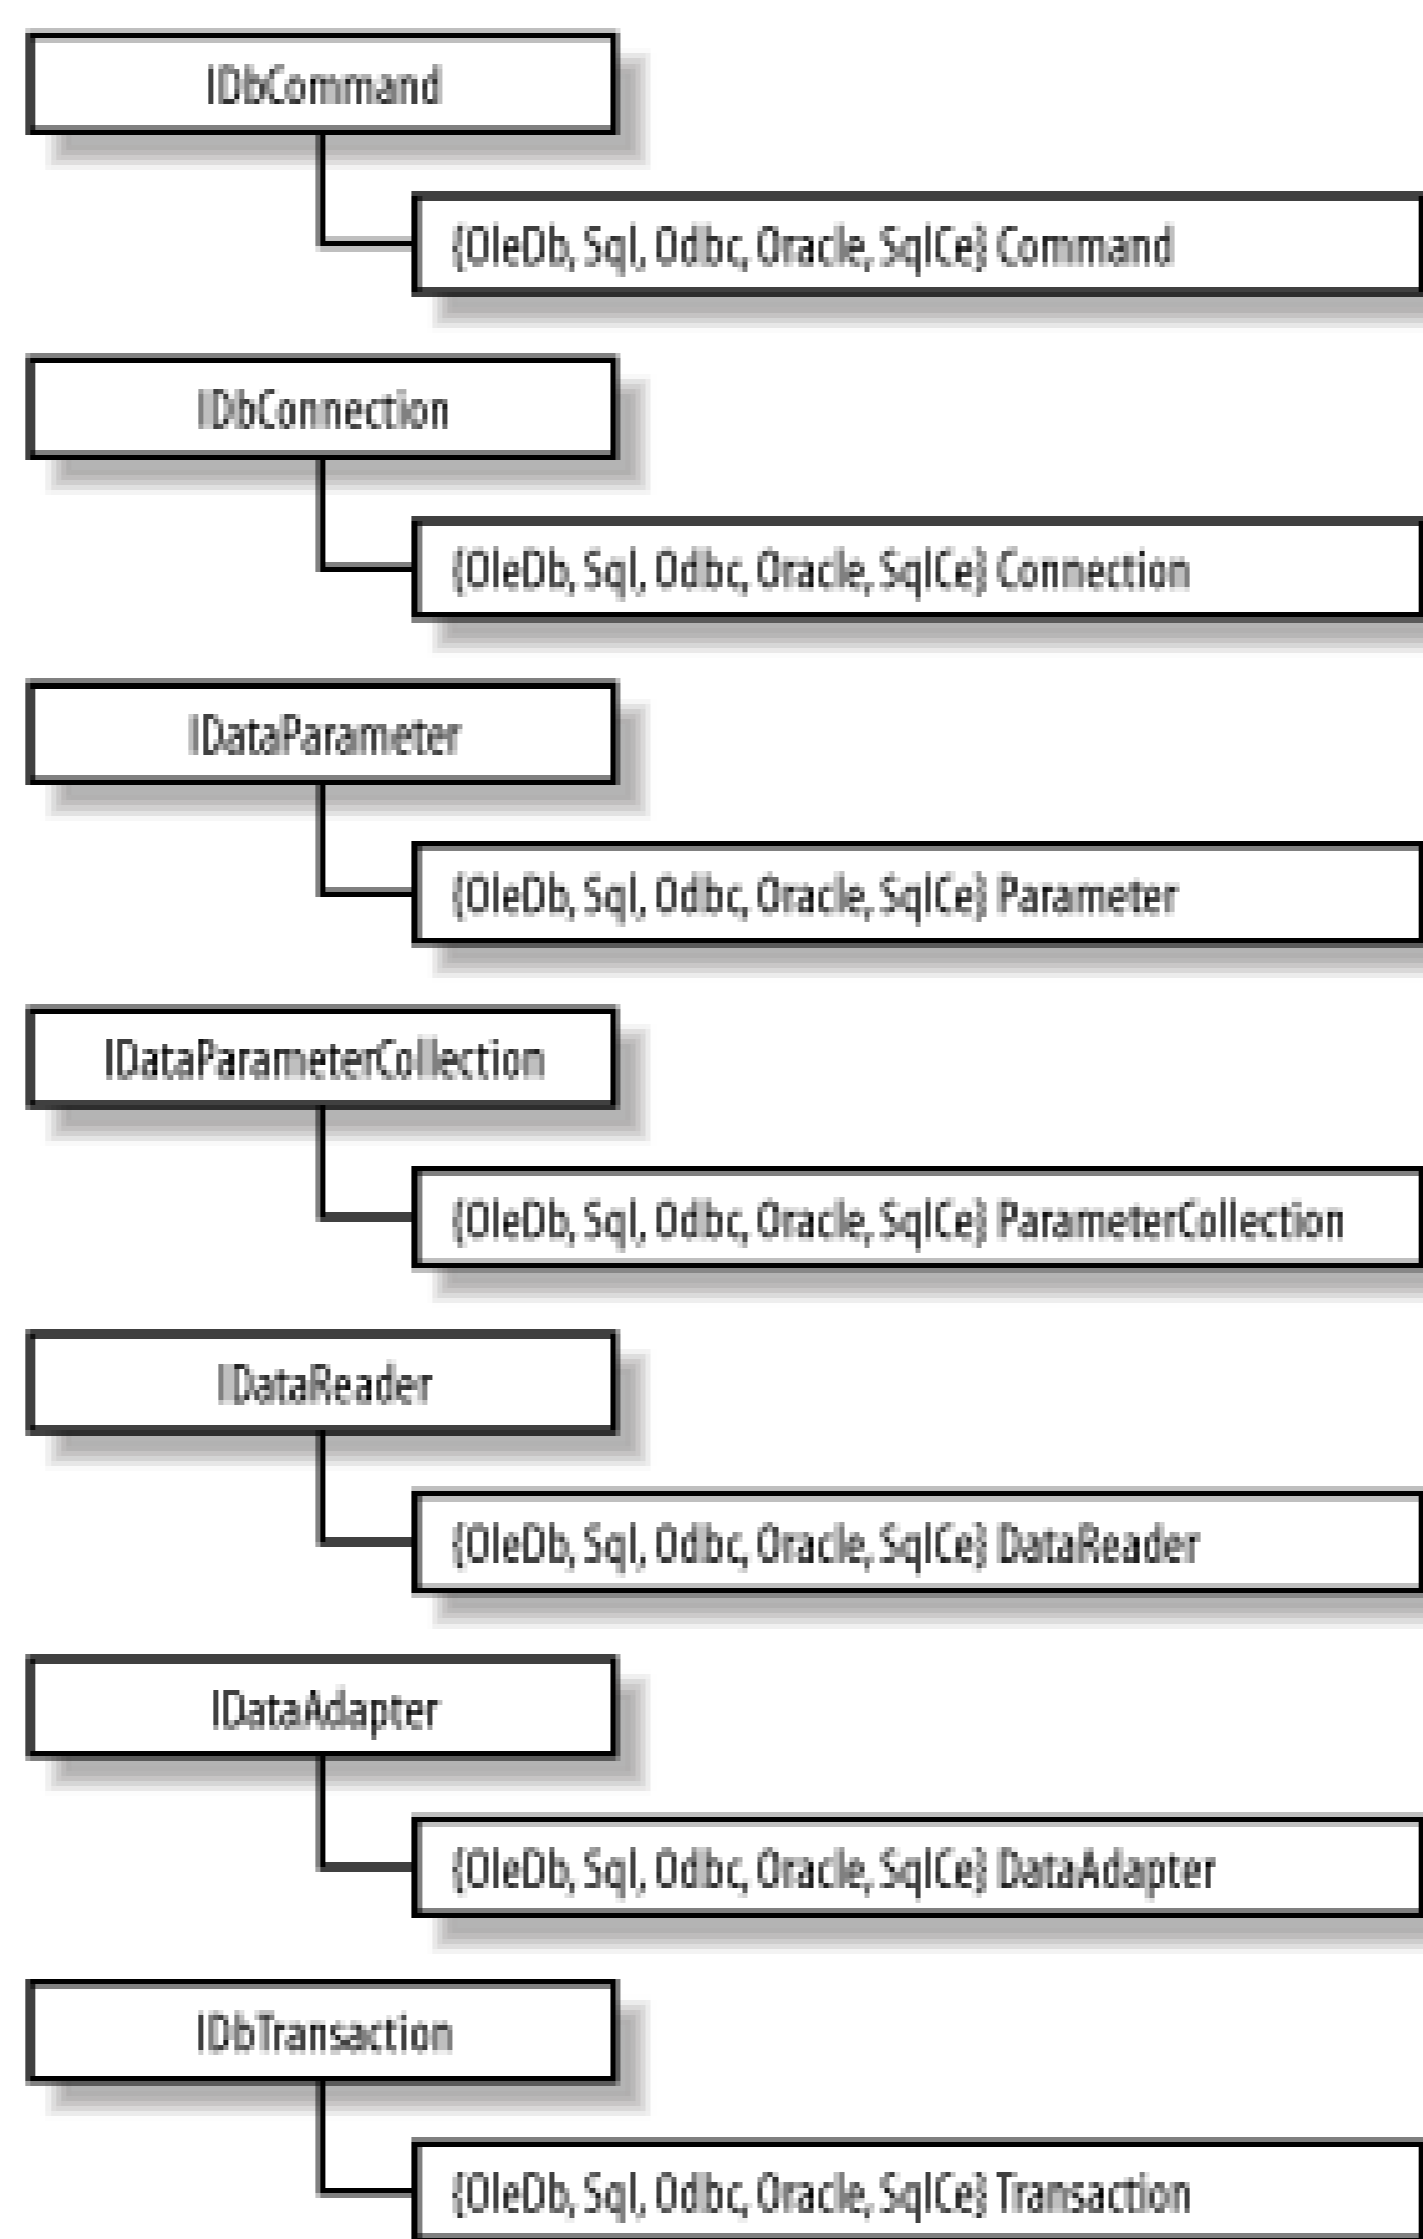
5.4 .NET Framework Data Providers

.NET Framework Data Provider (a.k.a. Managed Provider) is a term used for a group of .NET components that implement a fixed set of functionality put forth by the ADO.NET architecture. This enforces a common interface for accessing data. In order to build our own data provider, we must provide our own implementation of `System.Data.Common.DbDataAdapter` objects and implement interfaces such as `IDbCommand`, `IDbConnection`, and `IDataReader`. We are not building our own data provider here; [\[5\]](#) however, we do dive into each of these classes and interfaces in this section.

[5] There is a reference implementation of a .NET Framework Data Provider included in the .NET Framework documentation for any other type of data. In the near future, we are sure that the list of .NET Framework Data Providers will grow to cover even more different data sources.

Most of the time, developers don't have to know how to implement data providers, even though this might increase their productivity with regard to ADO.NET. Understanding how to use the stock data providers alone is sufficient to develop your enterprise application. Microsoft provides the following data providers in its current release of ADO.NET: OLE DB and SQL (Version 1 of .NET Framework) and ODBC, Oracle, SQL CE. The OLE DB data provider comes with `OleDbConnection`, `OleDbCommand`, `OleDbParameter`, and `OleDbDataReader`. The SQL Server data provider comes with a similar set of objects, whose names start with `SqlClient` instead of `OleDb`, and so on, as illustrated in [Figure 5-3](#). The implementation of this core function set for data providers is packaged in the `System.Data` namespace. The assemblies are: `System.Data.Odbc`, `System.Data.OleDb`, `System.Data.OracleClient`, `System.Data.SqlClient`, and `System.Data.SqlServerCe`.

Figure 5-3. Data provider framework class hierarchy



All of the included data providers implement a set of interfaces that access the appropriate data store. The OLE DB provider relies on OLE DB as an access layer to a broad variety of data sources, including Microsoft SQL Server. For performance reasons, the SQL data provider uses a proprietary protocol to communicate directly with SQL Server. In Version 1.1 of the .NET framework, ODBC, Oracle, and SQL CE data providers are added to provide better performance for these data store. Regardless of how the data is obtained, the resulting dataset remains the same. This clean separation of data providers and the XML-based dataset helps ADO.NET achieve portable data.

[Figure 5-3](#) shows the base classes and the all implementations of data provider. Because all data providers, adhere to a fixed, common set of interfaces (IDbCommand, IDbConnection, IDataParameterCollection, IDataReader, and IDataAdapter), you can easily adapt your application to switch data providers as the need arises.

5.4.1 Connection

All Connection classes implement `System.Data.IDbConnection` and, thus, inherit properties such as the connection string and the state of the connection. They implement the core set of methods specified by `IDbConnection`, including `Open` and `Close`.

Unlike the ADO Connection object, transaction support for the ADO.NET connection object has been moved to a Transaction object (such as `OleDbTransaction` and `SqlTransaction`). The reason for this is that we cannot assume that all data providers implement transaction the same way, so it's better for

the Connection object not to have transaction-related functionalities. To create a new transaction, execute the `BeginTransaction()` method of the Connection object. This returns an `IDbTransaction` implementation that supports transaction-oriented functionality such as `Commit` and `Rollback`. The `SqlTransaction` is currently the only provider that also supports saving checkpoints so that we can rollback to a specific checkpoint instead of rolling back the whole transaction. Again, if you examine the list of methods that any Connection class (such as `OleDbConnection` and `SqlConnection`) supports, you will find that the functionality is very much the same as the old ADO Connection object's. However, none of the Connection classes allows SQL statements or provider-specific text statements to be executed directly any more. In other words, `Execute()` is no longer supported by the Connection object. This is a better way for distributing functionality between classes. All execution is done through the Command object, which is discussed in the next section along with how to initiate a connection.

5.4.2 The Command and Data Reader Objects

Fortunately for ADO developers, ADO.NET's Command objects behave like ADO's Command object; however, the Command objects are the only way we can make execution requests to insert, update, and delete data in ADO.NET. This makes it easier to learn the object model. Developers are not faced with as many ways of doing the same things, as in the case (with ADO) of whether to execute the query through a Connection, Command, or even a Recordset object.

5.4.2.1 Command execution

All commands are associated with a connection object through the Command's `Connection` property. Think of the connection object as the pipeline between the data-reading component and the database back end. In order to execute a command, the active connection has to be opened. The command object also accepts parameters to execute a stored procedure at the back end. The top left of [Figure 5-5](#) shows the relationships between command, connection, and parameters objects.

There are two types of execution. The first type is a query command, which returns an `IDataReader` implementation. It is implemented by the `ExecuteReader()` method. The second type of command typically performs an update, insert, or deletion of rows in a database table. This type of execution is implemented by the `ExecuteNonQuery()` method.

One of the main differences between ADO.NET's Command objects and ADO's Command object is the return data. In ADO, the result of executing a query command is a recordset, which contains the return data in tabular form.^[6] In ADO.NET, however, recordsets are no longer supported. The result of executing a query command is now a data reader object (see the following section). This data reader object can be an `OleDbDataReader` for OLE DB, `SqlDataReader` for SQL Server (as of v.1 of .NET Framework), or any class implementing the `IDataReader` for custom reading needs. Once you've obtained a valid data reader object, you can perform a `Read` operation on it to get to your data.

[6] Disconnected record set.

Employing the command, connection, and data reader objects is a low-level, direct way to work with the data provider. As you will find out a little later, the data adapter encapsulates all this low-level plumbing as a more direct way to get the data from the data source to your disconnected dataset.

5.4.2.2 The data reader object

The data reader is a brand new concept to ADO developers, but it is straightforward. A data reader is similar to a stream object in object-oriented programming (OOP). If you need to access records in a forward-only, sequential order, use a data reader because it is very efficient. Since this is a server-side cursor, the connection to the server is open throughout the reading of data. Because of this continually open connection, we recommend that you exercise this option with care and not have the data reader linger around longer than it should. Otherwise, it might affect the scalability of your application.

The following code demonstrates basic use of `OleDbConnection`, `OleDbCommand`, and `OleDbDataReader`. Though we're using the OLE DB data provider here, the connection string is identical to the one we used earlier for ADO:[\[7\]](#)

[7] In addition, you can create a `Command` object from the current connection by using this instead: `oCmd = oConn.CreateCommand();`

```
using System;
using System.Data;
using System.Data.OleDb;

public class pubsdemo {

    public static void Main( ) {

        /* An OLE DB connection string. */
        String sConn =
            "provider=sqloledb;server=(local);database=pubs; Integrated Security=SSPI";

        /* An SQL statement. */
        String sSQL = "select au_fname, au_lname, phone from authors";

        /* Create and open a new connection. */
        OleDbConnection oConn = new OleDbConnection(sConn);
        oConn.Open( );

        /* Create a new command and execute the SQL statement. */
        OleDbCommand oCmd = new OleDbCommand(sSQL, oConn);
        OleDbDataReader oReader = oCmd.ExecuteReader( );

        /* Find the index of the columns we're interested in. */
        int idxFirstName = oReader.GetOrdinal("au_fname");
        int idxLastName = oReader.GetOrdinal("au_lname");
        int idxPhone = oReader.GetOrdinal("phone");

        /* Retrieve and display each column using their column index. */
        while(oReader.Read( )) {
            Console.WriteLine("{0} {1} {2}",
                oReader.GetValue(idxFirstName),
                oReader.GetValue(idxLastName),
                oReader.GetValue(idxPhone));
        }
    }
}
```

```

    }
}

```

The code opens a connection to the local SQL Server (using integrated security)^[8] and issues a query for first name, last name, and phone number from the authors table in the pubs database.

[8] Please be aware that database connection pooling relies on the uniqueness of the connection strings. When using the integrated security model of SQL Server, if you make the data access code run under the security context of each of the logged-in users, database connection pooling will suffer. You must create a small set of Windows accounts to overcome this problem; we don't discuss security in great depth in this book, due to its compact size.

If you don't have the pubs database installed on your system, you can load and run *instpubs.sql* in Query Analyzer (*instpubs.sql* can be found under the *MSSQL\Install* directory on your machine). For those that install the VS.NET Quickstart examples, change the server parameter of the connection string to `server=(local)\NetSDK` because the Quickstart examples installation lays down the NetSDK SQL Server instance that also include the infamous Pubs database. The following example uses `SqlClient` to get the same information. This time, instead of obtaining the indices for the column and getting the values based on the indices, this example indexes the column directly using the column names:

```

using System;
using System.Data;
using System.Data.SqlClient;

public class pubsdemo {

    public static void Main( ) {

        /* A SQL Server connection string. */
        String sConn = "server=(local);database=pubs;Integrated Security=SSPI";

        /* An SQL statement. */
        String sSQL = "select au_fname, au_lname, phone from authors";

        /* Create and open a new connection. */
        SqlConnection oConn = new SqlConnection(sConn);
        oConn.Open( );

        /* Create a new command and execute the SQL statement. */
        SqlCommand oCmd = new SqlCommand(sSQL, oConn);
        SqlDataReader oReader = oCmd.ExecuteReader( );

        /* Retrieve and display each column using the column names. */
        while(oReader.Read( )) {
            Console.WriteLine("{0} {1} {2}",
                oReader["au_fname"],
                oReader["au_lname"],
                oReader["phone"]);
        }
    }
}

```

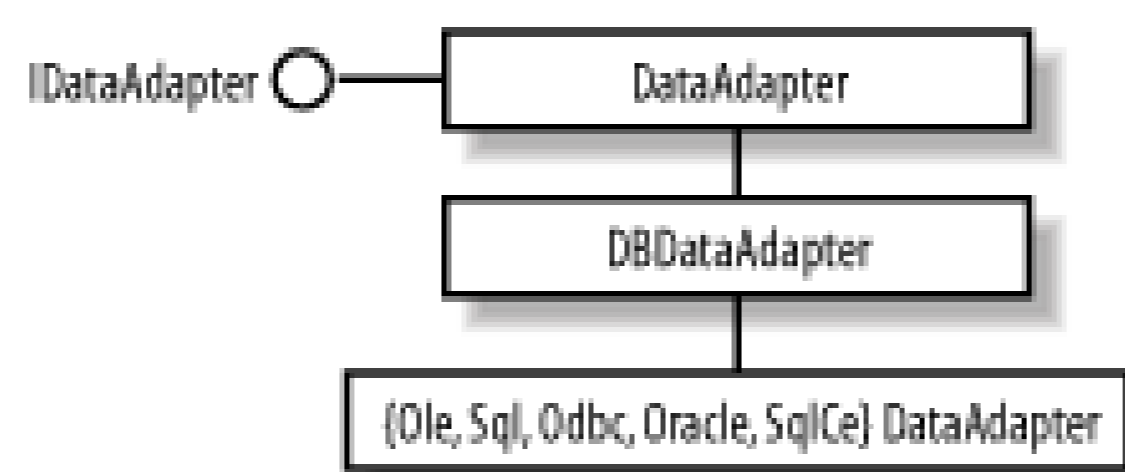
}

We leave the example code utilizing other data providers to the readers as an exercise.

5.4.3 The DataAdapter Object

Along with the introduction of data reader, ADO.NET also brings the DataAdapter object, which acts as the bridge between the data source and the disconnected DataSet. It contains a connection and a number of commands for retrieving the data from the data store into one DataTable in the DataSet and updating the data in the data store with the changes currently cached in the DataSet. Although each DataAdapter maps only one DataTable in the DataSet, you can have multiple adapters to fill the DataSet object with multiple DataTables. The class hierarchy of DataAdapter is shown in [Figure 5-4](#). All Data Adapters are derived from DbDataAdapter, which in turn is derived from the DataAdapter abstract class. This DataAdapter abstract class implements the IDataAdapter interface, which specifies that it supports Fill and Update. IDataAdapter is specified in the System.Data namespace, and is the DataSet itself.

Figure 5-4. DataSetCommand class hierarchy

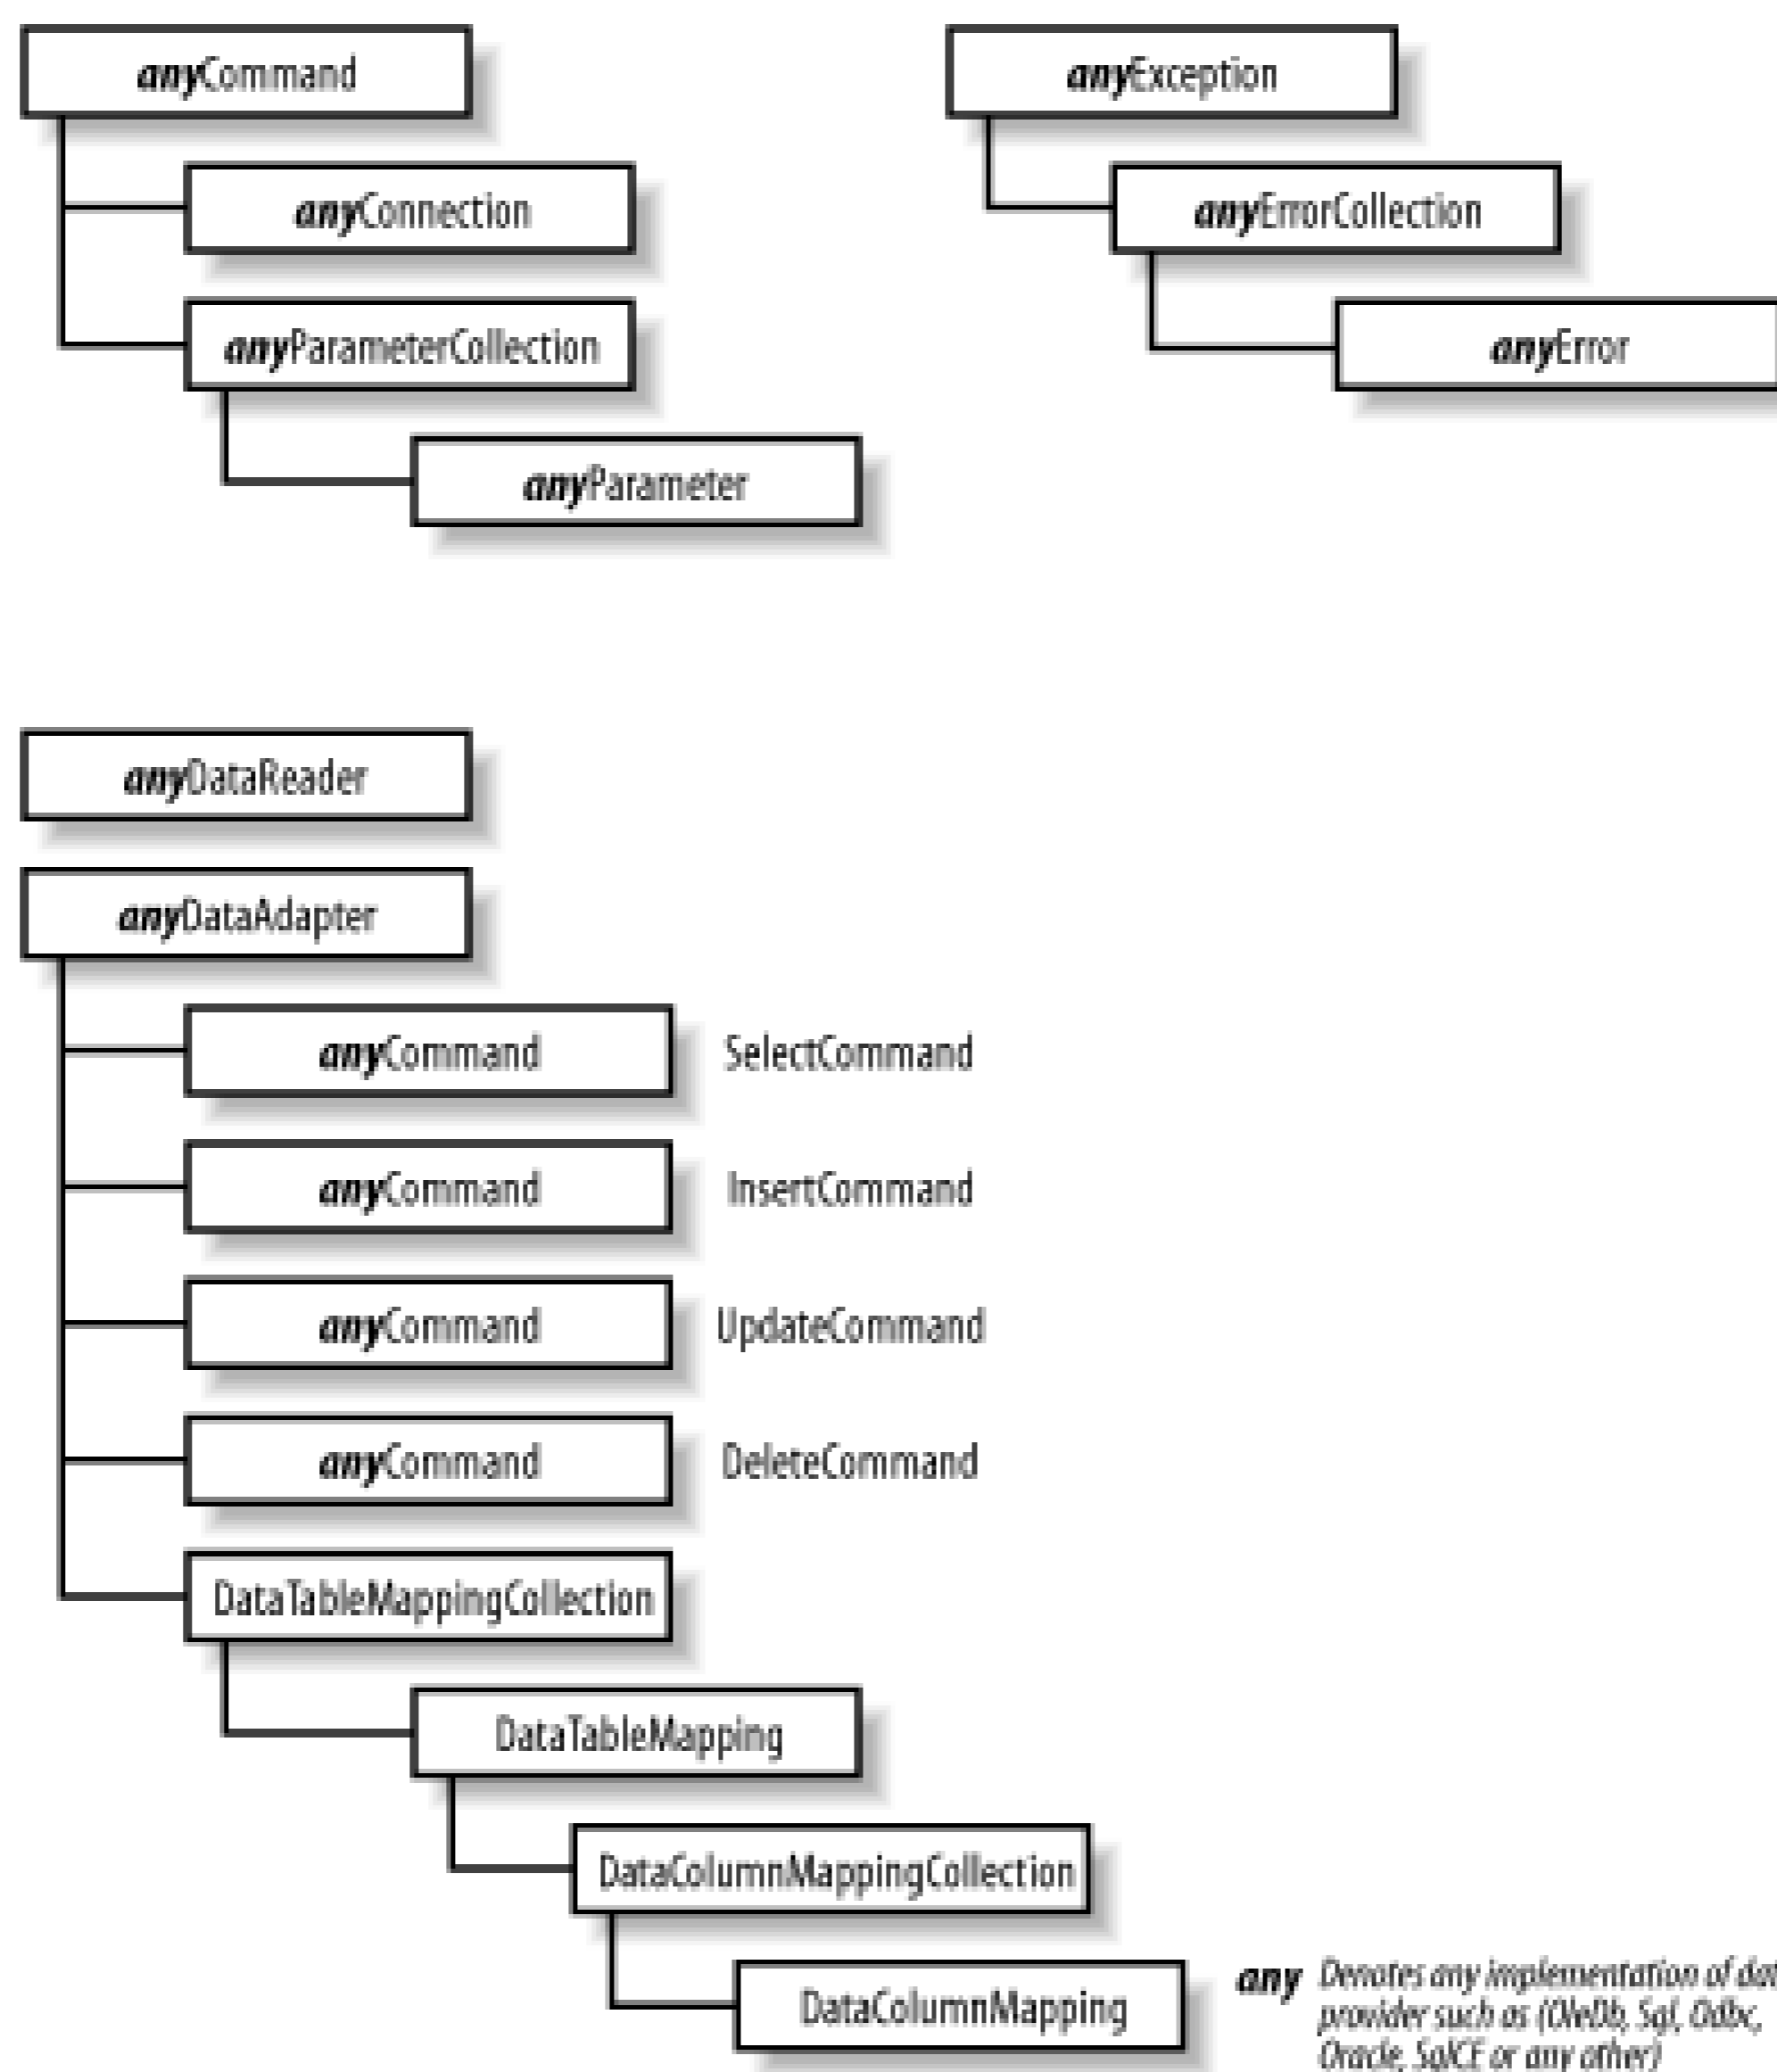


The data adapter can fill a DataSet with rows and update the data source when you make changes to the dataset. For example, you can use OleDbDataAdapter to move data from an OLE DB provider into a DataSet using the OleDbDataAdapter.Fill() method. Then you can modify the DataSet and commit the changes you made to the underlying database using the OleDbDataAdapterUpdate() method. These adapters act as the middleman bridging the data between the database back end and the disconnected DataSet.

For data retrieval, a data adapter uses the SQL **SELECT** command (exposed as the SelectCommand property). This **SELECT** command is used in the implementation of the IDataAdapter interface's Fill method. For updating data, a data adapter uses the SQL **UPDATE**, **INSERT**, and **DELETE** commands (exposed as the UpdateCommand, InsertCommand, and DeleteCommand properties).

Along with the Fill and Update methods from DbDataAdapter class, All data adapters also inherit the TableMappings property, a collection of TableMapping objects that enable the mapping of actual database column names to user-friendly column names. This further isolates the DataSet from the source where the actual data comes from. Even table names and column names can be mapped to more readable names, making it easier to use the DataSet. The application developer can be more productive at what he does best, which is to implement business logic and not to decipher cryptic database column names. [Figure 5-5](#) shows the relationship between data provider components.

Figure 5-5. Data adapter and supporting classes



Out of the four commands in the IDbDataAdapter object, only the **SELECT** command is required. The rest of the commands are optional since they can be generated automatically by the system. However, the auto-generation of these commands only works when certain conditions are met. For example, if your data adapter fills the data set from some database view that includes more than one table, you will have to explicitly define all four commands. Another example is when your adapter does not return key fields from the table, the system won't be able to generate the insert, update, or delete command. A typical usage of the data adapter involves the following steps:

- Create a data-adapter object.
- Set up the query string for the internal SelectCommand object.
- Set up the connection string for the SelectCommand's Connection object.
- Set up the InsertCommand, UpdateCommand, or DeleteCommand query strings and connections (Recommended).
- Call Fill() to fill the given dataset with the results from the query string.
- Make changes and call the adapter's Update() method with the changed DataSet (Optional).

The following block of code demonstrates these steps:

```

static DataSet GenerateDS( ) {

    /* Create the DataSet object. */
    DataSet ds = new DataSet("DBDataSet");
    String sConn =
        "provider=SQLOLEDB;server=(local);database=pubs; Integrated Security=SSPI ";

    /* Create the DataSet adapters. */
    OleDbDataAdapter dsAdapter1 =
        new OleDbDataAdapter("select * from authors", sConn);

    OleDbDataAdapter dsAdapter2 =
        new OleDbDataAdapter("select * from titles", sConn);

    OleDbDataAdapter dsAdapter3 =
        new OleDbDataAdapter("select * from titleauthor", sConn);

    /* Fill the data set with three tables. */
    dsAdapter1.Fill(ds, "authors");
    dsAdapter2.Fill(ds, "titles");
    dsAdapter3.Fill(ds, "titleauthor");

    // Add the two relations between the three tables. */
    ds.Relations.Add("authors2titleauthor",
        ds.Tables["authors"].Columns["au_id"],
        ds.Tables["titleauthor"].Columns["au_id"]);

    ds.Relations.Add("titles2titleauthor",
        ds.Tables["titles"].Columns["title_id"],
        ds.Tables["titleauthor"].Columns["title_id"]);

    // Return the DataSet.
    return ds;
}

```

This is a demonstration of constructing a dataset with three tables from the sample *pubs* database. The DataSet also contains two relationships that tie the three tables together. Let's take a look at the dataset in XML by trying out the next couple lines of code:

```

DataSet ds = GenerateDS( );
ds.WriteXml("DBDataSet.xml", XmlWriteMode.WriteSchema);

```

The content of *DBDataSet.xml* (with some omission for brevity) is shown next:

```

<?xml version="1.0" standalone="yes"?>
<DBDataSet>
  <xsd:schema id="DBDataSet" targetNamespace="" xmlns=""
    xmlns:xsd="http://www.w3.org/2001/XMLSchema"
    xmlns:msdata="urn:schemas-microsoft-com:xml-msdata">
    <xsd:element name="DBDataSet" msdata:IsDataSet="true">
      <xsd:complexType>

```

```

<xsd:choice maxOccurs="unbounded">
  <xsd:element name="authors">
    <xsd:complexType>
      <xsd:sequence>
        <!-- columns simplified for brevity -->
        <xsd:element name="au_id" type="xsd:string" />
        <xsd:element name="au_lname" type="xsd:string" />
        <xsd:element name="au_fname" type="xsd:string" />
        <xsd:element name="phone" type="xsd:string" />
        <xsd:element name="address" type="xsd:string" />
        <xsd:element name="city" type="xsd:string" />
        <xsd:element name="state" type="xsd:string" />
        <xsd:element name="zip" type="xsd:string" />
        <xsd:element name="contract" type="xsd:boolean" />
      </xsd:sequence>
    </xsd:complexType>
  </xsd:element>

  <!-- titles and titleauthor omitted for brevity -->

</xsd:choice>
</xsd:complexType>

<xsd:unique name="Constraint1">
  <xsd:selector xpath="//authors" />
  <xsd:field xpath="au_id" />
</xsd:unique>

<xsd:unique name="titles_Constraint1"
  msdata:ConstraintName="Constraint1">
  <xsd:selector xpath="//titles" />
  <xsd:field xpath="title_id" />
</xsd:unique>

<xsd:keyref name="titles2titleauthor"
  refer="titles_Constraint1">
  <xsd:selector xpath="//titleauthor" />
  <xsd:field xpath="title_id" />
</xsd:keyref>

<xsd:keyref name="authors2titleauthor"
  refer="Constraint1">
  <xsd:selector xpath="//titleauthor" />
  <xsd:field xpath="au_id" />
</xsd:keyref>

</xsd:element>
</xsd:schema>

<!-- Most rows removed for brevity -->

```

```

<authors>
  <au_id>899-46-2035</au_id>
  <au_lname>Ringer</au_lname>
  <au_fname>Anne</au_fname>
  <phone>801 826-0752</phone>
  <address>67 Seventh Av.</address>
  <city>Salt Lake City</city>
  <state>UT</state>
  <zip>84152</zip>
  <contract>>true</contract>
</authors>

<titles>
  <title_id>PS2091</title_id>
  <title>Is Anger the Enemy?</title>
  <type>psychology </type>
  <pub_id>0736</pub_id>
  <price>10.95</price>
  <advance>2275</advance>
  <royalty>12</royalty>
  <ytd_sales>2045</ytd_sales>
  <notes>Carefully researched study of the effects of strong
emotions on the body. Metabolic charts included.</notes>
  <pubdate>1991-06-15T00:00:00.0000</pubdate>
</titles>
  <title_id>MC3021</title_id>
  <title>The Gourmet Microwave</title>
  <type>mod_cook</type>
  <pub_id>0877</pub_id>
  <price>2.99</price>
  <advance>15000</advance>
  <royalty>24</royalty>
  <ytd_sales>22246</ytd_sales>
  <notes>Traditional French gourmet recipes adapted for modern
microwave cooking.</notes>
  <pubdate>1991-06-18T00:00:00.0000</pubdate>
</titles>

<titleauthor>
  <au_id>899-46-2035</au_id>
  <title_id>MC3021</title_id>
  <au_ord>2</au_ord>
  <royaltyper>25</royaltyper>
</titleauthor>
<titleauthor>
  <au_id>899-46-2035</au_id>
  <title_id>PS2091</title_id>
  <au_ord>2</au_ord>
  <royaltyper>50</royaltyper>
</titleauthor>

</DBDataSet>

```

The tables are represented as `<xsd:element name="table name"> . . . </xsd:element>` tag pairs that contain column definitions. In addition to one `xsd:element` for each table, we have one `xsd:unique` for each key and one `xsd:keyref` for each relationship. The `xsd:unique` specifies the key of the parent table in a relationship. The tag `xsd:keyref` is used for child tables in a relationship. This `xsd:keyref` serves as the foreign key and refers to the key in the parent table.

For brevity, we've stripped down the data portion of the XML to contain just one author, `Anne Ringer`, and two books she authored.

We can have many different DataAdapters populating the DataSet. Each of these DataAdapters can be going against a completely different data source or data server. In other words, you can construct a DataSet object filled with data that is distributed across multiple servers. In the previous example, we have three different DataAdapters; however, all of them are going to the same server.

[\[Team LiB \]](#)

[Team LiB]

5.5 XML in the .NET Framework

XML has rapidly gained popularity. Enterprise applications are using XML as the main data format for data exchanges.

ADO.NET breaks away from the COM-based recordset and employs XML as its transport data format. Because XML is platform independent, ADO.NET extends the reach to include anyone who is able to encode/decode XML. This is a big advantage over ADO because a COM-based recordset is not platform independent.

5.5.1 XML Parsers

Even though XML is text-based and readable by humans, you still should have some way of programmatically reading, inspecting, and changing XML. This is the job of XML parsers. There are two kinds of XML parsers: tree-based and stream-based. Depending on your needs, these two types of parsers should complement each other and serve you well.

Tree-based XML parsers read the XML file (or stream) in its entirety to construct a tree of XML nodes. Think of these XML nodes as your XML tag:

```
<car>
  <vin>VI00000383148374</vin>
  <make>Acura</make>
  <model>Integra</model>
  <year>1995</year>
</car>
```

When parsed into a tree, this information would have one root node: `car` ; under `car`, there are four nodes: `vin`, `make`, `model`, and `year`. As you might have suspected, if the XML stream is very large in nature, then a tree-based XML parser might not be a good idea. The tree would be too large and consume a lot of memory.

A Stream-based XML parser reads the XML stream as it goes. SAX (Simple API for XML) is a specification for this kind of parsing. The parser raises events as it reads the data, notifying the application of the tag or text the parser just read. It does not attempt to create the complete tree of all XML nodes as does the tree-based parser. Therefore, memory consumption is minimal. This kind of XML parser is ideal for going through large XML files to look for small pieces of data. The .NET framework introduces another stream-based XML parser: the `XmlReader`. While SAX pushes events at the application as it reads the data, the `XmlReader` allows the application to pull data from the stream.

Microsoft implements both types of parsers in its XML parser. Because XML is so powerful, Microsoft, among other industry leaders, incorporates XML usage in almost all the things they do. That includes, but is not limited to, the following areas:

- XML+HTTP in SOAP

- XML+SQL in SQL2000
- XML in BizTalk
- XML+DataSet in ADO.NET
- XML in web services and Web Services Discovery (DISCO) (see Chapter 6)

In this chapter, we will discuss XML+Dataset in ADO.NET, and XML in web services will be examined in the next chapter. Because XML is used everywhere in the .NET architecture, we also provide a high-level survey of the XML classes.

5.5.2 XML Classes

To understand the tree-based Microsoft XML parser, which supports the Document Object Model (DOM Level 2 Core standard), there are only a handful of objects you should know:

- XmlNode and its derivatives
- XmlNodeList, as collection XmlNode
- XmlNamedNodeMap, as a collection of XmlAttribute

We will walk through a simple XML example to see how XML nodes are mapped into these objects in the XML DOM.

5.5.2.1 XmlNode and its derivatives

XmlNode is a base class that represents a single node in the XML document. In the object model, almost everything derives from XmlNode (directly or indirectly). This includes: XmlAttribute, XmlDocument, XmlElement, and XmlText, among other XML node types.

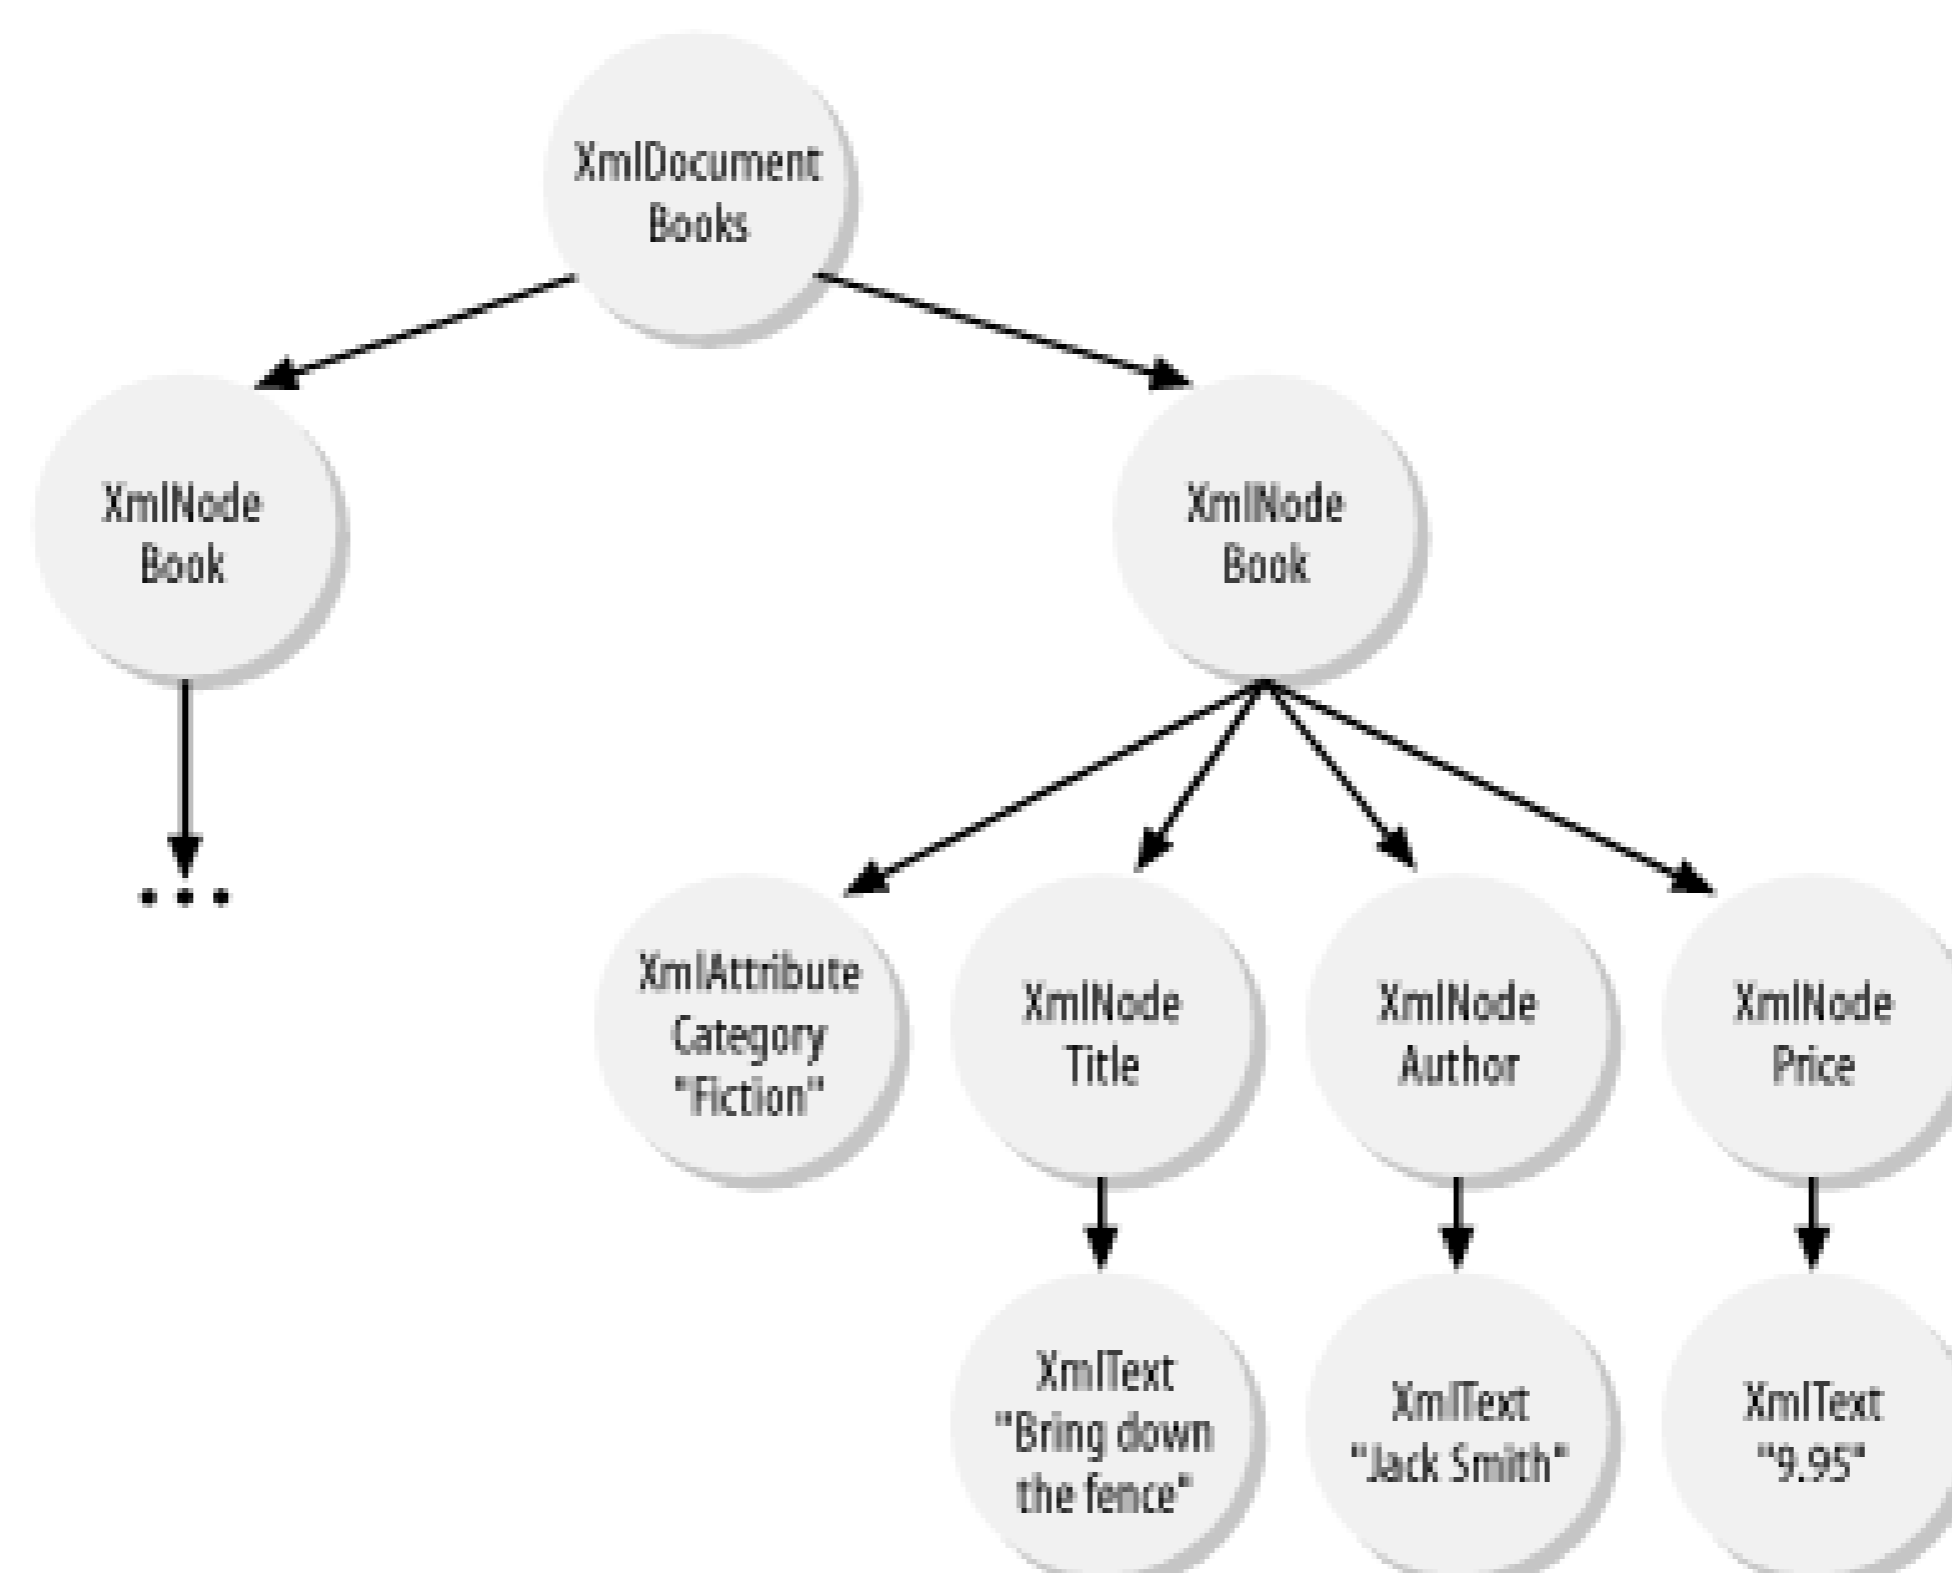
The following XML excerpt demonstrates mapping of XML tags to the node types in the DOM tree:

```
<books>
  <book category="How To">
    <title>How to drive in DC metropolitan</title>
    <author>Jack Daniel</author>
    <price>19.95</price>
  </book>
  <book category="Fiction">
    <title>Bring down the fence</title>
    <author>Jack Smith</author>
    <price>9.95</price>
  </book>
</books>
```

After parsing this XML stream, you end up with the tree depicted in Figure 5-6. It contains one root node, which is just a derivative of XmlNode. This root node is of type XmlDocument. Under this `books` root node, you have two children, also derivatives of XmlNode. This time, they are of type XmlElement. Under each `book` element node, there are four children. The first child is `category`. This `category` node

is of type `XmlAttribute`, a derivative of `XmlNode`. The next three children are of type `XmlElement`: `title`, `author`, and `price`. Each of these elements has one child of type `XmlText`.

Figure 5-6. Tree representation of an XML document



As a base class, `XmlNode` supports a number of methods that aid in the constructing of the XML document tree. These methods include `AppendChild()`, `PrependChild()`, `InsertBefore()`, `InsertAfter()`, and `Clone()`.

`XmlNode` also supports a group of properties that aid in navigation within the XML document tree. These properties include `FirstChild`, `NextSibling`, `PreviousSibling`, `LastChild`, `ChildNodes`, and `ParentNode`. You can use the `ChildNodes` property to navigate down from the root of the tree. For traversing backward, use the `ParentNode` property from any node on the tree.

5.5.2.2 XmlNodeList

Just as an `XmlNode` represents a single XML element, `XmlNodeList` represents a collection of zero or more `XmlNode`s. The `ChildNodes` property of the `XmlNode` is of type `XmlNodeList`. Looking at the root node `books`, we see that its `ChildNodes` property would be a collection of two `XmlNode`s. `XmlNodeList` supports enumeration, so we can iterate over the collection to get to each of the `XmlNode` objects. We can also index into the collection through a zero-based index.

Each of the book `XmlElement` objects would have a `ChildNodes` collection that iterates over `title`, `author`, and `price` `XmlElement`s.

5.5.2.3 XmlNamedNodeMap

Similar to `XmlNodeList`, `XmlNamedNodeMap` is also a collection object. `XmlNamedNodeMap` is a collection of `XmlAttribute` objects that enable both enumeration and indexing of attributes by name. Each `XmlNode` has a property named `Attributes`. In the case of the book elements, these collections

contain only one attribute, which is `category` .

5.5.2.4 XmlDocument

In addition to all methods and properties supported by `XmlNode`, this derivative of `XmlNode` adds or restricts methods and properties. Here, we inspect only `XmlDocument` as an example of a derivative of `XmlNode`.

`XmlDocument` extends `XmlNode` and adds a number of helper functions. These helper functions are used to create other types of `XmlNodes`, such as `XmlAttribute`, `XmlComment`, `XmlElement`, and `XmlText`. In addition to allowing for the creation of other XML node types, `XmlDocument` also provides the mechanism to load and save XML contents.

The following code demonstrates how an `XmlDocument` is programmatically generated with DOM:

```
using System;
using System.Xml;

public class XmlDemo {

    public static void Main( ) {

        // Code that demonstrates how to create XmlDocument programmatically
        XmlDocument xmlDom = new XmlDocument( );
        xmlDom.AppendChild(xmlDom.CreateElement("", "books", ""));
        XmlElement xmlRoot = xmlDom.DocumentElement;
        XmlElement xmlBook;
        XmlElement xmlTitle, xmlAuthor, xmlPrice;
        XmlText xmlText;

        xmlBook= xmlDom.CreateElement("", "book", "");
        xmlBook.SetAttribute("category", "", "How To");

        xmlTitle = xmlDom.CreateElement("", "title", "");
        xmlText = xmlDom.CreateTextNode("How to drive in DC metropolitan");
        xmlTitle.AppendChild(xmlText);
        xmlBook.AppendChild(xmlTitle);

        xmlAuthor = xmlDom.CreateElement("", "author", "");
        xmlText = xmlDom.CreateTextNode("Jack Daniel");
        xmlAuthor.AppendChild(xmlText);
        xmlBook.AppendChild(xmlAuthor);

        xmlPrice = xmlDom.CreateElement("", "price", "");
        xmlText = xmlDom.CreateTextNode("19.95");
        xmlPrice.AppendChild(xmlText);
        xmlBook.AppendChild(xmlPrice);

        xmlRoot.AppendChild(xmlBook);

        xmlBook= xmlDom.CreateElement("", "book", "");
        xmlBook.SetAttribute("category", "", "Fiction");
```

```

xmlTitle = xmlDoc.CreateElement("", "title", "");
xmlText = xmlDoc.CreateTextNode("Bring down the fence");
xmlTitle.AppendChild(xmlText);
xmlBook.AppendChild(xmlTitle);

xmlAuthor = xmlDoc.CreateElement("", "author", "");
xmlText = xmlDoc.CreateTextNode("Jack Smith");
xmlAuthor.AppendChild(xmlText);
xmlBook.AppendChild(xmlAuthor);

xmlPrice = xmlDoc.CreateElement("", "price", "");
xmlText = xmlDoc.CreateTextNode("9.95");
xmlPrice.AppendChild(xmlText);
xmlBook.AppendChild(xmlPrice);

xmlRoot.AppendChild(xmlBook);

Console.WriteLine(xmlDoc.InnerXml);
}
}

```

The XmlDocument also supports LoadXml and Load methods, which build the whole XML tree from the input parameter. LoadXml takes a string in XML format, whereas Load can take a stream, a filename or a URL, a TextReader, or an XmlReader. The following example continues where the previous one left off. The XML tree is saved to a file named *books.xml*. Then this file is loaded back into a different XML tree. This new tree outputs the same XML stream as the previous one:

```

. . .
xmlDoc.Save("books.xml");
XmlDocument xmlDoc2 = new XmlDocument( );
xmlDoc2.Load("books.xml");
Console.WriteLine(xmlDoc2.InnerXml);

```

5.5.2.5 XmlReader

The XmlReader object is a fast, noncached, forward-only way of accessing streamed XML data. There are two derivatives of XmlReader: XmlTextReader and XmlNodeReader. Both of these readers read XML one tag at a time. The only difference between the two is the input to each reader. As the name implies XmlTextReader reads a stream of pure XML text. XmlNodeReader reads a stream of nodes from an XmlDocument. The stream can start at the beginning of the XML file for the whole XmlDocument or only at a specific node of the XmlDocument for partial reading.

Consider the following XML excerpt for order processing. If this file is large, it is not reasonable to load it into an XmlDocument and perform parsing on it. Instead, we should read only nodes or attributes we are interested in and ignore the rest. We can use XmlReader derived classes to do so:

```

<Orders>
<Order id="ABC001" . . . >
<Item code="101" qty="3" price="299.00" . . . >17in Monitor</Item>

```

```

<Item code="102" qty="1" price="15.99" . . . >Keyboard</Item>
<Item code="103" qty="2" price="395.95" . . . >CPU</Item>
</Order>
<Order id="ABC002" . . . >
<Item code="101b" qty="1" price="499.00" . . . >21in Monitor</Item>
<Item code="102" qty="1" price="15.99" . . . >Keyboard</Item>
</Order>
< . . . >
</Orders>

```

The following block of code traverses and processes each order from the `largeOrders.xml` input file:

```

using System;
using System.IO;
using System.Xml;

class TestXMLReader
{

static void Main(string[] args)
{
    TestXMLReader tstObj = new TestXMLReader( );
    StreamReader myStream = new StreamReader("Orders.xml");
    XmlTextReader xmlTxtRdr = new XmlTextReader(myStream);
    while(xmlTxtRdr.Read( ))
    {
        if(xmlTxtRdr.NodeType == XmlNodeType.Element
            && xmlTxtRdr.Name == "Order")
        {
            tstObj.ProcessOrder(xmlTxtRdr);
        }
    }
}

public void ProcessOrder(XmlTextReader reader)
{
    Console.WriteLine("Start processing order: " +
        reader.GetAttribute("id"));
    while(!(reader.NodeType == XmlNodeType.EndElement
        && reader.Name == "Order")
        && reader.Read( ))
    {
        // Process Content of Order
        if(reader.NodeType == XmlNodeType.Element
            && reader.Name == "Item")
        {
            Console.WriteLine("itemcode:" + reader.GetAttribute("code") +
                ". Qty: " + reader.GetAttribute("qty"));
        }
    }
}
}

```

}

Let's take a closer look at what is going on. Once we have established the `XmlTextReader` object with the stream of data from the string, all we have to do is loop through and perform a `Read()` operation until there is nothing else to read. While we are reading, we start to process the order only when we come across a node of type `XmlElement` and a node named `Order`. Inside the `ProcessOrder` function, we read and process all items inside an order until we encounter the end tag of `Order`. In this case, we return from the function and go back to looking for the next `Order` tag to process the next order.

`XmlNodeReader` is similar to `XmlTextReader` because they both allow processing of XML sequentially. However, `XmlNodeReader` reads XML nodes from a complete or fragment of an XML tree. This means `XmlNodeReader` is not helpful when processing large XML files.

5.5.2.6 XmlWriter

The `XmlWriter` object is a fast, noncached way of writing streamed XML data. It also supports namespaces. The only derivative of `XmlWriter` is `XmlTextWriter`.

`XmlWriter` supports namespaces by providing a number of overloaded functions that take a namespace to associate with the element. If this namespace is already defined and there is an existing prefix, `XmlWriter` automatically writes the element name with the defined prefix. Almost all element-writing methods are overloaded to support namespaces.

The following code shows how to use an `XmlTextWriter` object to write a valid XML file:

```
XmlTextWriter writer =
    new XmlTextWriter("test.xml", new System.Text.AsciiEncoding( ));
writer.Formatting = Formatting.Indented;
writer.Indentation = 4;
writer.WriteStartDocument( );
writer.WriteComment("Comment");
writer.WriteStartElement("ElementName", "myns");
writer.WriteStartAttribute("prefix", "attrName", "myns");
writer.WriteEndAttribute( );
writer.WriteElementString("ElementName", "myns", "value");
writer.WriteEndElement( );
writer.WriteEndDocument( );
writer.Flush( );
writer.Close( );
```

This produces the following XML document in *test.xml*:

```
<?xml version="1.0" encoding="us-ascii"?>
<!--Comment-->
<ElementName prefix:attrName="" xmlns:prefix="myns" xmlns="myns">
  <prefix:ElementName>value</prefix:ElementName>
</ElementName>
```

5.5.2.7 XsltTransform

`XsltTransform` converts XML from one format to another. It is typically used in data-conversion programs

or to convert XML to HTML for the purpose of presenting XML data in a browser. The following code demonstrates how such a conversion takes place:

```
using System;
using System.Xml;           // XmlTextWriter
using System.Xml.Xsl;      // XslTransform
using System.Xml.XPath;    // XPathDocument
using System.IO;           // StreamReader

public class XSLDemo {
    public static void Main( ) {
        XslTransform xslt = new XslTransform( );
        xslt.Load("XSLTemplate.xsl");
        XPathDocument xDoc = new XPathDocument("Books.xml");
        XmlTextWriter writer = new XmlTextWriter("Books.html", null);
        xslt.Transform(xDoc, null, writer, new XmlUrlResolver( ));
        writer.Close( );
        StreamReader stream = new StreamReader("Books.html");
        Console.Write(stream.ReadToEnd( ));
    }
}
```

The code basically transforms the XML in the *Books.xml* file, which we've seen earlier, into HTML to be displayed in a browser. Even though you can replace the `XPathDocument` with `XmlDocument` in the previous code, `XPathDocument` is the preferred class in this case because it is optimized for XSLT processing.^[9]

^[9] `XPathDocument` loads data faster than `XmlDocument` because it does not maintain node identity and it does not perform rule checking. One catch to this advantage is that the content is read-only.

Figure 5-7 and Figure 5-8 show the source XML and the output HTML when viewed in a browser.

Figure 5-7. Books.xml shown in IE

Figure 5-8. Books.html shown in IE



The template XSL file that was used to transform the XML is:

```
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
<xsl:template match = "/" >

<html>
<head><title>A list of books</title></head>
<style>
.hdr { background-color=#ffeedd; font-weight:bold; }
</style>
<body>
<B>List of books</B>
<table style="border-collapse:collapse" border="1">
<tr>
  <td class="hdr">Title</td>
  <td class="hdr">Author</td>
  <td class="hdr">Price</td>
</tr>
<xsl:for-each select="//books/book">
<tr>
  <td><xsl:value-of select="title"/></td>
  <td><xsl:value-of select="author"/></td>
  <td><xsl:value-of select="price"/></td>
</tr>
</xsl:for-each>
</table>
</body>
</html>

</xsl:template>
</xsl:stylesheet>
```

5.5.2.8 XmlDocument

One of the most important points in ADO.NET is the tight integration of DataSet with XML. DataSet can easily be streamed into XML and vice versa, making it easy to exchange data with any other components in the enterprise system. The schema of the DataSet can be loaded and saved as XML Schema Definition (XSD), as described earlier.

XmlDocument can be associated with DataSet. The following code excerpt illustrates how such an

association takes place:

```
using System;
using System.Data;
using System.Data.OleDb;
using System.Xml;

class TestXMLDataDocument
{

static void Main(string[] args)
{
    TestXMLDataDocument tstObj = new TestXMLDataDocument( );

    // Construct the XmlDataDocument with the DataSet.
    XmlDataDocument doc = tstObj.GenerateXmlDataDocument( );

    XmlNodeReader myXMLReader = new XmlNodeReader(doc);
    while (myXMLReader.Read( ))
    {
        if(myXMLReader.NodeType == XmlNodeType.Element
            && myXMLReader.Name == "Orders")
        {
            tstObj.ProcessOrder(myXMLReader);
        }
    }
}

public void ProcessOrder(XmlNodeReader reader)
{
    Console.WriteLine("Start processing order: ");
    while(!(reader.NodeType == XmlNodeType.EndElement
        && reader.Name == "Orders")
        && reader.Read( ))
    {
        if(reader.NodeType == XmlNodeType.Element
            && reader.Name == "OrderID")
        {
            reader.Read( );
            Console.WriteLine(reader.Value);
        }
        if(reader.NodeType == XmlNodeType.Element
            && reader.Name == "OrderDetails")
        {
            ProcessLine(reader);
        }
    }
}

public void ProcessLine(XmlNodeReader reader)
{
    while(!(reader.NodeType == XmlNodeType.EndElement
```

```

        && reader.Name == "OrderDetails")
        && reader.Read( ))
    {
        if(reader.NodeType == XmlNodeType.Element && reader.Name == "ProductID")
        {
            reader.Read( );
            Console.WriteLine(". ItemCode: " + reader.Value);
        }
        if(reader.NodeType == XmlNodeType.Element && reader.Name == "Quantity")
        {
            reader.Read( );
            Console.WriteLine(". Quantity: " + reader.Value);
        }
    }
}
public XmlDataDocument GenerateXmlDataDocument( )
{
    /* Create the DataSet object. */
    DataSet ds = new DataSet("DBDataSet");
    String sConn =
        "provider=SQLOLEDB;server=(local);database=NorthWind;Integrated Security=SSPI";

    /* Create the DataSet adapters. */
    OleDbDataAdapter dsAdapter1 =
        new OleDbDataAdapter("select * from Orders", sConn);

    OleDbDataAdapter dsAdapter2 =
        new OleDbDataAdapter("select * from [Order Details]", sConn);

    /* Fill the data set with three tables. */
    dsAdapter1.Fill(ds, "Orders");
    dsAdapter2.Fill(ds, "OrderDetails");

    DataColumn[] keys = new DataColumn[1];
    keys[0] = ds.Tables["Orders"].Columns["OrderID"];
    ds.Tables["Orders"].PrimaryKey = keys;

    // Add the two relations between the three tables. */
    ds.Relations.Add("Orders_OrderDetails",
        ds.Tables["Orders"].Columns["OrderID"],
        ds.Tables["OrderDetails"].Columns["OrderID"]);

    ds.Relations["Orders_OrderDetails"].Nested = true;
    //ds.WriteXml("NorthWindOrders.xml");

    return new XmlDataDocument(ds);
}
}

```

The previous section describing DataSet has already shown you that once we have a DataSet, we can persist the data inside the DataSet into an XML string or file. This time, we demonstrated how to convert the DataSet into an XmlDocument that we can manipulate in memory.

[Team LiB]

[\[Team LiB \]](#)

5.6 Summary

This chapter describes the core of ADO.NET. Having focused on the disconnected dataset, ADO.NET enables us not only to build high-performance, scalable solutions for e-commerce, but also allows the applications to reach other platforms through the use of XML. This chapter serves as a high-level survey into the classes that make up ADO.NET and serves to familiarize you with the System.Xml library.^[10] In the next chapter, we delve into building software as services. We will make use of ADO.NET as the data-access and exchange mechanism in our software services.

^[10] For additional information on ADO.NET, see O'Reilly's *ADO.NET in a Nutshell*, by Bill Hamilton and Matthew MacDonald, and *ADO.NET Cookbook*, by Bill Hamilton.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Chapter 6. Web Services

Web services allow access to software components through standard web protocols such as HTTP and SMTP. Using the Internet and XML, we can now create software components that communicate with others, regardless of language, platform, or culture. Until now, software developers have progressed toward this goal by adopting proprietary componentized software methodologies, such as DCOM; however, because each vendor provides its own interface protocol, integration of different vendors' components is a nightmare. By substituting the Internet for proprietary transport formats and adopting standard protocols such as SOAP, web services help software developers create building blocks of software, which can be reused and integrated regardless of their location.

In this chapter, we describe the .NET web services architecture and provide examples of a web service provider and several web service consumers.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

6.1 Web Services in Practice

You may have heard the phrase "software as a service" and wondered about its meaning. The term *service*, in day-to-day usage, refers to what you get from a service provider. For example, you bring your dirty clothing to a cleaner to use its cleaning *service*. Software, on the other hand, is commonly thought of as an application, either an off-the-shelf product, or a custom application developed by a software firm. You typically buy the software (or in our case, build the software). It usually resides on some sort of media such as floppy diskette or CD and is sold in a shrink-wrapped package through retail outlets, or, in the case of a web application, the software application is not distributed, but is accessed through a browser.

How can software be viewed as a service? The example we are about to describe might seem far-fetched; however, it is possible with current technology. Imagine the following. As you grow more comfortable with the Internet, you might choose to replace your computer at home with something like an Internet Device, a large-screen PDA designed for use with the Internet. Let's call the device an iDev. Let's suppose that with this device, you can be on the Internet immediately through your cell phone, WiFi, or some other means. When you want to do word processing, your iDev is configured to print to a Microsoft Word service somewhere in Redmond, so you can type away without the need to install word-processing software. When you are done, the document can be saved at an iStore server where you can later retrieve it. Notice that for you to do this, the iStore server must host a software service that allows you to store documents. Microsoft might charge you a service fee based on the amount of time your word processor is running and the features you use (such as the grammar and spell checkers). The iStore service charges might vary based on the size of your document and how long it is stored. Of course, none of these charges would come in the mail but rather through an escrow service where the money would be withdrawn from your bank account or credit card.

As long as your document is in a standard format, such as XML, you're free to switch word processors at any time. Of course, the document that you store at the iStore server is already in a standard data format. Since iStore utilizes the iMaxSecure software service from a company called iNNSA (Internet Not National Security Agency), the security of your files is assured. And because you use the document storage service at iStore, you also benefit from having your document authenticated and decrypted upon viewing, as well as encrypted at storing time.

While this particular vision of software as a service has yet to be realized, a variety of for-fee and free services have begun to appear. In early 2001, Microsoft announced plans for an integrated collection of consumer-oriented services (known first by their codename, "Hailstorm," and later as ".NET My Services") but was forced to abandon the initiative for a variety of reasons, some technical and others legal, political, or market-related. Today, Microsoft offers a variety of user-centric services for identification and authentication, email, instant messaging, automated alerts, calendar, address book, and personal information storage. These are available through its MSN online services, and through Passport (<http://www.passport.net>), Alerts (<http://alerts.microsoft.com>), MSN Wallet (<http://wallet.msn.com>), and Hotmail.

Of greater interest to developers, however, is the availability of these services for use as building blocks in third-party web applications. Hosted by Microsoft and known as Microsoft .NET Services, Passport, Alerts, and MSN Wallet can each be licensed and incorporated into any application that

adheres to XML web services standards, regardless of platform. A more recently announced web service is Microsoft MapPoint .NET, which is a set of services that allows you to incorporate maps, driving directions, distance calculations, proximity searches, and other location intelligence into your applications.

In addition to Microsoft, other companies are beginning to offer information and functionality as services over the Web. A recent poster child is Salesforce.com, which offers customer relations management (CRM) software over the Web, either as a standalone product or as a set of services that can be incorporated into third-party applications. The Liberty Alliance is at work defining an authentication service that can be offered as an alternative to Microsoft Passport. And both Google and Amazon now make portions of their business information available through public web service interfaces.

The potential for consumer-oriented and business-to-business web services like Microsoft .NET Services is great, although there are serious and well-founded concerns about security and privacy. In the mean time, web services can be great in interoperability areas where there are needs to expose legacy functionalities or to enable interaction between multiple heterogeneous systems. In one form or another, though, web services are here to stay, so let's dive in and see what's underneath.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

6.2 Web Services Framework

Web services combine the best of both distributed componentization and the World Wide Web, extending distributed computing to broader ranges of client applications. The best thing is that this is done by seamlessly marrying and enhancing existing technologies.

6.2.1 Web Services Architecture

Web services are distributed software components accessible through standard web protocols. The first part of the definition is similar to that for COM/DCOM components. However, it is the second part that distinguishes web services from the crowd. Web services enable software to interoperate with a much broader range of clients. While COM-aware clients can understand only COM components, web services can be consumed by any application that understands how to parse an XML-formatted stream transmitted through HTTP channels. XML is the key technology used in web services and is used in the following areas of the Microsoft .NET web services framework:

Web service wire formats

The technology enabling universal understanding of how to perform data exchanges between the service provider and consumer; the format of data for the request and response.

Web service description in Web Services Description Language (WSDL)

The language describing how the service can be used. Think of this as the instructions on the washing machine at the laundromat telling you where to put quarters, what buttons to push, etc.

Web service discovery

The process of advertising or publishing a piece of software as a service and allowing for the discovery of this service.

[Figure 6-1](#) depicts the architecture of web applications using Windows DNA, while [Figure 6-2](#) shows .NET-enabled web applications architecture. As you can see, communication between components of a web application does not have to be within an intranet. Furthermore, intercomponent communication can also use HTTP/XML.

Figure 6-1. Windows Distributed interNet Architecture (DNA)


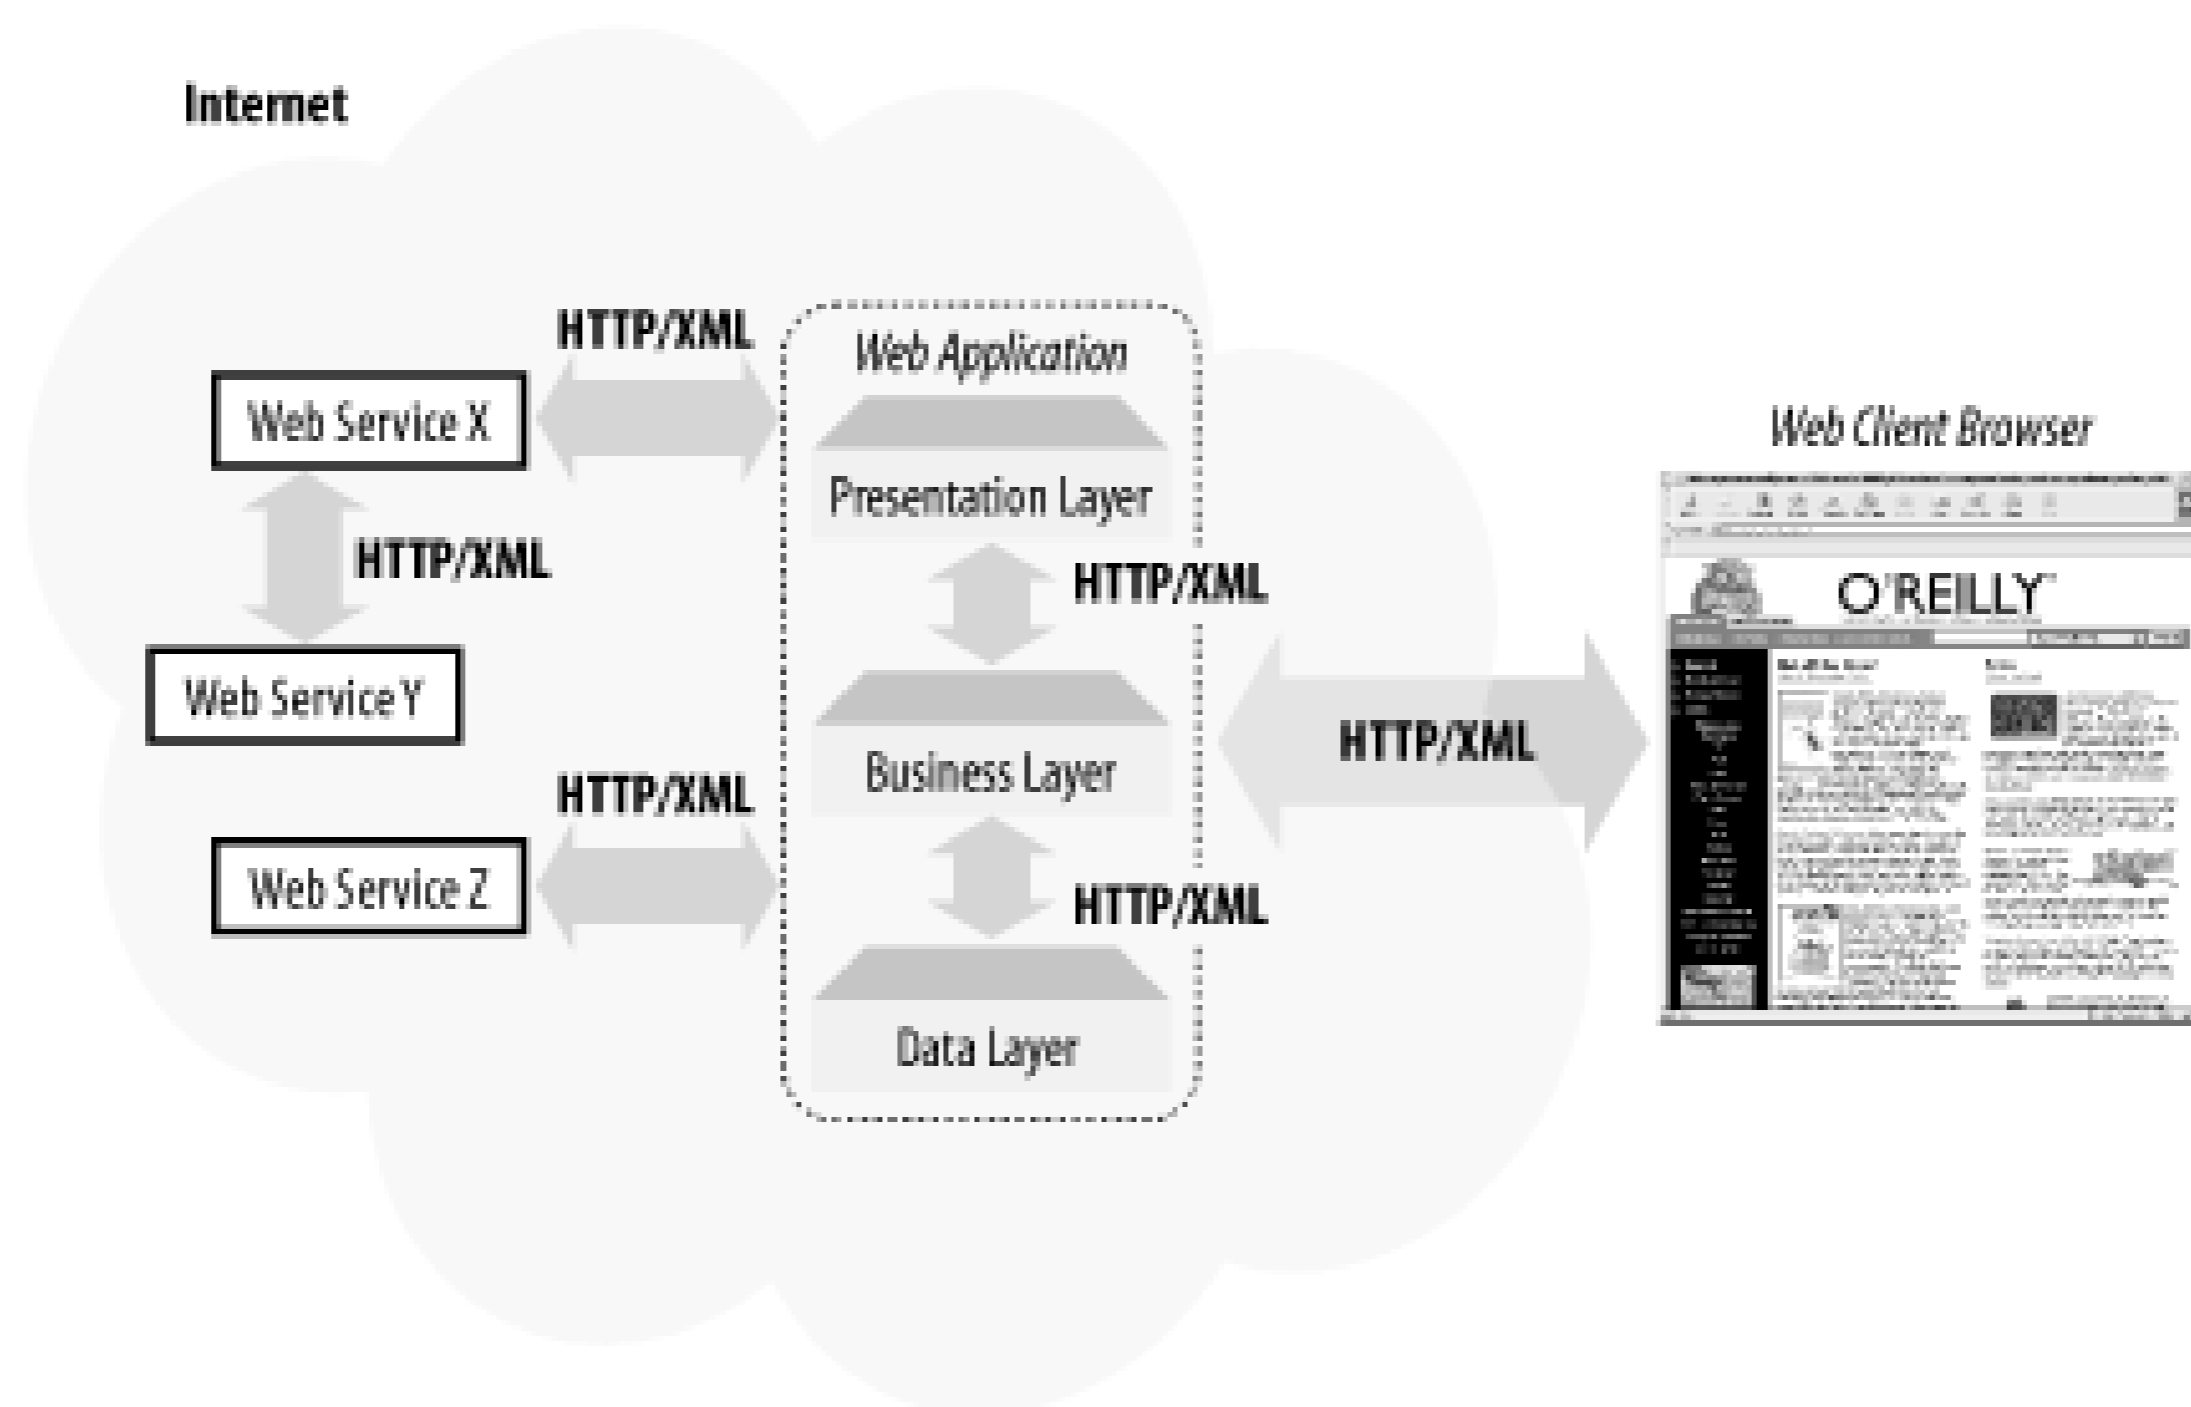


Figure 6-2. .NET-enabled web application framework



6.2.2 Web Services Wire Formats

You may have heard the phrase "DCOM is COM over the wire." Web services are similar to DCOM except that the wire is no longer a proprietary communication protocol. With web services, the wire formats rely on more open Internet protocols such as HTTP or SMTP.

A web service is more or less a stateless component running on the web server, exposed to the world through standard Internet protocols. Microsoft .NET web services currently supports three protocols: HTTP GET, HTTP POST, and SOAP over HTTP (Simple Object Access Protocol), explained in the next sections. Because these protocols are standard protocols for the Web, it is very easy for the client applications to use the services provided by the server.

6.2.2.1 HTTP GET and HTTP POST

As their names imply, both HTTP GET and HTTP POST use HTTP as their underlying protocol. The GET and POST methods of the HTTP protocol have been widely used in ASP (Active Server Pages), CGI, and other server-side architectures for many years now. Both of these methods encode request

parameters as name/value pairs in the HTTP request. The GET method creates a query string and appends it to the script's URL on the server that handles the request. For the POST method, the name/value pairs are passed in the body of the HTTP request message.

6.2.2.2 SOAP

Similar to HTTP GET and HTTP POST, SOAP serves as a mechanism for passing messages between the clients and servers. In this context, the clients are web services consumers, and the servers are the web services. The clients simply send an XML-formatted request message to the server to get the service over an HTTP channel. The server responds by sending back yet another XML-formatted message. The SOAP specification describes the format of these XML requests and responses. It is simple, yet it is extensible, because it is based on XML.

SOAP is different than HTTP GET and HTTP POST because it uses XML to format its payload. The messages being sent back and forth have a better structure and can convey more complex information compared to simple name/value pairs in HTTP GET/POST protocols. Another difference is that SOAP can be used on top of other transport protocols, such as SMTP in addition to HTTP.^[1]

[1] SOAP is recognized as the cross-platform standard protocol to use for web services both inside and outside the Microsoft circle.

6.2.3 Web Services Description (WSDL)

For web service clients to understand how to interact with a web service, there must be a description of the method calls, or the interface that the web service supports. This web service description document is found in an XML schema called *Web Services Description Language (WSDL)*. Remember that type libraries and IDL scripts are used to describe a COM component. Both IDL and WSDL files describe an interface's method calls and the list of in and out parameters for the particular call. The only major difference between the two description languages is that all descriptions in the WSDL file are done in XML.

In theory, any WSDL-capable SOAP client can use the WSDL file to get a description of your web service. It can then use the information contained in that file to understand the interface and invoke your web service's methods.

6.2.3.1 WSDL structure

The root of any web service description file is the `<definitions>` element. Within this element, the following elements provide both the abstract and concrete description of the service:

Types

A container for data type definitions.

Message

An abstract, typed definition of the data being exchanged between the web service providers and consumers. Each web method has two messages: input and output. The input describes the parameters for the web method; the output describes the return data from the web

method. Each message contains zero or more `<part>` parameters. Each parameter associates with a concrete type defined in the `<types>` container element.

Port type

An abstract set of operations supported by one or more endpoints.

Operation

An abstract description of an action supported by the service. Each operation specifies the input and output messages defined as `<message>` elements.

Binding

A concrete protocol and data-format specification for a particular port type. Similar to port type, the binding contains operations, as well as the input and output for each operation. The main difference is that with binding, we are now talking about actual transport type and how the input and output are formatted.

Service

A collection of network endpoints-ports. Each of the web service wire formats defined earlier constitutes a port of the service (HTTP GET, HTTP POST, and SOAP ports).

Port

A single endpoint defined by associating a binding and a network address. In other words, it describes the protocol and data-format specification to be used as well as the network address of where the web service clients can bind to for the service.

The following shows a typical WSDL file structure:

```
<definitions name="" targetNamespace="" xmlns: . . . >

  <types> . . . </types>

  <message name=""> . . . </message>
  . . .

  <portType name="">
    <operation name="">
      <input message="" />
      <output message="" />
    </operation>
    . . .
  </portType>
  . . .

  <binding name="">
    <protocol:binding . . . >
    <operation name="">
      <protocol:operation . . . >
      <input> . . . </input>
      <output> . . . </output>
    </operation>
    . . .
  </binding>
  . . .
```

```

<service name="">
  <port name="" binding="">
    <protocol:address location="" />
  </port>
  . . .
</service>
</definitions>

```

The `<types>` element contains physical type descriptions defined in XML Schema (XSD). These types are being referred to from the `<message>` elements.

For each of the web methods in the web service, there are two messages defined for a particular port: input and output. This means if a web service supports all three protocols: SOAP, HTTP GET, and HTTP POST, there will be six `<message>` elements defined, one pair for each port. The naming convention used by the Microsoft .NET autogenerated WSDL is:

MethodName + Protocol + {In, Out}

For example, a web method called `GetBooks()` has the following messages:

```

<message name="GetBooksSoapIn"> . . . </message>
<message name="GetBooksSoapOut"> . . . </message>
<message name="GetBooksHttpGetIn"> . . . </message>
<message name="GetBooksHttpGetOut"> . . . </message>
<message name="GetBooksHttpPostIn"> . . . </message>
<message name="GetBooksHttpPostOut"> . . . </message>

```

For each protocol that the web service supports, there is one `<portType>` element defined. Within each `<portType>` element, all operations are specified as `<operation>` elements. The naming convention for the port type is:

WebServiceName + Protocol

To continue our example, here are the port types associated with the web service that we build later in this chapter, PubsWS:

```

<portType name="PubsWSSoap">
  <operation name="GetBooks">
    <input message="GetBooksSoapIn" />
    <output message="GetBooksSoapOut" />
  </operation>
</portType>

<portType name="PubsWSHttpGet">
  <operation name="GetBooks">
    <input message="GetBooksHttpGetIn" />
    <output message="GetBooksHttpGetOut" />
  </operation>
</portType>

<portType name="PubsWSHttpPost">
  <operation name="GetBooks">

```

```

    <input message="GetBooksHttpPostIn" />
    <output message="GetBooksHttpPostOut" />
  </operation>
</portType>

```

We have removed namespaces from the example to make it easier to read.

While the port types are abstract operations for each port, the bindings provide concrete information on what protocol is being used, how the data is being transported, and where the service is located. Again, there is a `<binding>` element for each protocol supported by the web service:

```

<binding name="PubsWSSoap" type="s0:PubsWSSoap">
  <soap:binding transport="http://schemas.xmlsoap.org/soap/http"
    style="document" />
  <operation name="GetBooks">
    <soap:operation soapAction="http://tempuri.org/GetBooks"
      style="document" />
    <input>
      <soap:body use="literal" />
    </input>
    <output>
      <soap:body use="literal" />
    </output>
  </operation>
</binding>

```

```

<binding name="PubsWSHttpGet" type="s0:PubsWSHttpGet">
  <http:binding verb="GET" />
  <operation name="GetBooks">
    <http:operation location="/GetBooks" />
    <input>
      <http:urlEncoded />
    </input>
    <output>
      <mime:mimeType part="Body" />
    </output>
  </operation>
</binding>

```

```

<binding name="PubsWSHttpPost" type="s0:PubsWSHttpPost">
  <http:binding verb="POST" />
  <operation name="GetBooks">
    <http:operation location="/GetBooks" />
    <input>
      <mime:content type="application/x-www-form-urlencoded" />
    </input>
    <output>
      <mime:mimeType part="Body" />
    </output>
  </operation>
</binding>

```

For SOAP protocol, the binding is `<soap:binding>`, and the transport is SOAP messages on top of HTTP protocol. The `<soap:operation>` element defines the HTTP header `soapAction`, which points to the web method. Both input and output of the SOAP call are SOAP messages.

For the HTTP GET and HTTP POST protocols, the binding is `<http:binding>` with the verb being GET and POST, respectively. Because the GET and POST verbs are part of the HTTP protocol, there is no need for the extended HTTP header (like `soapAction` for SOAP protocol). The only thing we need is the URL that points to the web method; in this case, the `<soap:operation>` element contains the attribute `location`, which is set to `/GetBooks`.

The only real difference between the HTTP GET and POST protocols is the way the parameters are passed to the web server. HTTP GET sends the parameters in the query string, while HTTP POST sends the parameters in the form data. This difference is reflected in the `<input>` elements of the operation `GetBooks` for the two HTTP protocols. For the HTTP GET protocol, the input is specified as `<http:urlEncoded />`, whereas for the HTTP POST protocol, the input is `<mime:contentType="application/x-www-form-urlencoded" />`.

Looking back at the template of the WSDL document, we see that the only thing left to discuss is the `<service>` element, which defines the ports supported by this web service. For each of the supported protocol, there is one `<port>` element:

```
<service name="PubsWS">

  <port name="PubsWSSoap" binding=s0:PubsWSSoap">
    <soap:address
      location="http:// . . . /PubsWs.asmx" />
  </port>

  <port name="PubsWSHttpGet" binding="s0:PubsWSHttpGet">
    <http:address
      location="http:// . . . /PubsWs.asmx" />
  </port>

  <port name="PubsWSHttpPost" binding="s0:PubsWSHttpPost">
    <http:address
      location="http:// . . . /PubsWs.asmx" />
  </port>

</service>
```

Even though the three different ports look similar, their `binding` attributes associate the address of the service with a `binding` element defined earlier. Web service clients now have enough information on where to access the service, through which port to access the web service method, and how the communication messages are defined.

Although it is possible to read the WSDL and manually construct the HTTP^[2] conversation with the server to get to a particular web service, there are tools that autogenerate client-side proxy source code to do the same thing. We show such a tool in "web services Consumers" later in this chapter.

[2] Current Microsoft .NET SOAP implementation runs on top of HTTP.

6.2.4 Web Services Discovery

Even though advertising of a web service is important, it is optional. Web services can be private as well as public. Depending on the business model, some business-to-business (B2B) services would not normally be advertised publicly. Instead, the web service owners would provide specific instructions on accessing and using their service only to the business partner.

To advertise web services publicly, authors post discovery files on the Internet. Potential web service clients can browse to these files for information about how to use the web services—the WSDL. Think of it as the yellow pages for the web service. All it does is point you to where the actual web services reside and to the description of those web services.

The process of looking up a service and checking out the service description is called *web service discovery*. Currently in .NET, there are two ways of advertising the service: static and dynamic. In both of these, XML conveys the locations of web services.

6.2.4.1 Static discovery

Static discovery is easier to understand because it is explicit in nature. If you want to advertise your web service, you must explicitly create the .disco discovery file and point it to the WSDL.^[3]

^[3] If you use Visual Studio .NET to create your web service, the discovery file is created automatically.

All *.disco* files contain a root element `discovery`, as shown in the following code sample. Note that `discovery` is in the namespace `http://schemas.xmlsoap.org/disco/`, which is referred to as `disco` in this sample.

```
<?xml version="1.0" ?>
<disco:discovery xmlns:disco="http://schemas.xmlsoap.org/disco/">
</disco:discovery>
```

Inside the `discovery` element, there can be one or more of `contractRef` or `discoveryRef` elements. Both of these elements are described in the namespace `http://schemas.xmlsoap.org/disco/scl/`. The `contractRef` tag is used to reference an actual web service URL that would return the WSDL or the description of the actual web service contract. The `discoveryRef` tag, on the other hand, references another discovery document.

This XML document contains a link to one web service and a link to another discovery document:

```
<?xml version="1.0" ?>
<disco:discovery
  xmlns:disco="http://schemas.xmlsoap.org/disco/"
  xmlns:scl="http://schemas.xmlsoap.org/disco/scl/">
  <scl:contractRef ref="http://yourWebServer/yourWebService.asmx?WSDL"/>
  <scl:discoveryRef ref="http://yourBrotherSite/hisWebServiceDirectory.disco"/>
```



```
</disco:discovery>
```

This sample *disco* file specifies two different namespaces: *disco*, which is a nickname for the namespace `http://schemas.xmlsoap.org/disco/`; and *scl*, short for `http://schemas.xmlsoap.org/disco/scl/`. The `contractRef` element specifies the URL where *yourWebService* WSDL can be obtained. Right below that is the `discoveryRef` element, which links to the discovery file on *yourBrotherSite* web site. This linkage allows for structuring networks of related discovery documents.

6.2.4.2 Dynamic discovery

As opposed to explicitly specifying the URL for all web services your site supports, you can enable *dynamic discovery*, which enables all web services underneath a specific URL on your web site to be listed automatically. For your web site, you might want to group related web services under many different directories and then provide a single dynamic discovery file in each of the directories. The root tag of the dynamic discovery file is *dynamicDiscovery* instead of *discovery*.

```
<?xml version="1.0" encoding="utf-8"?>
<dynamicDiscovery xmlns="urn://schemas-dynamic:disco.2000-03-17" />
```

You can optionally specify exclude paths so that the dynamic mechanism does not have to look for web services in all subdirectories underneath the location of the dynamic discovery file. Exclude paths are in the following form:

```
<exclude path="pathname" />
```

If you run IIS as your web server, you'd probably have something like the following for a dynamic discovery file: [\[4\]](#)

[4] VS.NET uses vsdisco as the extension for its dynamic discovery files.

```
<?xml version="1.0" encoding="utf-8"?>
<dynamicDiscovery xmlns="urn://schemas-dynamic:disco.2000-03-17">
  <exclude path="_vti_cnf" />
  <exclude path="_vti_pvt" />
  <exclude path="_vti_log" />
  <exclude path="_vti_script" />
  <exclude path="_vti_txt" />
  <exclude path="Web References" />
</dynamicDiscovery>
```

6.2.4.3 Discovery setting in practice

A combination of dynamic and static discovery makes a very flexible configuration. For example, you can provide static discovery documents at each of the directories that contain web services. At the root of the web server, provide a dynamic discovery document with links to all static discovery documents in all subdirectories. To exclude web services from public viewing, provide the `exclude` argument to XML nodes to exclude their directories from the dynamic discovery document.

6.2.4.4 UDDI

Universal Description, Discovery, and Integration (UDDI) Business Registry is like a yellow pages of web services. It allows businesses to publish their services and locate web services published by partner organizations so that they can conduct transactions quickly, easily, and dynamically with their trading partner.

Through UDDI APIs, businesses can find services over the web that match their criteria (e.g., cheapest fare), that offer the service they request (e.g., delivery on Sunday), and so on. Currently backed by software giants such as Microsoft, IBM, and Ariba, UDDI is important to web services because it enables access to businesses from a single place.^[5]

^[5] UDDI SDK can be downloaded from Microsoft as a .NET assembly.

6.2.5 The System.Web.Services Namespace

Now that we have run through the basic framework of Microsoft .NET web services, let us take a look inside what the .NET SDK provides us in the System.Web.Services namespace.

There are only a handful of classes in the System.Web.Services namespace and the most important ones for general use are:

WebService

The base class for all web services.

WebServiceAttribute

An attribute that can be associated with a web service-derived class.

WebMethodAttribute

An attribute that can be associated with public methods within a web service-derived class.

The two essential classes for creating web services are the WebService base class and WebMethodAttribute. We make use of these classes in the next section, where we implement a web service provider and several web service consumers. WebService is the base class from which all web services inherit. It provides properties inherent to legacy ASP programming such as Application, Server, Session, and a new property, Context, which now includes Request and Response.

The WebMethodAttribute class allows you to apply attributes to each public method of your web service. Using this class, you can assign specific values to the following attributes: description, session state enabling flag, message name, transaction mode, and caching. See the following section for an example of attribute setting in C# and VB.

The WebServiceAttribute class is used to provide more attributes about the web service itself. You can display a description of the web service, as well as the namespace to which this web service belongs.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

6.3 Web Services Provider

In this section, we describe how to develop a web service, first from the point of view of service providers and then of the consumers. Web services providers implement web services and advertise them so that the clients can discover and make use of the services. Because web services run on top of HTTP, there must be a web server application of some sort on the machine that hosts the web services. This web server application can be Microsoft Internet Information Services (IIS), Apache, or any other program that can understand and process the HTTP protocol. In our examples, we use Microsoft IIS, since that is the only web server currently supported by .NET.

6.3.1 Web Service Provider Example

We will be building a web service called *PubsWS* to let consumers get information from the sample *Pubs* database. All data access will be done through ADO.NET, so read [Chapter 5](#) before attempting the examples.

Creating a web service is a three-step process:

1. Create a new *asmx* file for the web service. This must contain the `<% webservice . . . %>` directive, as well as the class that provides the web service implementation. To the web service clients, this *asmx* file is the entry point to your web service. You need to put this in a virtual directory that has the `executescripts` permission turned on.
2. Inherit from the `WebService` class of the `System.Web.Services` namespace. This allows the derived class to access all the normal ASP.NET objects exposed in the `WebService` base class such as `Application`, `Session`, `Server`, `Request`, and `Response`.^[6] It is highly recommended that you specify a namespace for your web service before publishing it publicly because the default namespace, <http://tempuri.org/>, will not uniquely identify your web service from other web services. To do this, tag the `WebService` class with the `Namespace` attribute, specifying your own namespace.

^[6] Access to the `Request` and `Response` objects *can* be done through the `Context` property of the `WebService` class.

3. Tag the public methods with `WebMethod` attributes to make *web methods*—public methods of a distributed component that are accessible via the Web. You don't have to tag a method as `WebMethod` unless you want that method to be published as a web method.

The following C# code demonstrates a simple web service^[7] that exposes four methods to Internet clients.

^[7] For security reasons, the current release of ASP.NET runs as the account `ASPNET`. If you are using integrated security to access database resources, you must grant database access to the `ASPNET` account. You can also enable impersonation in the *web.config* or *machine.config* file:

```
<system.web>
```

```
<identity impersonate="true" userName="." password="" />
</system.web>
```

If you set impersonate to true but leave UserName and password blank, the application will run as *MachineName\USR_MachineName*, so make sure to grant this user (or whatever userName you specify) database access.

We emphasize "Internet" because anyone that can access this *asmx* file on the web server can access these methods, as opposed to your COM component, which can be accessed only by COM clients:

```
<%@ WebService Language="C#" Class="PubsWS.PubsWS" %>
```

```
namespace PubsWS
{
    using System;
    using System.Data;
    using System.Data.OleDb;
    using System.Web;
    using System.Web.Services;

    [WebService(Namespace="http://Oreilly/DotNetEssentials/")]
    public class PubsWS : WebService
    {
        private static string m_sConnStr =
"provider=sqloledb;server=(local);database=pubs; Integrated Security=SSPI";

        [WebMethod(Description="Returns a DataSet containing all authors.")]

        public DataSet GetAuthors( )
        {
            OleDbDataAdapter oDBAdapter;
            DataSet oDS;

            oDBAdapter = new OleDbDataAdapter("select * from authors",
                                             m_sConnStr);

            oDS = new DataSet( );
            oDBAdapter.Fill(oDS, "Authors");
            return oDS;
        }

        [WebMethod]

        public DataSet GetAuthor(string sSSN)
        {
            OleDbDataAdapter oDBAdapter;
            DataSet oDS;

            oDBAdapter = new OleDbDataAdapter(
```

```

        "select * from authors where au_id =" +
        + sSSN + "'", m_sConnStr);
    oDS = new DataSet( );
    oDBAdapter.Fill(oDS, "SelectedAuthor");
    return oDS;
}

```

```
[WebMethod(MessageName="GetBooksByAuthor" ,
```

```

    Description="Find books by author's SSN.")]
```

```

public DataSet GetBooks(string sAuthorSSN)
{
    OleDbDataAdapter oDBAdapter;
    DataSet oDS;

    oDBAdapter = new OleDbDataAdapter(
        "select * from titles inner join titleauthor on " +
        "titles.title_id=titleauthor.title_id " +
        "where au_id='" + sAuthorSSN + "'", m_sConnStr);
    oDS = new DataSet( );
    oDBAdapter.Fill(oDS, "Books");
    oDBAdapter = new OleDbDataAdapter("select * from authors " +
        "where au_id='" + sAuthorSSN + "'", m_sConnStr);
    oDBAdapter.Fill(oDS, "Author");

    return oDS;
}

```

```
[WebMethod]
```

```

public DataSet GetBooks( )
{
    OleDbDataAdapter oDBAdapter;
    DataSet oDS;

    oDBAdapter = new OleDbDataAdapter("select * from titles" ,
        m_sConnStr);

    oDS = new DataSet( );
    oDBAdapter.Fill(oDS, "Books");
    return oDS;
}

} // End PubsWS
}

```

If you are familiar with ASP, you may recognize the usage of the @ symbol in front of keyword `WebService`. This WebService directive specifies the language of the web service so that ASP.NET can compile the web service with the correct compiler. This directive also specifies the class that

implements the web service so it can load the correct class and employ reflection to generate the WSDL for the web service.

Because *PubsWS* also uses ADO.NET's OLE DB provider for its data-access needs, we have to add a reference to System.Data and System.Data.OleDb, in addition to the System, System.Web, and System.Web.Services namespaces.

Class PubsWS inherits from WebService with the colon syntax that should be familiar to C++ or C# developers:

```
public class PubsWS : WebService
```

The four methods that are tagged with WebMethod attributes are GetAuthors(), GetAuthor(), GetBooks(*string*), and GetBooks(). In C#, you can tag public methods with a WebMethod attribute using the [] syntax. In VB, you must use <>. For example, in VB, the second method would be declared as:

```
<WebMethod( )> Public Function GetAuthor(sSSN As String) As DataSet
```

By adding [WebMethod] in front of your public method, you make the public method callable from any Internet client. What goes on behind the scenes is that your public method is associated with an *attribute*, which is implemented as a WebMethodAttribute class. WebMethodAttribute has six properties:

BufferResponse (boolean)

Controls whether or not to buffer the method's response.

CacheDuration

Specifies the length of time in seconds to keep the method response in cache; the default is not to hold the method response in cache (0 seconds). A cache hit is when an identical call with identical parameters is requested. The cached response is used to avoid re-processing.

Description

Provides additional information about a particular web method.

EnableSession (boolean)

Enables or disables session state. If you don't want to use session state for the web method, you should make sure that this flag is disabled so the web server doesn't have to generate and manage session IDs for each user accessing the method.

MessageName

Distinguishes web methods with the same names. For example, if you have two different methods called GetBooks (one method retrieves all books while the other method retrieves only books written by a certain author) and you want to publish both of these methods as web methods, the system will have a problem trying to distinguish the two methods since their names are duplicated. You have to use the MessageName property to make sure all service signatures within the WSDL are unique. If the protocol is SOAP, MessageName is mapped to the SOAPAction request header and nested within the soap:Body element. For HTTP GET and HTTP POST, it is the PathInfo portion of the URI (as in *http://localhost//PubsWS/PubsWS.asmx/GetBooksByAuthor*).

TransactionOption

Can be one of five modes: Disabled, NotSupported, Supported, Required, and RequiresNew. Even though there are five modes, web methods can only participate as the root object in a transaction. This means both Required and RequiresNew result in a new transaction being created for the web method. The Disabled, NotSupported, and Supported settings result in no transaction being used for the web method. The TransactionOption property of a web method is set to Disabled by default.

To set up these properties, pass the property name and its value as a *name = value* pair:

```
[WebMethod(Description="Returns a DataSet containing all authors.")]
public DataSet GetAuthors( )
```

You can separate multiple properties with a comma:

```
[WebMethod(MessageName="GetBooksByAuthor",
            Description="Find books by author's SSN.")]
public DataSet GetBooks(string sAuthorSSN)
```

6.3.1.1 Web.Config

If you set up your web services from scratch, you should also need to provide the configuration file (web.config) in the same directory as your asmx file. This configuration file allows you to control various application settings about the virtual directory. Here, we set the authentication mode to None to make our web services development and testing a little easier. When you release your web services to the public, you should change this setting to Windows, Forms, or Passport instead of None:

```
<configuration>
  <system.web>
    <authentication mode="None" />
  </system.web>
</configuration>
```

The following list shows the different modes of authentication:

Forms

Basic Forms authentication is where unauthenticated requests are redirected to a login form.

Windows

Authentication is performed by IIS in one of three ways: basic, digest, or Integrated Windows Authentication.

Passport

Unauthenticated requests to the resource are redirected to Microsoft's centralized authentication service. When authenticated, a token is passed back and used by subsequent requests.

In v1.1 of .NET Framework, for security reasons, HttpGet and HttpPost protocol are disabled by default. To override the default and enable or disable any particular protocols, the web.config can be modified as the following:

```
<configuration>
```

```
. . .  
<webServices>  
  <protocols>  
    <add name="HttpGet" />  
    <add name="HttpPost" />  
  </protocols>  
</webServices>  
. . .  
</configuration>
```

6.3.1.2 Discover files

After creating the web service, you can provide the supporting files to help in the discovery of the service. The static discovery disco file is:[\[8\]](#)

[8] This code snippet assumes the virtual directory you set up is /PubsWS on your local web server.

```
<?xml version="1.0" ?>  
<disco:discovery xmlns:disco="http://schemas.xmlsoap.org/disco/"  
  xmlns:scl="http://schemas.xmlsoap.org/disco/scl/">  
<scl:contractRef ref="http://localhost/PubsWS/PubsWS.asmx?WSDL" />  
</disco:discovery>
```

[\[Team LiB \]](#)

[\[Team LiB \]](#)

6.4 Web Services Consumers

Now that you have successfully created a web service, let's take a look at how this web service is used by web clients. Web services clients communicate with web services through standard web protocols. They send and receive XML-encoded messages to and from the web services. This means any application on any platform can access the web services as long as it uses standard web protocols and understands the XML-encoded messages. As mentioned earlier, there are three protocols that the web clients can employ to communicate with the servers (web services): HTTP GET, HTTP POST, and SOAP. We demonstrate next how to build client applications that utilize each of these protocols. These web services-client applications are done in legacy languages such as VB6 and Perl,^[9] and .NET languages, such as C# and VB.NET, to demonstrate the cross-language/cross-platform benefits of web services.

^[9] We use SOAP: :Lite PERL modules. See <http://www.soaplite.com/>.

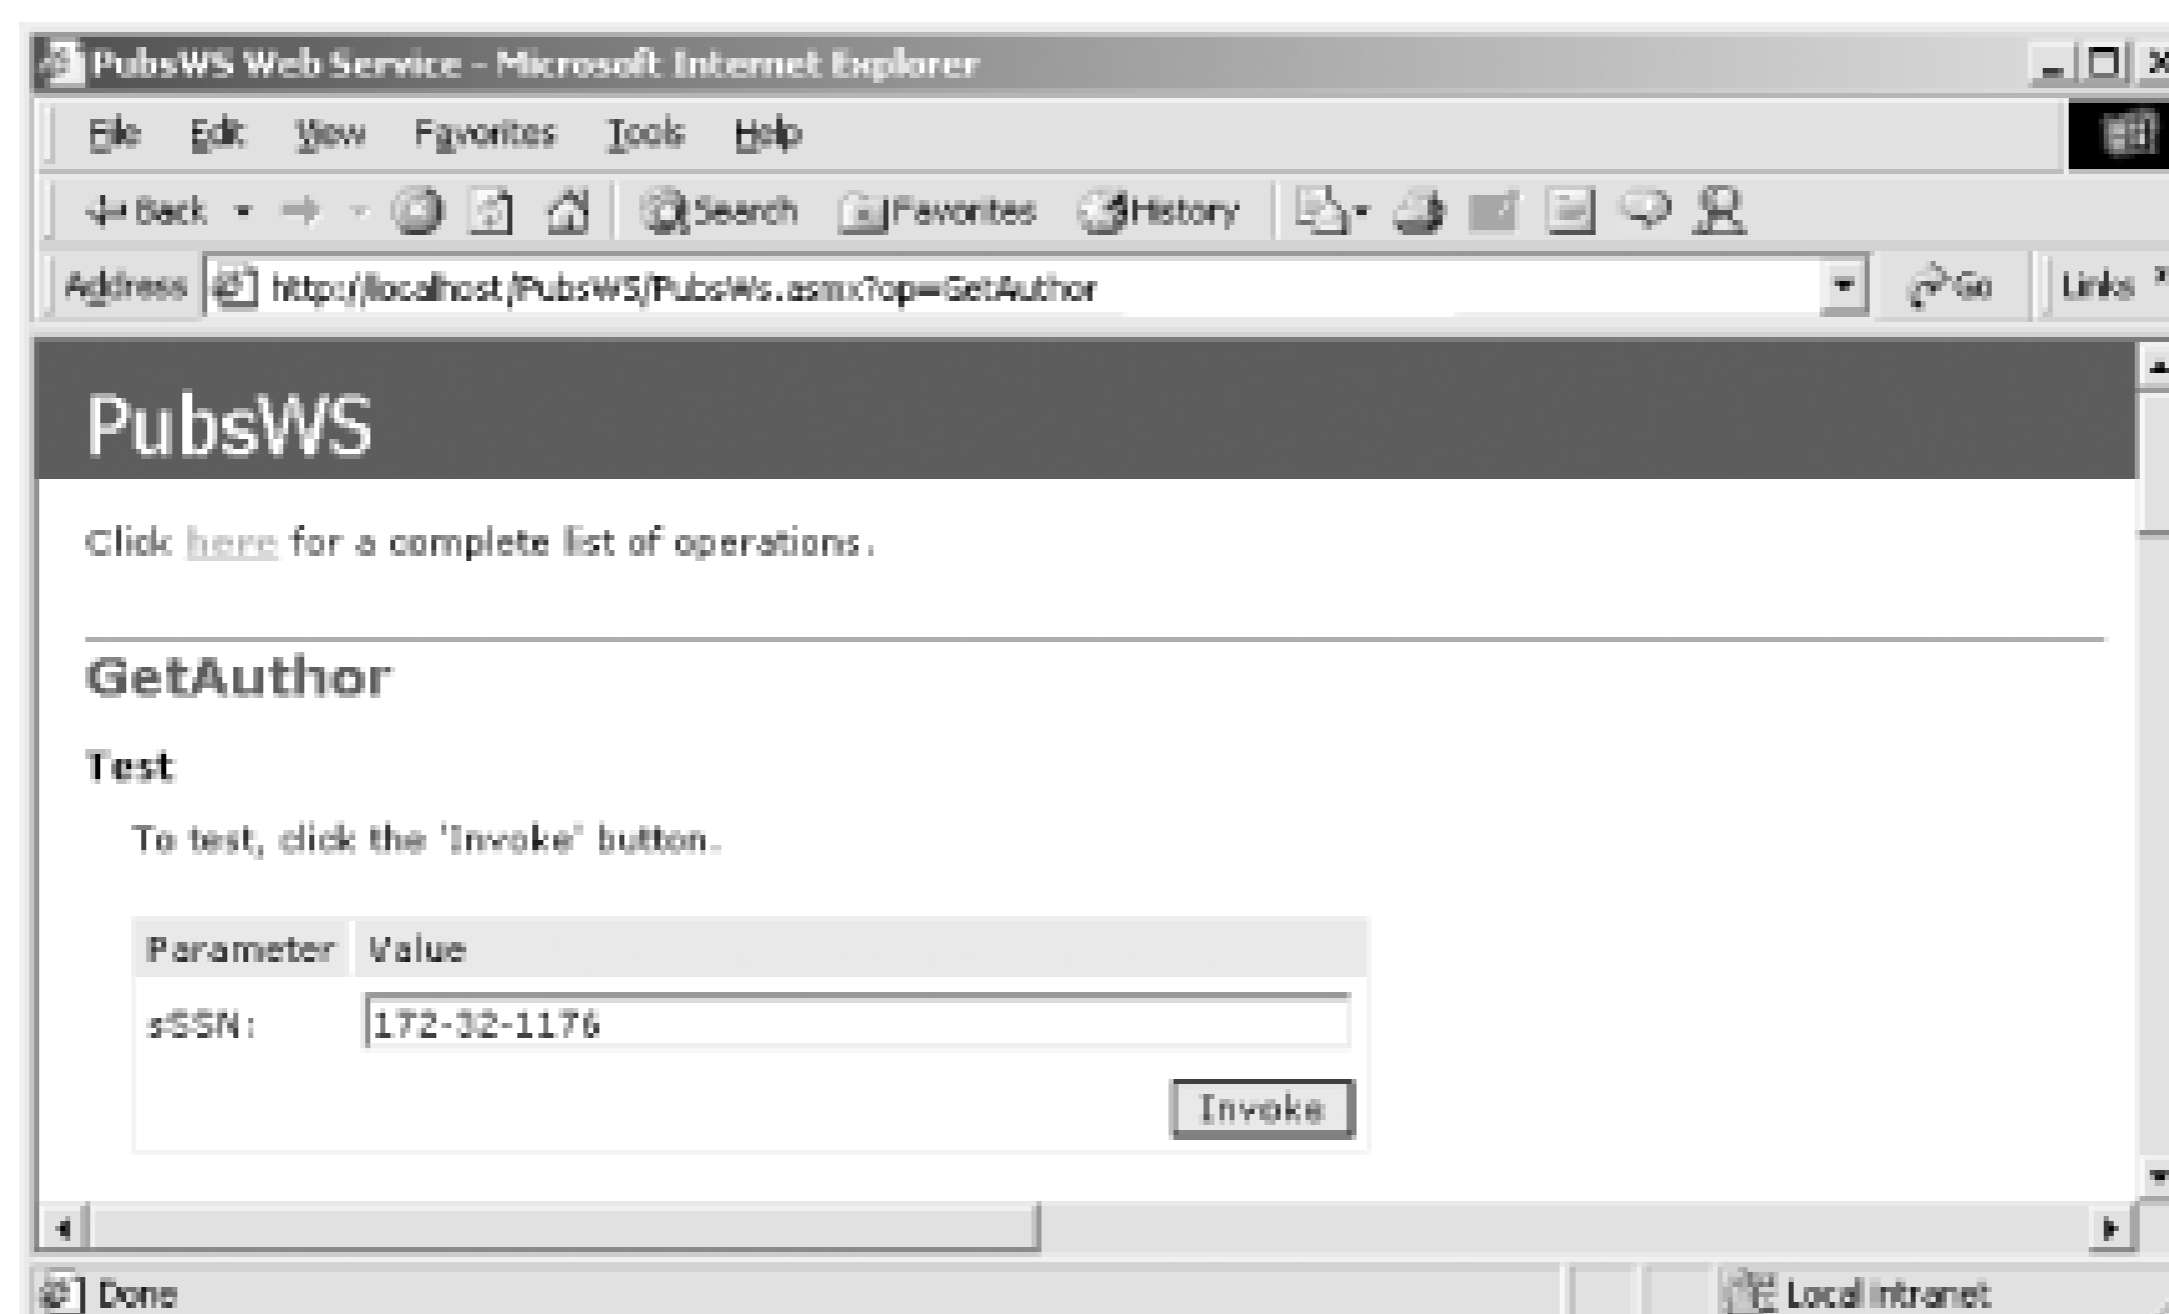
6.4.1 HTTP GET Consumer

Let's look at how it is done using HTTP GET first, since it is the simplest. In the examples that follow, we use *localhost* as the name of the web server running the service and *PubsWS* as the virtual directory. If you have deployed the sample web service on a remote server, you'll need to substitute the name of the server and virtual directory as appropriate.

If you point your web browser at the web service URL (<http://localhost/PubsWS/PubsWS.asmx>), it will give you a list of supported methods. To find out more about these methods, click one of them. This brings up a default web service consumer. This consumer, autogenerated through the use of reflection, is great for testing your web services' methods.^[10] It uses the HTTP GET protocol to communicate with the web service. This consumer features a form that lets you test the method (see [Figure 6-3](#)), as well as descriptions of how to access the method via SOAP, HTTP GET, or HTTP POST.

^[10] A simple reflection example can be found in [Section 4.3.1.3](#).

Figure 6-3. An autogenerated web services consumer



Here is the description of the GET request and response supplied by the default consumer:

The following is a sample HTTP GET request and response. The **placeholders shown need to be replaced with actual values**.

```
GET /PubsWS/PubsWS.asmx/GetAuthor?sSSN=string HTTP/1.1
Host: localhost
```

```
HTTP/1.1 200 OK
Content-Type: text/xml; charset=utf-8
Content-Length: length
```

```
<?xml version="1.0" encoding="utf-8"?>
<DataSet xmlns="http://Oreilly/DotNetEssentials/">
  <schema xmlns="http://www.w3.org/2001/XMLSchema">schema</schema>xml
</DataSet>
```

Using HTTP GET protocol, the complete URL to invoke the web method, along with parameters, can be the following:

```
http://localhost/PubsWS/PubsWS.asmx/GetAuthor?sSSN=172-32-1176
```

Here is the response; including HTTP response headers and the raw XML (note how the response includes the serialized schema and data from the DataSet object):

```
Cache-Control: private, max-age=0
Date: Tue, 08 May 2001 20:53:16 GMT
Server: Microsoft-IIS/5.0
Content-Length: 2450
Content-Type: text/xml; charset=utf-8
Client-Date: Tue, 08 May 2001 20:53:16 GMT
Client-Peer: 127.0.0.1:80
```

```
<?xml version="1.0" encoding="utf-8"?>
```

```

<DataSet xmlns="http://Oreilly/DotNetEssentials/">
  <xs:schema id="NewDataSet"
    xmlns=""
    xmlns:xs="http://www.w3.org/2001/XMLSchema"
    xmlns:msdata="urn:schemas-microsoft-com:xml-msdata">
    <xs:element name="NewDataSet" msdata:IsDataSet="true">
      <xs:complexType>
        <xs:choice maxOccurs="unbounded">
          <xs:element name="SelectedAuthor">
            <xs:complexType>
              <xs:sequence>
                <xs:element name="au_id" type="xs:string"
                  minOccurs="0" />
                <xs:element name="au_lname" type="xs:string"
                  minOccurs="0" />
                <xs:element name="au_fname" type="xs:string"
                  minOccurs="0" />
                <xs:element name="phone" type="xs:string"
                  minOccurs="0" />
                <xs:element name="address" type="xs:string"
                  minOccurs="0" />
                <xs:element name="city" type="xs:string"
                  minOccurs="0" />
                <xs:element name="state" type="xs:string"
                  minOccurs="0" />
                <xs:element name="zip" type="xs:string"
                  minOccurs="0" />
                <xs:element name="contract" type="xs:boolean"
                  minOccurs="0" />
              </xs:sequence>
            </xs:complexType>
          </xs:element>
        </xs:choice>
      </xs:complexType>
    </xs:element>
  </xs:schema>
  <diffgr:diffgram
    xmlns:msdata="urn:schemas-microsoft-com:xml-msdata"
    xmlns:diffgr="urn:schemas-microsoft-com:xml-diffgram-v1">
    <NewDataSet xmlns="">
      <SelectedAuthor diffgr:id="SelectedAuthor1" msdata:rowOrder="0">
        <au_id>172-32-1176</au_id>
        <au_lname>White</au_lname>
        <au_fname>Johnson</au_fname>
        <phone>408 496-7223</phone>
        <address>10932 Bigge Rd.</address>
        <city>Menlo Park</city>
        <state>CA</state>
        <zip>94025</zip>
        <contract>true</contract>
      </SelectedAuthor>
    </NewDataSet>
  </diffgr:diffgram>
</DataSet>

```

```
</diffgr:diffgram>
</DataSet>
```

6.4.2 HTTP POST Consumer

In the section "HTTP GET Consumer," we saw the automatic creation of a web services consumer by merely hitting the URL of the web services, *http://localhost/PubsWS/PubsWS.asmx*. It is now time for us to see how a web client can use HTTP POST and SOAP to access a web service. This time around, we are going to write a C# web service consumer.

The Microsoft .NET SDK has a rich set of tools to simplify the process of creating or consuming web services. We are going to use one of these tools, `wSDL`, to generate source code for the proxies to the actual web services:[\[11\]](#)

[11] `wSDL.exe` generates the proxy source code similar to the way IDL compilers generate source files for DCOM proxies. The only difference is that WSDL is the language that describes the interface of the software component, which is XML-based.

```
wSDL /l:CS /protocol:HttpPost http://localhost/PubsWS/PubsWS.asmx?WSDL
```

This command line creates a proxy for the PubsWS web service from the WSDL document from the URL *http://localhost/PubsWS/PubsWS.asmx?WSDL*. The proxy uses HTTP POST as its protocol to talk to the web service; it is generated as a C# source file. The `wSDL` tool can also take a WSDL file as its input instead of a URL pointing to the location where the WSDL can be obtained.

This C# proxy source file represents the proxy class for the PubsWS web service that the clients can compile against. This generated C# file contains a proxy class `PubsWS` that derives from `HttpPostClientProtocol` class. If you use the `/protocol:HttpGet` or `/protocol:SOAP2`[\[12\]](#) parameters, the `PubsWS` derives from either the `HttpGetClientProtocol` or `SoapHttpClientProtocol` class.

[12] In .NET Framework 1.1, SOAP 1.2 is also supported via `/protocol:SOAP12`. All this does is generate the SOAP derived proxy for SOAP Version 1.2.

After generating the C# source file *PubsWS.cs*, we have two choices for how this proxy can be used. One way is to include this source file in the client application project using Visual Studio .NET. The project has to be a C# project if you choose this route.[\[13\]](#) To make use of the proxy, you also have to add to your project any references that the proxy depends on. In this example, the necessary references for the proxy file are `System.Web.Services`, `System.Web.Services.Protocols`, `System.Xml.Serialization`, and `System.Data`.

[13] For other languages, use `wSDL` with the `/l` option to specify the language. See [Appendix D](#) for more details.

The other way to use the proxy is more flexible. You can compile the C# source file into a dynamic link library (DLL) and then add a reference to this DLL to any project you want to create. This way you can even have a VB project use the DLL.

Below is the command line used to compile the C# proxy source into a DLL. Notice that the three references are linked to *PubsWS.cs* so that the resulting *PubsWS.DLL* is self-contained (type the entire command on one line):

```
csc /t:library
/r:system.web.services.dll
/r:system.xml.dll
/r:system.data.dll
```

PubsWS.cs

Regardless of how you choose to use the proxy, the client application code will still look the same. Consider the next two code examples containing C# and VB code. For both languages, the first lines create an instance of the proxy to the web service, PubsWS. The second lines invoke the GetAuthors web method to get a DataSet as the result. The remaining lines bind the default view of the table *Authors* to the data grid, add the data grid to a form, and display the form. Note that these examples use the Windows Forms API, which we'll discuss in [Chapter 8](#). Here is the C# web service client, *TestProxy.cs*:

```
using System;
using System.Drawing;
using System.Windows.Forms;
using System.Data;

public class TestProxy
{
    public static void Main( )
    {
        /* Create a proxy. */
        PubsWS oProxy = new PubsWS( );

        /* Invoke GetAuthors( ) over HTTPPOST and get the data set. */
        DataSet oDS = oProxy.GetAuthors( );

        /* Create a data grid and connect it to the data set. */
        DataGrid dg = new DataGrid( );
        dg.Size = new Size(490, 270);
        dg.DataSource = oDS.Tables["Authors"].DefaultView;

        /* Set the properties of the form and add the data grid. */
        Form myForm = new Form( );
        myForm.Text = "DataGrid Sample";
        myForm.Size = new Size(500, 300);
        myForm.Controls.Add(dg);

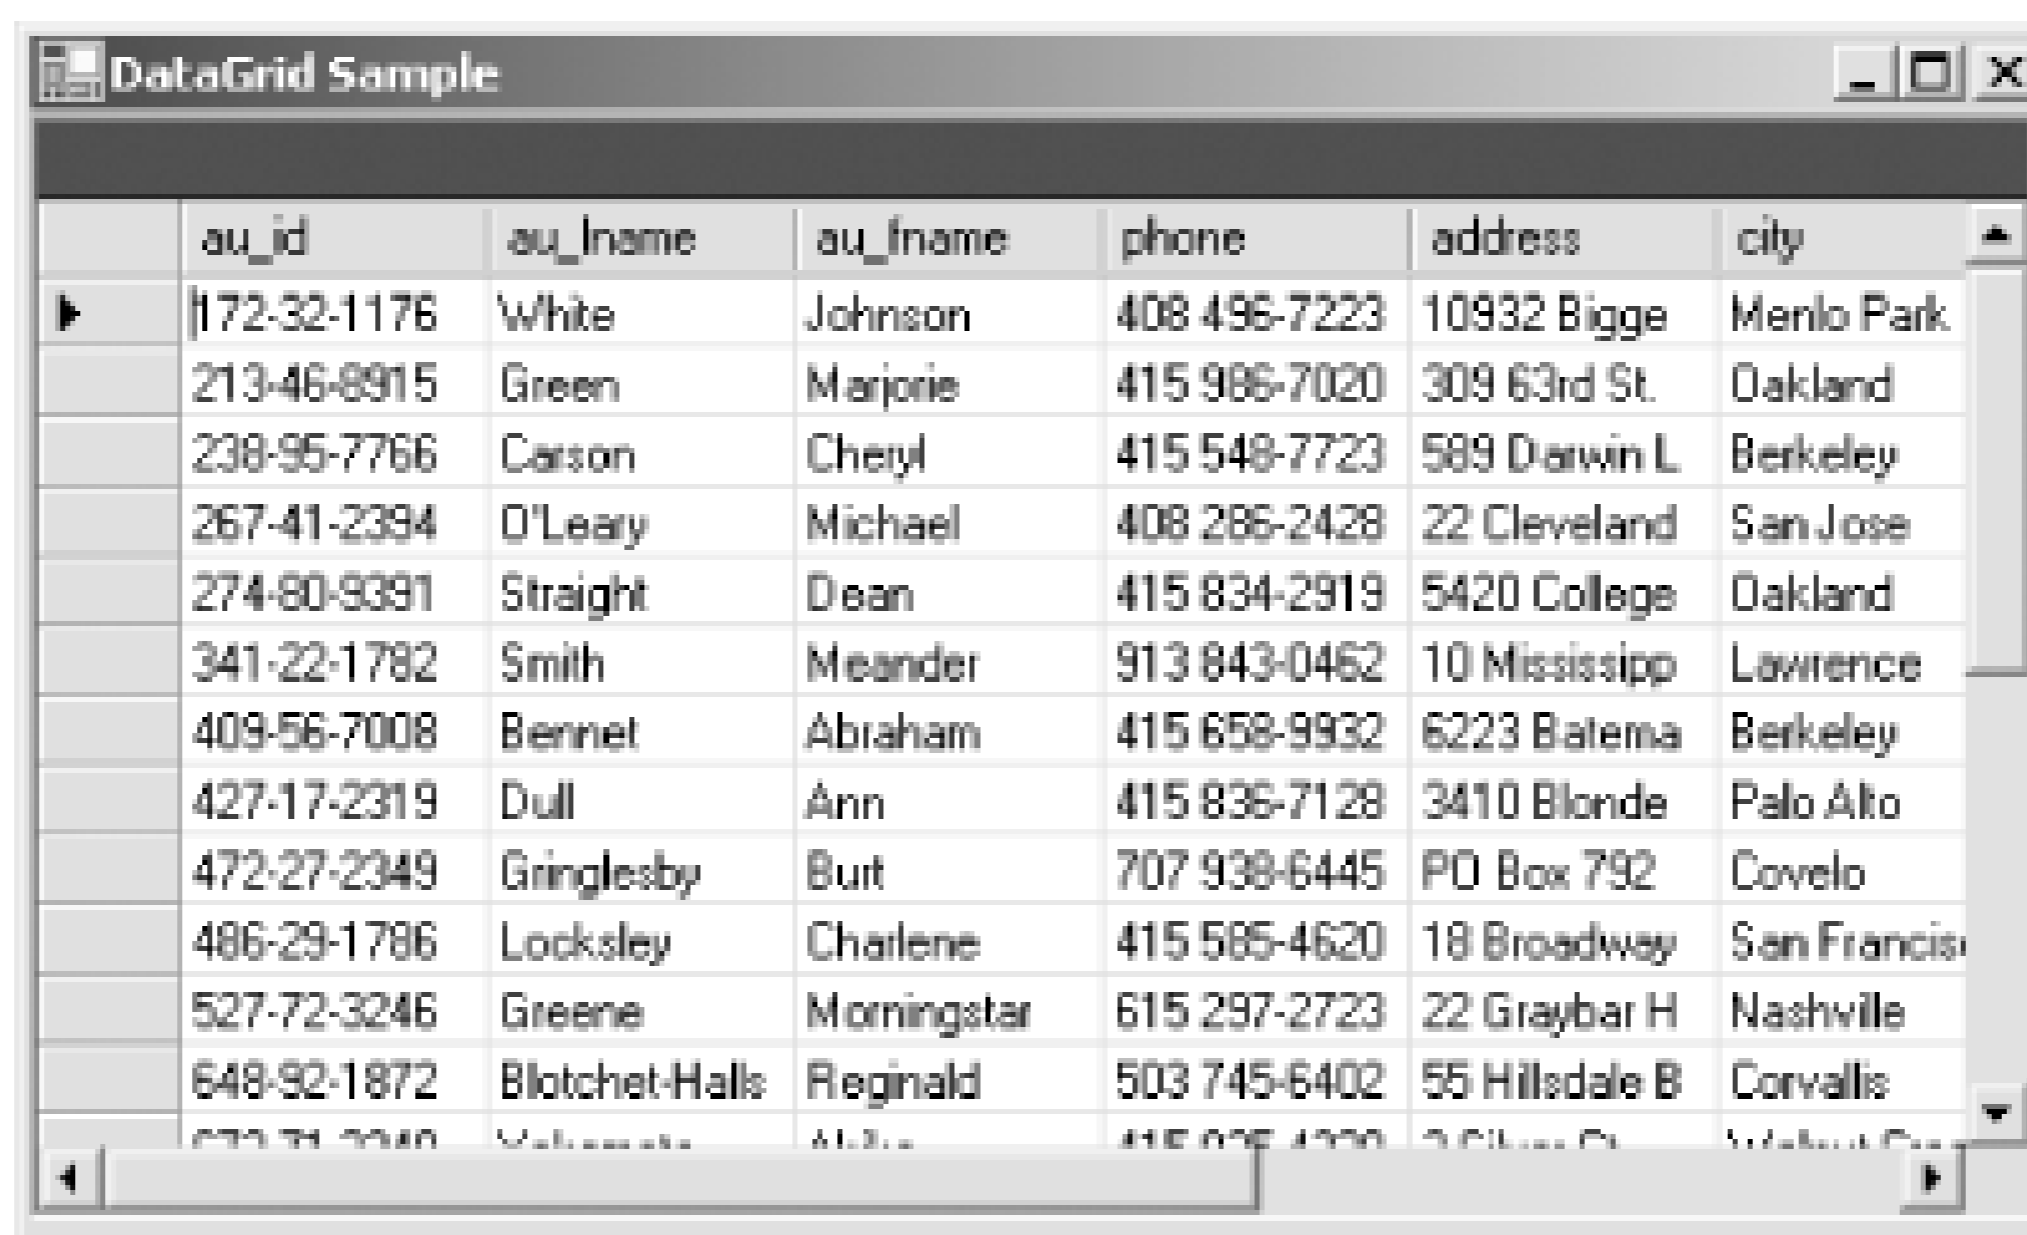
        /* Display the form. */
        System.Windows.Forms.Application.Run(myForm);
    }
}
```

If you created the DLL as previously directed, you can compile this with the following command:

```
csc TestProxy.cs /r:PubsWS.dll
```

This creates the executable *TestProxy.exe*, which gets a DataSet using a HTTP POST call, and displays a data grid containing that dataset. [Figure 6-4](#) shows the output of the C# client after obtaining the data from the PubsWS web service via HTTP POST protocol.

Figure 6-4. C# web service client after calling GetAuthors()



au_id	au_lname	au_fname	phone	address	city
172-32-1176	White	Johnson	408 496-7223	10932 Bigge	Menlo Park
213-46-8915	Green	Marjorie	415 986-7020	309 63rd St.	Oakland
238-95-7766	Carson	Cheryl	415 548-7723	589 Darwin L.	Berkeley
267-41-2394	O'Leary	Michael	408 286-2428	22 Cleveland	San Jose
274-80-9391	Straight	Dean	415 834-2919	5420 College	Oakland
341-22-1782	Smith	Meander	913 843-0462	10 Mississipp	Lawrence
409-56-7008	Bennet	Abraham	415 658-9932	6223 Batema	Berkeley
427-17-2319	Dull	Ann	415 836-7128	3410 Blonde	Palo Alto
472-27-2349	Gringlesby	Burt	707 938-6445	PO Box 792	Covelo
486-29-1786	Locksley	Charlene	415 585-4620	18 Broadway	San Francis
527-72-3246	Greene	Morningstar	615 297-2723	22 Graybar H	Nashville
648-92-1872	Blotchet-Halls	Reginald	503 745-6402	55 Hilldale B	Corvallis

Here is the VB web service client, *TestProxyVB.vb*.

```

imports System
imports System.Drawing
imports System.Windows.Forms
imports System.Data

Module TestProxyVB
  Sub Main( )
    ' Create a proxy.
    dim oProxy as PubsWS = new PubsWS( )

    ' Invoice GetAuthors( ) over SOAP and get the data set.
    dim oDS as DataSet = oProxy.GetAuthors( )

    ' Create a data grid and connect it to the data set.
    dim dg as DataGrid = new DataGrid( )
    dg.Size = new Size(490, 270)
    dg.DataSource = oDS.Tables("Authors").DefaultView

    ' Set the properties of the form and add the data grid.
    dim myForm as Form = new Form( )
    myForm.Text = "DataGrid Sample"
    myForm.Size = new Size(500, 300)
    myForm.Controls.Add(dg)

    ' Display the form.
    System.Windows.Forms.Application.Run(myForm)
  End Sub
End Module

```

You can compile the VB web service client with this command (type the entire command on one line)

```
vbc TestProxyVB.vb
```

```

/r:System.Drawing.dll
/r:System.Windows.Forms.dll
/r:System.Data.dll
/r:PubsWS.dll
/r:System.Web.Services.dll
/r:System.dll
/r:System.xml.dll

```

Instead of using `wsdl` to generate and include the proxy in your application, you can also rely on VS.NET to automate the whole process. In VS.NET, you can just add a Web Reference to your application. The process of adding a Web Reference to an application involves the discovery of the web service, obtaining the WSDL, generating the proxy, and including the proxy into the application. [\[14\]](#)

[14] You can find the proxy source file under Web References\ReferenceName as *reference.cs* (if you're working with C#). If you have not renamed the reference name, it is *localhost* by default. (You might have to select the option to "show all files" in VS.NET Solution Explorer.)

6.4.3 Non-.NET Consumers

This section shows how to develop non-.NET web service consumers using HTTP GET, HTTP POST, and SOAP protocols. Because we cannot just create the proxy class from the WSDL and compile it with the client code directly, we must look at the WSDL file to understand how to construct and interpret messages between the web service and the clients. We trimmed down the WSDL file for our PubsWS web service to show only types, messages, ports, operations, and bindings that we actually use in the next several web service-client examples. In particular, we will have our VB6 client access the following:

Web method	Protocol
GetBooks()	HTTP GET protocol
GetAuthor(ssn)	HTTP POST protocol
GetBooksByAuthor(ssn)	SOAP protocol

As a reference, here is the simplified version of the WSDL file while you experiment with the VB6 client application:

```

<?xml version="1.0" encoding="utf-8"?>
<definitions xmlns: . . .
  xmlns:s0="http://Oreilly/DotNetEssentials/"
  targetNamespace="http://Oreilly/DotNetEssentials/" >

  <types>
    <!-- This data type is used by the HTTP POST call -->
    <s:element name="GetAuthor">
      <s:complexType>
        <s:sequence>
          <s:element minOccurs="1" maxOccurs="1"
            name="sSSN" type="s:string" />

```

```

        </s:sequence>
    </s:complexType>
</s:element>
<!-- This data type is used by the HTTP POST call -->
<s:element name="GetAuthorResponse">
    <s:complexType>
        <s:sequence>
            <s:element minOccurs="1" maxOccurs="1"
                name="GetAuthorResult">
                <s:complexType>
                    <s:sequence>
                        <s:element ref="s:schema" />
                        <s:any />
                    </s:sequence>
                </s:complexType>
            </s:element>
        </s:sequence>
    </s:complexType>
</s:element>

<!-- This data type is used by the SOAP call -->
<s:element name="GetBooksByAuthor">
    <s:complexType>
        <s:sequence>
            <s:element minOccurs="1" maxOccurs="1"
                name="sAuthorSSN" type="s:string" />
        </s:sequence>
    </s:complexType>
</s:element>

<!-- This data type is used by the SOAP call -->
<s:element name="GetBooksByAuthorResponse">
    <s:complexType>
        <s:sequence>
            <s:element minOccurs="1" maxOccurs="1"
                name="GetBooksByAuthorResult">
                <s:complexType>
                    <s:sequence>
                        <s:element ref="s:schema" />
                        <s:any />
                    </s:sequence>
                </s:complexType>
            </s:element>
        </s:sequence>
    </s:complexType>
</s:element>

<!-- This data type is used by the HTTP GET call -->
<s:element name="GetBooks">
    <s:complexType />
</s:element>
<!-- This data type is used by the HTTP GET call -->
<s:element name="GetBooksResponse">

```



```

    <s:complexType>
      <s:sequence>
        <s:element minOccurs="1" maxOccurs="1"
          name="GetBooksResult">
          <s:complexType>
            <s:sequence>
              <s:element ref="s:schema" />
              <s:any />
            </s:sequence>
          </s:complexType>
        </s:element>
      </s:sequence>
    </s:complexType>
  </s:element>

  <!-- This data type is used by the HTTP GET/POST responses -->
  <s:element name="DataSet"
    <s:complexType>
      <s:sequence>
        <s:element ref="s:schema" />
        <s:any />
      </s:sequence>
    </s:complexType>
  </s:element>

</types>

<!-- These messages are used by the SOAP call -->
<message name="GetBooksByAuthorSoapIn">
  <part name="parameters" element="s0:GetBooksByAuthor" />
</message>
<message name="GetBooksByAuthorSoapOut">
  <part name="parameters" element="s0:GetBooksByAuthorResponse" />
</message>

<!-- These messages are used by the HTTP GET call -->
<message name="GetBooksHttpGetIn" />
<message name="GetBooksHttpGetOut">
  <part name="Body" element="s0:DataSet" />
</message>

<!-- These messages are used by the HTTP POST call -->
<message name="GetAuthorHttpPostIn">
  <part name="sSSN" type="s:string" />
</message>
<message name="GetAuthorHttpPostOut">
  <part name="Body" element="s0:DataSet" />
</message>

<!-- SOAP port -->
<portType name="PubsWSSoap">
  <operation name="GetBooks">

```

```

        <documentation>Find books by author's SSN.</documentation>
        <input name="GetBooksByAuthor"
            message="s0:GetBooksByAuthorSoapIn" />
        <output name="GetBooksByAuthor"
            message="s0:GetBooksByAuthorSoapOut" />
    </operation>
</portType>

<!-- HTTP GET port -->
<portType name="PubsWSHttpGet">
    <operation name="GetBooks">
        <input message="s0:GetBooksHttpGetIn" />
        <output message="s0:GetBooksHttpGetOut" />
    </operation>
</portType>

<!-- HTTP POST port -->
<portType name="PubsWSHttpPost">
    <operation name="GetAuthor">
        <input message="s0:GetAuthorHttpPostIn" />
        <output message="s0:GetAuthorHttpPostOut" />
    </operation>
</portType>

<!-- SOAP binding -->
<binding name="PubsWSSoap" type="s0:PubsWSSoap">
    <soap:binding
        transport="http://schemas.xmlsoap.org/soap/http"
        style="document" />
    <operation name="GetBooks">
        <soap:operation
            soapAction="http://Oreilly/DotNetEssentials/GetBooksByAuthor"
            style="document" />
        <input name="GetBooksByAuthor">
            <soap:body use="literal" />
        </input>
        <output name="GetBooksByAuthor">
            <soap:body use="literal" />
        </output>
    </operation>
</binding>

<!-- HTTP GET binding -->
<binding name="PubsWSHttpGet" type="s0:PubsWSHttpGet">
    <http:binding verb="GET" />
    <operation name="GetBooks">
        <http:operation location="/GetBooks" />
        <input>
            <http:urlEncoded />
        </input>
        <output>
            <mime:mimeXml part="Body" />
        </output>
    </operation>
</binding>

```

```

        </output>
    </operation>
</binding>

<!-- HTTP POST binding -->
<binding name="PubsWSHttpPost" type="s0:PubsWSHttpPost">
    <http:binding verb="POST" />
    <operation name="GetAuthor">
        <http:operation location="/GetAuthor" />
        <input>
            <mime:content type="application/x-www-form-urlencoded" />
        </input>
        <output>
            <mime:mimeXml part="Body" />
        </output>
    </operation>
</binding>

<!-- The whole web service and address bindings -->
<service name="PubsWS">

    <port name="PubsWSSoap" binding="s0:PubsWSSoap">
        <soap:address location="http://localhost/PubsWS/PubsWS.asmx" />
    </port>

    <port name="PubsWSHttpGet" binding="s0:PubsWSHttpGet">
        <http:address location="http://localhost/PubsWS/PubsWS.asmx" />
    </port>

    <port name="PubsWSHttpPost" binding="s0:PubsWSHttpPost">
        <http:address location="http://localhost/PubsWS/PubsWS.asmx" />
    </port>

</service>

</definitions>

```

In both the HTTP GET and HTTP POST protocols, you pass parameters to the web services as name/value pairs. With the HTTP GET protocol, you must pass parameters in the query string, whereas the HTTP POST protocol packs the parameters in the body of the request package. To demonstrate this point, we will construct a simple VB client using both HTTP GET and HTTP POST protocols to communicate with the PubsWS web service.

Let's first create a VB6 standard application. We need to add a reference to Microsoft XML, v3.0 (*msxml3.dll*), because we'll use the XMLHTTP object to help us communicate with the web services. For demonstrative purposes, we will also use the Microsoft Internet Controls component (*shdocvw.dll*) to display XML and HTML content.

First, add two buttons on the default form, `form1`, and give them the captions `GET` and `POST`, as well as the names `cmdGet` and `cmdPost`, respectively. After that, drag the WebBrowser object from the toolbar onto the form, and name the control `myWebBrowser`. If you make the WebBrowser navigate to `about:blank` initially, you will end up with something like [Figure 6-5](#).

Figure 6-5. VB client form to test web services



Now all we need is some code similar to the following to handle the two buttons' click events:

```

Private Sub cmdGet_Click( )
    Dim oXMLHTTP As XMLHTTP
    Dim oDOM As DOMDocument
    Dim oXSL As DOMDocument

    ' Call the web service to get an XML document
    Set oXMLHTTP = New XMLHTTP
    oXMLHTTP.open "GET", _
        "http://localhost/PubsWS/PubsWS.asmx/GetBooks", _
        False
    oXMLHTTP.send
    Set oDOM = oXMLHTTP.responseXML

    ' Create the XSL document to be used for transformation
    Set oXSL = New DOMDocument
    oXSL.Load App.Path & "\templateTitle.xsl"

    ' Transform the XML document into an HTML document and display
    myWebBrowser.Document.Write CStr(oDOM.transformNode(oXSL))
    myWebBrowser.Document.Close

    Set oXSL = Nothing
    Set oDOM = Nothing
    Set oXMLHTTP = Nothing
End Sub

Private Sub cmdPost_Click( )
    Dim oXMLHTTP As XMLHTTP
    Dim oDOM As DOMDocument
    Dim oXSL As DOMDocument

    ' Call the web service to get an XML document
    Set oXMLHTTP = New XMLHTTP
    oXMLHTTP.open "POST", _
        "http://localhost/PubsWS/PubsWS.asmx/GetAuthor", _
        False
    oXMLHTTP.setRequestHeader "Content-Type", _
        "application/x-www-form-urlencoded"
    oXMLHTTP.send "sSSN=172-32-1176 "

```

```

Set oDOM = oXMLHTTP.responseXML

' Create the XSL document to be used for transformation
Set oXSL = New DOMDocument
oXSL.Load App.Path & "\templateAuthor.xsl"

' Transform the XML document into an HTML document and display
myWebBrowser.Document.Write oDOM.transformNode(oXSL)
myWebBrowser.Document.Close

Set oXSL = Nothing
Set oDOM = Nothing
Set oXMLHTTP = Nothing
End Sub

```

The two subroutines are similar in structure, except that the first one uses the HTTP GET protocol and the second one uses the HTTP POST protocol to get to the PubsWS web service. Let's take a closer look at what the two subroutines do.

For the HTTP GET protocol, we use the XMLHTTP object to point to the URL for the web method, as specified in the WSDL. Since the GetBooks web method does not require any parameters, the query string in this case is empty. The method is invoked synchronously because the `async` parameter to XMLHTTP's `open` method is set to `false`. After the method invocation is done, we transform the XML result using `templateTitle.xsl` and display the HTML on the `myWebBrowser` instance on the form. [Figure 6-6](#) displays the screen of our web services testing application after invoking the GetBooks web method at URL `http://localhost/PubsWS/PubsWS.asmx/` through HTTP GET protocol.

Figure 6-6. VB client form after calling GetBooks

For the HTTP POST protocol, we also point the XMLHTTP object to the URL for the web method—in this case, method GetAuthor. Because this is a POST request, we have to specify in the HTTP header that the request is coming over as a form by setting the `Content-Type` header variable to `application/x-www-form-urlencoded`. If this variable is not set, XMLHTTP by default passes the data to the server in XML format.

Another difference worth noticing is that the GetAuthor method requires a single parameter, which is the SSN of the author as a string. Since this is a post request, we are going to send the name/value pair directly to the server in the body of the message. Because the `Content-Type` header has been set to `application/x-www-form-urlencoded`, the server will know how to get to the parameters and perform the work requested. This time, we use `templateAuthor.xsl` to transform the XML result to HTML and display it. [Figure 6-7](#) shows our application after invoking the GetAuthor web method of PubsWS web service through HTTP POST protocol.

Figure 6-7. VB client form after calling GetAuthor



The following code is the XSL used to transform the XML result from the GetBooks web method call to HTML to be displayed on the web browser instance on the VB form:

```
<html version="1.0" xmlns:xsl="http://www.w3.org/TR/WD-xsl">
<head><title>A list of books</title></head>
<style>
.hdr { background-color=#ffeedd; font-weight:bold; }
</style>
<body>
<B>List of books</B>
<table style="border-collapse:collapse" border="1">
<tr>
<td class="hdr">Title</td>
<td class="hdr">Type</td>
<td class="hdr">Price</td>
<td class="hdr">Notes</td>
</tr>
<xsl:for-each select="//Books">
<tr>
<td><xsl:value-of select="title"/></td>
<td><xsl:value-of select="type"/></td>
<td><xsl:value-of select="price"/></td>
<td><xsl:value-of select="notes"/></td>
</tr>
</xsl:for-each>
</table>
</body>
</html>
```

Here is the XSL used to transform the XML result from the GetAuthor web method call to HTML to be displayed on the web browser instance on the VB form:

```

<html version="1.0" xmlns:xsl="http://www.w3.org/TR/WD-xsl">
<head><title>Selected author</title></head>
<STYLE>
.hdr { background-color:'#ffeedd';
      text-align:'right'; vertical-align:'top';
      font-weight:bold; }
</STYLE>
<body>
<B>Selected author</B>
<xsl:for-each select="//SelectedAuthor">
<table style="border-collapse:'collapse'" border="1">
<tr><td class="hdr">ID</td>
      <td><xsl:value-of select="au_id"/></td></tr>
<tr><td class="hdr">Name</td>
      <td><xsl:value-of select="au_fname"/>
          <xsl:value-of select="au_lname"/></td></tr>
<tr><td class="hdr">Address</td>
      <td><xsl:value-of select="address"/><br>
          <xsl:value-of select="city"/>,
          <xsl:value-of select="state"/>
          <xsl:value-of select="zip"/></td></tr>
<tr><td class="hdr">Phone</td>
      <td><xsl:value-of select="phone"/></td></tr>
</table>
</xsl:for-each>
</body>
</html>

```

We can also use SOAP protocol to access the web service. Because the web service is exposed through HTTP and XML, any clients on any platform can access the service as long as they conform to the specification of the service. Again, this specification is the WSDL file. By inspecting the WSDL file-specifically, the SOAP section-we can use XMLHTTP again to communicate in SOAP dialog. Let's see how this can be done.

Let's go back to the example of consumer web services using VB6 and XMLHTTP. Add another button on the form, and call it `cmdSOAP` with caption `SOAP`. This time, we will ask the web service to return all books written by a particular author:

```

Private Sub cmdSOAP_Click( )
    Dim oXMLHTTP As XMLHTTP
    Dim oDOM As DOMDocument
    Dim oXSL As DOMDocument

    ' Call the web service to get an XML document
    Set oXMLHTTP = New XMLHTTP
    oXMLHTTP.open "POST", "http://localhost/PubsWS/PubsWS.asmx", False

    Dim sBody As String

    sBody = "" & _
        "<soap:Envelope" & _
        " xmlns:xsi=""http://www.w3.org/2001/XMLSchema-instance"" & _

```

```

" xmlns:xsd=""http://www.w3.org/2001/XMLSchema"" & _
" xmlns:soap=""http://schemas.xmlsoap.org/soap/envelope/"">" & _
"<soap:Body>" & _
"<GetBooksByAuthor xmlns=""http://Oreilly/DotNetEssentials/"">" & _
"<sAuthorSSN>213-46-8915</sAuthorSSN>" & _
"</GetBooksByAuthor>" & _
"</soap:Body>" & _
"</soap:Envelope>"

oXMLHTTP.setRequestHeader "Content-Type", "text/xml "
oXMLHTTP.setRequestHeader "SOAPAction",
    "http://Oreilly/DotNetEssentials/GetBooksByAuthor"

oXMLHTTP.send sBody

Set oDOM = oXMLHTTP.responseXML

' Create the XSL document to be used for transformation
Set oXSL = New DOMDocument
oXSL.Load App.Path & "\templateAuthorTitle.xsl"

' Transform the XML document into an HTML document
myWebBrowser.Document.Write oDOM.transformNode(oXSL)
myWebBrowser.Document.Close

Set oXSL = Nothing
Set oDOM = Nothing
Set oXMLHTTP = Nothing
End Sub

```

This method is structurally similar to the ones used for HTTP GET and HTTP POST; however, it has some very important differences. In SOAP, you have to set the `Content-Type` to `text/xml` instead of `application/x-www-form-urlencoded` as for the HTTP POST. By this time, it should be clear to you that only HTTP POST and SOAP care about the `Content-Type` because they send the data in the body of the HTTP request. The HTTP GET protocol does not really care about the `Content-Type` because all of the parameters are packaged into the query string. In addition to the difference in format of the data content, you also have to refer to the WSDL to set the `SOAPAction` header variable to the call you want. Looking back at the SOAP section of the WSDL, if you want to call the `GetBooks(sAuthorSSN)` method of the web service, you will set the `SOAPAction` header variable to `http://Oreilly/DotNetEssentials/GetBooksByAuthor`. On the other hand, if you want to call the `GetBooks()` method instead, the `SOAPAction` variable has to be set to `http://Oreilly/DotNetEssentials/GetBooks`. The reason the namespace is `http://Oreilly/DotNetEssentials/` is because we set it up as the attribute of the PubsWS web service class.

After setting up the header variables, pass the parameters to the server in the body of the message. While HTTP POST passes the parameters in name/value pairs, SOAP passes the parameters in a well-defined XML structure:

```

<soap:Envelope . . . namespace omitted . . . >
  <soap:Body>
    <GetBooksByAuthor xmlns=""http://Oreilly/DotNetEssentials/"">

```



```

    <sAuthorSSN>213-46-8915</sAuthorSSN>
  </GetBooksByAuthor>
</soap:Body>
</soap:Envelope>

```

Both the SOAP request and response messages are packaged within a **Body** inside an **Envelope**. With the previously specified request, the response SOAP message looks like this:

```

<?xml version="1.0"?>
<soap:Envelope . . . namespace omitted . . . >
  <soap:Body>
    <GetBooksByAuthorResult xmlns="http://Oreilly/DotNetEssentials/">
      <result>
        <xsd:schema id="NewDataSet" . . . >
          < . . . content omitted . . . >
        </xsd:schema>
        <NewDataSet xmlns="">
          <Books>
            <title_id>BU1032</title_id>
            <title>The Busy Executive's Database Guide</title>
            < . . . more . . . >
          </Books>
          <Books>
            <title_id>BU2075</title_id>
            <title>You Can Combat Computer Stress!</title>
            < . . . more . . . >
          </Books>
          <Author>
            <au_id>213-46-8915</au_id>
            <au_lname>Green</au_lname>
            <au_fname>Marjorie</au_fname>
            <phone>415 986-7020</phone>
            <address>309 63rd St. #411</address>
            <city>Oakland</city>
            <state>CA</state>
            <zip>94618</zip>
            <contract>True</contract>
          </Author>
        </NewDataSet>
      </result>
    </GetBooksByAuthorResult>
  </soap:Body>
</soap:Envelope>

```

[Figure 6-8](#) shows the result of the test form after invoking the `GetBooksByAuthor` web method using the SOAP protocol.

Figure 6-8. VB client form after calling `GetBooksByAuthor`

List of books written by <i>Marjorie Green (Oakland, CA)</i>			
Title	Type	Price	Notes
The Busy Executive's Database Guide	business	19.99	An overview of available database systems with emphasis on common business applications. Illustrated.
You Can Combat Computer Stress!	business	2.99	The latest medical and psychological techniques for living with the electronic office. Easy-to-understand explanations.

The XSL stylesheet used for transformation of the resulting XML to HTML is included here for your reference. Notice that since GetBooksByAuthor returns two tables in the dataset, author and books, we can display both the author information and the books that this author wrote:

```
<html version="1.0" xmlns:xsl="http://www.w3.org/TR/WD-xsl">
<head><title>A list of books</title></head>
<style>
.hdr { background-color=#ffeedd; font-weight:bold; }
</style>
<body>
<B>List of books written by
  <I><xsl:value-of select="//Author/au_fname"/>
    <xsl:value-of select="//Author/au_lname"/>
    (<xsl:value-of select="//Author/city"/>,
    <xsl:value-of select="//Author/state"/>)
  </I>
</B>
<table style="border-collapse:collapse" border="1">
<tr>
  <td class="hdr">Title</td>
  <td class="hdr">Type</td>
  <td class="hdr">Price</td>
  <td class="hdr">Notes</td>
</tr>
<xsl:for-each select="//Books">
<tr>
  <td><xsl:value-of select="title"/></td>
  <td><xsl:value-of select="type"/></td>
  <td><xsl:value-of select="price"/></td>
  <td><xsl:value-of select="notes"/></td>
</tr>
</xsl:for-each>
</table>
</body>
</html>
```

It's ok to go through the previous exercise to understand how to write a VB application as a front-end client for web services. However, for your real-life applications, you should use a SOAP toolkit (such as the Microsoft SOAP toolkit) to make things easier.

The next couple of examples take a step further. This time, we will use Perl to access our sample web service PubWS with the help of ActiveState Perl and SOAP::Lite Perl Library (author: Paul Kulchenko).[\[15\]](#)

[15] <http://www.ActiveState.com/> and <http://www.soaplite.com/>, respectively.

With SOAP::Lite Perl Library, all you have to do is to create a SOAP::Lite object (similar to the proxy object in the C# example) and setup its uri and proxy properties. Once this is done, you can ask the proxy object to run the remote web methods. The uri property maps to the first part of the soapAction and the web method name maps to the second part. The proxy property maps to the physical location of the web service itself.

In our first Perl example, we want to call the GetAuthors web method. This method does not take any parameter and returns a DataSet object. The SOAP portion of the WSDL states that the soapAction for this method is *http://Oreilly/DotNetEssentials/GetAuthors*. By default, SOAP::Lite library constructs the soapAction as *<uri string> + "#" + <web method name>*. This does not agree with web services written on the .NET Framework, where the soapAction is *uri_string + "/" + web method name*. Fortunately, the SOAP::Lite library provides a callback-like kind of feature so that we can plug in a sub-routine to override the default construction of the soapAction string. The highlighted line of PERL script below basically just concatenates the uri string and the web method name. For simplicity we rely on the fact that the uri string already has the trailing slash^[16]

[16] The trace for the SOAP::Lite library turns out to be extremely useful because it shows the SOAP request and response as well as other vital information. Insert `trace => all` at the end of the first line of the example.

To run the program, just type `perl <program name>`:

```
use SOAP::Lite
    on_action => sub {sprintf '%s%s', @_};

my $proxy = SOAP::Lite
    -> uri('http://Oreilly/DotNetEssentials/')
    -> proxy('http://localhost/PubsWS/PubsWS.asmx');

my $method = SOAP::Data->name('GetAuthors')
    ->attr({xmlns => 'http://Oreilly/DotNetEssentials/'});
my $xmlDataSet = $proxy->call($method);
# my $xmlDataSet = $proxy->GetAuthors( ); # You can also do this
my @authorsNodes = $xmlDataSet->valueof('//Authors');
foreach $author (@authorsNodes) {
    print $author->{'au_lname'}, ", ", $author->{'au_fname'}, "\n";
    print "\t", $author->{'address'}, "\n";
    print "\t", $author->{'city'}, ", ",
        $author->{'state'}, " ", $author->{'zip'}, "\n";
    print "\t", $author->{'phone'}, "\n";
}
```

The output of this program is the following:

```
White, Johnson
    10932 Bigge Rd.
    Menlo Park, CA 94025
    408 496-7223
Green, Marjorie
    309 63rd St. #411
    Oakland, CA 94618
```

415 986-7020

. . .

Once the \$proxy object is created, we call the GetAuthors web method knowing that the generated soapAction will be ``http://Oreilly/DotNetEssentials/` + `GetAuthors`` and the location of the web service where SOAP::Lite library will try to contact is `http://localhost/PubsWS/PubsWS.asmx`. Because this web method returns a DataSet, which translates to an XML document, we can parse it with XPATH-like syntax to obtain the list of authors and traverse the list to display the information for each author.

Now we want to call the GetBooksByAuthor web method as the second Perl program. This method takes one string parameter, the SSN for the author, and returns a DataSet object containing the author's name and all books he has published. The SOAP portion of the WSDL document states that the soapAction for this method is named `http://Oreilly/DotNetEssentials/GetBooksByAuthor`, and the parameter is named sAuthorSSN:

```
use SOAP::Lite
    on_action => sub { sprintf '%s%s', @_ };

my $proxy = SOAP::Lite
    -> uri('http://Oreilly/DotNetEssentials/')
    -> proxy('http://localhost/PubsWS/PubsWS.asmx');

my $method = SOAP::Data->name('GetBooksByAuthor')
    ->attr({xmlns => 'http://Oreilly/DotNetEssentials/'});

my @params = (SOAP::Data->name(sAuthorSSN => '998-72-3567'));

my $xmlDataSet = $proxy->call($method => @params);

my $author = $xmlDataSet->valueof('//Author');
print "Books by author:\n";
print $author->{'au_lname'}, ", ", $author->{'au_fname'}, "\n";
print $author->{'address'}, "\n";
print $author->{'city'}, ", ",
    $author->{'state'}, " ", $author->{'zip'}, "\n";
print $author->{'phone'}, "\n\n";

my @books = $xmlDataSet->valueof('//Books');
foreach $book (@books) {
    print "Type : ", $book->{'type'}, "\n";
    print "Title : ", $book->{'title'}, "\n";
    print "Price : \$", $book->{'price'}, "\n";
    print "Notes : ", $book->{'notes'}, "\n";
    print "\n";
}
```

Type `perl yourPerlFile.pl` to run the program. The output is:

```
Books by author:
Ringer, Albert
67 Seventh Av.
```

Salt Lake City, UT 84152
801 826-0752

Type : psychology

Title : Is Anger the Enemy?

Price : \$10.95

Notes : Carefully researched study of the effects of strong emotions on the body. Metabolic charts included.

Type : psychology

Title : Life Without Fear

Price : \$7

Notes : New exercise, meditation, and nutritional techniques that can reduce the shock of daily interactions. Popular audience. Sample menus included, exercise video available separately.

As you can see, you can easily use any type of web service client to access a .NET web service. Clients of web services need to know how to communicate in HTTP and understand the WSDL. By the same token, you can develop a web service in any language and on any platform as long as it adheres to its WSDL specification.

Now that you have seen how simple it is to create and consume web services, we should let you in on a small but important point-that XML web services is not the solution for everything. As responsible developers, we have to evaluate the requirements as well as the applicability of certain technology before applying it. On one hand, it's great to use web services as a way to expose legacy functionalities, or to enable integration of disparate systems on different locations, different platforms, or different companies. But there is no real reason to make a web service out of something that can be implemented as a simple component, or through .NET Remoting, which is discussed in [Chapter 4](#). The decision is up to you. Again, the technology is just a tool; how you use it is another story.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

6.5 Async Web Services

All of the web service client examples we've used so far make a call to the web service and wait for data before continuing work. While this approach might be appropriate for small, quick calls, such synchronous calls might yield an unacceptable response time. The web service has to perform a time consuming task. Asynchronous web service calls are the answer in this case. This section explores how the .NET Framework web services enable clients to call web methods asynchronously, yielding a better overall response time for an application.

If you look back at the generated source for the web service proxy in the previous section, you will find that for every web method there are three separate proxy methods generated. The following shows the three proxy methods: `GetAuthors()`, `BeginGetAuthors()`, and `EndGetAuthors()`:

```
public System.Data.DataSet GetAuthors( ) {
    object[] results = this.Invoke("GetAuthors", new object[0]);
    return ((System.Data.DataSet)(results[0]));
}

public System.IAsyncResult BeginGetAuthors(System.AsyncCallback callback,
    object asyncState) {
    return this.BeginInvoke("GetAuthors",new object[0],callback,asyncState);
}

public System.Data.DataSet EndGetAuthors(System.IAsyncResult asyncResult) {
    object[] results = this.EndInvoke(asyncResult);
    return ((System.Data.DataSet)(results[0]));
}
```

In the earlier web service consumer example, we've used the proxy method `GetAuthors()` to call the web method on the other side of the HTTP channel. The following example shows how the other two methods can be used to perform an asynchronous web method call.

As it turns out, there are a couple of different ways to use the `BeginMethodName()` and `EndMethodName()` to call the web method asynchronously. We will show you each one and describe how they are different from one another.

The first way will be calling `BeginMethodName()` passing in `null` for both the callback and the `asyncState` objects. We will poll on the `IAsyncResult` object until it is completed while simulating other works. When the call is done, we call `EndMethodName()` to complete the web method call and obtain the `DataSet` result:^[17]

^[17] The proxy in this example was generated with a namespace "WSPubsWS" instead of the default global namespace.

```
WSPubsWS.PubsWS oProxy = new WSPubsWS.PubsWS( );
IAsyncResult result;
// Async call polling (cannot run this in debugger)
Console.Write("Async Polling");
```

```

result = oProxy.BeginGetAuthors(null, null);
while(!result.IsCompleted) {
    Console.WriteLine(" . . . "); Thread.Sleep(5);
}
DataSet oDSAsync = oProxy.EndGetAuthors(result);
oDSAsync.WriteXml("outputpolling.xml");
Console.WriteLine("done");

```

Although this might not be the best way, it demonstrates one way of performing an asynchronous call to a web method.

The second way still does not use the callback mechanism. Be patient. Again, we use the `IAsyncResult` object to perform a block waiting on the remote call. Block waiting put the current thread to sleep, giving up the CPU to other processes or threads until the wait handle is signaled. Depending on your application, this might be how you would want to do it:

```

// Async call with waithandle
Console.WriteLine("Async, processing then wait for handle");
result = oProxy.BeginGetAuthors(null, null);
Console.WriteLine(" . . . Processing some more . . . ");
result.AsyncWaitHandle.WaitOne( ); // block on handle
DataSet oDSAsyncWait = oProxy.EndGetAuthors(result);
oDSAsyncWait.WriteXml("outputWaiting.xml");
Console.WriteLine("done");

```

The third and fourth ways use the callback mechanism so we will know exactly when to complete the web method call with the `EndMethodName()`. We make the call to the `BeginMethodName()` passing in the callback delegate. When the web method call completes, we will get called back. Since this example is done as a console application, a little more setting up is required; otherwise, the application might exit before the callback occurs. This is done through the `AutoResetEvent` object (a synchronization primitive similar to a mutex). Basically, we want to set up a wait object, call the async web method, and wait for the signal from the callback before continuing with the application:

```

// Async call with callback 1
AutoResetEvent oWait = new AutoResetEvent(false);
Console.WriteLine("Async, processing then wait for callback");
Callback cb = new Callback(oProxy, oWait);
Result =oProxy.BeginGetAuthors(new AsyncCallback(cb.CallbackMethod), null);
Console.WriteLine(" . . . Processing some more . . . ");
Console.WriteLine("Application waits for callback rendezvous");
oWait.WaitOne( );
Console.WriteLine("done");

```

The `Callback` class needs to have the proxy to complete the web method call. It also needs the `AutoResetEvent` to signal the main application flow to continue. The following is the definition of the `Callback` class at this moment:

```

public class Callback {
    public Callback(WSPubsWS.PubsWS oProxy, AutoResetEvent oWait) {
        m_oProxy = oProxy;
        m_oWait = oWait;
    }
    public void CallbackMethod(IAsyncResult result) {

```

```

        DataSet oDSAsyncCB;
        oDSAsyncCB = m_oProxy.EndGetAuthors(result);
        oDSAsyncCB.WriteXml("outputCallback.xml");
        m_oWait.Set( );
    }
    private WSPubsWS.PubsWS m_oProxy;
    private AutoResetEvent m_oWait;
}

```

In the next example, we pass both the callback and the asyncState object to the `BeginMethodName()` method. We already know all about the callback object. What about this asyncState thing? It actually can be *anything* you want to pass to the callback object. We will pass the proxy object itself via this asyncState object so we will know how to complete the web method call when the callback is invoked. The following code segment shows how this is done:

```

// async call with callback 2
oWait.Reset( );
Console.WriteLine("Async, processing then wait for callback also");
CallBack cbWithData = new CallBack(oWait);
result=oProxy.BeginGetAuthors(new AsyncCallback(cbWithData.CallBackMethod), oProxy);
Console.WriteLine("Processing some more data");
Console.WriteLine("Application waits for callback rendezvous");
oWait.WaitOne( );
Console.WriteLine("done");

```

This time, the CallBack object is instantiated with only one parameter. Changes to the CallBack class are highlighted in the following code block. When the callback method is invoked, you can cast the IAsyncResult object's AsyncState back to whatever type you passed into the second parameter of the `BeginMethodName()` method. In this case, we cast the AsyncState object back to the proxy in order to complete the web method call with `EndMethodName()` to obtain the resulting DataSet:

```

public class CallBack {
    public CallBack(WSPubsWS.PubsWS oProxy, AutoResetEvent oWait) {
        m_oProxy = oProxy;
        m_oWait = oWait;
    }
    public CallBack(AutoResetEvent oWait) {
        m_oWait = oWait;
    }
    public void CallBackMethod(IAsyncResult result) {
        DataSet oDSAsyncCB;
        if(m_oProxy == null) {
            WSPubsWS.PubsWS oTmp = (WSPubsWS.PubsWS)result.AsyncState;
            oDSAsyncCB = oTmp.EndGetAuthors(result);
            oDSAsyncCB.WriteXml("outputCallbackWithData.xml");
        } else {
            oDSAsyncCB = m_oProxy.EndGetAuthors(result);
            oDSAsyncCB.WriteXml("outputCallback.xml");
        }
        m_oWait.Set( );
    }
    private WSPubsWS.PubsWS m_oProxy;
}

```



```
private AutoResetEvent m_oWait;  
}
```

In a Windows Form application, you can start the web service call via a button clicked or a menu clicked and still can have other parts of the application be responsive to the user. We will leave that example to you as an exercise.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

6.6 SOAP Header in Web Services

As you have seen from the web service consumer using SOAP, the data being passed back and forth as is structured XML inside the body of the HTTP package. In particular it has the following format:

```

<Envelope>
  <Header>
    < . . . >
  </Header>
  <Body>
    < . . . >
  </Body>
</Envelope>

```

We have seen how to pass data from the client to the web service through the Body of the Envelope. In this section, we show how to use the Header of the Envelope.

Through this optional header node, the developers can pass information that does not relate to any particular web method or information that is common to all web methods. Of course, you can pass information that does not related to any particular web method in a method call itself, such as `InitiateService(param)` enforcing a usage rule that this method has to be called first. This is not an elegant solution. Just add the `param` of the `InitiateService` to the header. On the other hand, if all web methods of your service require a common parameter, wouldn't it be nice if you could just set up the header and didn't have to worry about this parameter for each of the web methods. Examples of header information are a security token, a payment token, priority, and so on; there is no reason to pass these pieces of common information via every web method.

Once you construct your web service to use a SOAP header, the WSDL will instruct the client that a header node is in place so that a web service client knows how to set up the header information before making a call to the web service methods. The following example should provide some clarification:

```

<%@ WebService Language="C#" Class="TestHeader.PayWS" %>

namespace TestHeader {
    using System;
    using System.Web;
    using System.Web.Services;
    using System.Web.Services.Protocols;

    public class Payment : SoapHeader {
        public string CreditCardNumber;
    }

    [WebService(Namespace="http://Oreilly/DotNetEssentials/")]
    public class PayWS : WebService {

```

```

public Payment clientPayment;

[WebMethod, SoapHeader("clientPayment")]
public string Task1( ) {
    return string.Format("Task1 performed. " +
        "Charging $25,000.00 on credit card: {0}",
        clientPayment.CreditCardNumber);
}

[WebMethod, SoapHeader("clientPayment")]
public string Task2( ) {
    return string.Format("Task2 performed. " +
        "Charging $4,500.00 on credit card: {0}",
        clientPayment.CreditCardNumber);
}
}
}

```

In this example, we create a web service with two methods, Task1 and Task2. Both of these methods use the Payment object in the soap header.

The SOAP request and response format for Task1 follows:

```

POST /SOAPHeader/PayWS.asmx HTTP/1.1
Host: localhost
Content-Type: text/xml; charset=utf-8
Content-Length: length
SOAPAction: "http://Oreilly/DotNetEssentials/Task1"

<?xml version="1.0" encoding="utf-8"?>
<soap:Envelope
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xmlns:xsd="http://www.w3.org/2001/XMLSchema"
    xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
  <soap:Header>
    <Payment xmlns="http://Oreilly/DotNetEssentials/">
      <CreditCardNumber>string</CreditCardNumber>
    </Payment>
  </soap:Header>
  <soap:Body>
    <Task1 xmlns="http://Oreilly/DotNetEssentials/" />
  </soap:Body>
</soap:Envelope>

```

```

HTTP/1.1 200 OK
Content-Type: text/xml; charset=utf-8
Content-Length: length

```

```

<?xml version="1.0" encoding="utf-8"?>
<soap:Envelope
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xmlns:xsd="http://www.w3.org/2001/XMLSchema"

```

```

    xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
<soap:Body>
  <Task1Response xmlns="http://Oreilly/DotNetEssentials/">
    <Task1Result>string</Task1Result>
  </Task1Response>
</soap:Body>
</soap:Envelope>

```

This is different than the soap envelope we've seen for PubsWS where there was no SOAP header. In this case, the header is an object of type Payment and this Payment class has a string property called CreditCardNumber.

Our example expects the clients to this web service to pass in the Payment object. The following is the client code. Of course, you will have to generate the proxy class for the PayWS web service using *wsdl.exe* and compile this proxy along with the client code:

```

public class ClientWS {
    public static void Main( ) {
        // Create a proxy
        PayWS oProxy = new PayWS( );

        // Create the payment header.
        Payment pmt = new Payment( );
        pmt.CreditCardNumber = "1234567890123456";

        // Attach the payment header to the proxy.
        oProxy.PaymentValue = pmt;

        // Call Task1.
        System.Console.WriteLine(oProxy.Task1( ));

        // Call Task2.
        System.Console.WriteLine(oProxy.Task2( ));
    }
}

```

The output of the client:

```

Task1 performed.   Charging $25,000.00 on credit card: 1234567890123456
Task2 performed.   Charging $4,500.00 on credit card: 1234567890123456

```

The above is a trivial usage of SOAP header. In reality, the SOAP headers are probably not as simple as our Payment class.

Microsoft, IBM, and a number of other companies have been working on the Global XML web services Architecture (GXA) that defines a framework on how to build standardized web services incorporating security, reliability, referral, routing, transaction, and so on. A number of GXA specifications use SOAP Header as the means for this infrastructure.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

6.7 Web Services and Security

This section demonstrates how to incorporate security into your web service. We will do so in two ways: system security and application security. System-level security allows for restricting access to the web services from unauthorized clients. It is done in a declarative fashion, whereas application-level security is more flexible. With system-level security, you will most likely have the list of authorized clients' IP addresses that you will let access your web service through the use of some configuration-management tools. With application-level security, you will incorporate the authentication into your web service, thus providing a more flexible configuration.^[18]

^[18] As always, consult your security expert on this decision.

6.7.1 System Security

Because web services communication is done through HTTP, you can apply system-level security on web services just as you do for other web pages or resources on your web site.

There are a number of different ways you can secure your web services. For a business-to-business (B2B) solution, you can use the IIS Administration Tool to restrict or grant permission to a set of IP addresses, using the Internet Protocol Security (IPSec) to make sure that the IP address in the TCP/IP header is authenticated. When you rely only on the client to provide the IP in the TCP/IP header, hackers can still impersonate other host IPs when accessing your web services. IPSec authenticates the host addresses using the Kerberos authentication protocol. You can also use a firewall to restrict access to your web services for your partner companies. For a business-to-consumer (B2C) scenario, you can take advantage of the authentication features of the HTTP protocol.

To show you how to use the authentication feature of the HTTP protocol to secure your web services, let's revisit the example web service we have in this chapter, PubsWS. All we have to do to secure PubsWS web service is go to the IIS Admin Tool and choose to edit the **File Security** properties for the *PubsWS.asmx*. Instead of keeping the default setting, which leaves this file accessible to all anonymous users, we change this setting to "Basic Authentication" only, which means unchecking "Anonymous Access" and checking only "Basic Authentication" in the Authenticated Access frame. After this change, only users that pass the authentication can make use of the web service.

For real-life situations, of course, we are not going to use just the Basic Authentication method, because it sends the username and password in clear text through the HTTP channel. We would choose other methods, such as Secure Sockets Layer (SSL) underneath Basic Authentication, so that the data passed back and forth is secure. Available methods include:

Basic Authentication

Sends the username and password to the Web Server in clear text. IIS authenticates the login against the database of users for the domain.

Basic over SSL Authentication

Similar to Basic Authentication, except that the username and password are sent with SSL encryption.

Digest Authentication

Uses a hashing technique, as opposed to SSL encryption, to send client credentials securely to the server.

Integrated Windows Authentication

Good for intranet scenarios only. Uses the login information of the client for authentication.

Client Certificates Authentication

Requires each of the clients to obtain a certificate that is mapped to a user account. The use of client-side digital certificates is not widespread at this time.

6.7.2 Application Security

A less systematic way of securing your web services involves taking security into your own hands. You can program your web services so that all of their methods require an access token, which can be obtained from the web service after sending in the client's username and password. The client credentials can be sent to the server through SSL, which eliminates the risk of sending clear-text passwords across the wire. Through this SSL channel, the server returns an access token to the caller, who can use it to invoke all other web service methods. Of course, all of the other web methods that you publish must have one parameter as the token. A simple pseudocode example of a bank account web service can be the following:

web service Bank Account

Web Methods:

Login(user id, password) returns **access token** or nothing

Deposit(**access token**, account number, amount, balance) returns boolean

Withdraw(**access token**, account number, amount, balance) returns boolean

The only method that should be on SSL is the Login method. Once the token is obtained, it can be used for other web methods. Of course, you should be able to make sure that subsequent calls using this token are coming from the same IP as the Login() call. You can also incorporate an expiration timestamp on this access token to ensure that the token only exists in a certain time frame until a renewal of the access token is needed.

You can also use public/private keys (asymmetric) encryption for better key- management. The following scenario might suit your needs.

The client gets the server's public key and uses it to encrypt the requesting data (possibly including the client's private symmetric encryption key) before calling the web service. This ensures that the requesting data is encrypted and only the server can decrypt it. The server decrypts the data using the private key, figures out what the request is and uses the client's private key to encrypt the response before sending it back to the client. This time, the response data is encrypted. The client then decrypts the response package with its private symmetric key to view clear data. Since asymmetric cryptography operations are always much slower than symmetric, it's probably best if you only use asymmetric cryptography for key distribution. Once the key is communicated, you can use the symmetric operation for better performance. Rolling your own security is always harder than using a standard security solution. There are numerous areas you will have to manage yourself such as identity authentication, message authentication, and so on. Be forewarned if you choose this

route.

The Microsoft .NET Cryptographic Services can be very useful if you choose to apply application security for your web services. DES, RC2, RC4, TripleDES, and RSA cryptography algorithms are supported along with hashing methods such as SHA and MD5. These implementations in the .NET library enable developers to avoid low-level grunt work and focus on the application logic.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

6.8 Summary

In this chapter, we've introduced you to the new paradigm of applications-the enterprise application. You are no longer restricted to homogeneous platforms for implementing your solutions. With Microsoft web services, your solutions can span many different platforms because the communication between web services is done through standard Internet protocols such as HTTP and XML. The distributed components in Windows DNA with which you may be familiar are now replaced by web services. Using web services as components in a distributed environment allows for a heterogeneous system. The web services in your system can not only be implemented in different languages, but can even be on different platforms. Because of this greater interoperability, web services are eminently suitable for B2B integration.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Chapter 7. ASP.NET

This chapter introduces the next technology for providing dynamic and interactive web pages. ASP.NET takes Active Server Pages (ASP) a step further to simplify the development process of web pages and enhance scalability of web applications. First, we review the conventional way of developing web applications using ASP to uncover some of the pitfalls that ASP.NET overcomes. We then discuss the benefits of ASP.NET and provide a high-level survey of the classes in the ASP.NET, such as control and page, as well as the complete syntax of ASP.NET Web Forms. To wrap up the chapter, we discuss the many aspects of ASP.NET development: how ASP.NET supports web services development, how to use custom server controls, and how session management has been improved to provide more scalable web solutions. This overview chapter is designed to provide experienced developers with a solid introduction to ASP.NET and Web Forms; for additional in-depth information, see *Programming ASP.NET* by Liberty and Hurwitz (O'Reilly) and *ASP.NET in a Nutshell* by Duthie and MacDonald (O'Reilly).

[\[Team LiB \]](#)

[\[Team LiB \]](#)

7.1 ASP

Microsoft ASP is a server-side scripting technology enabling dynamic web pages. An ASP page contains HTML markup and server-side scripts that generate HTML content dynamically. The server-side scripts run when a request for the ASP page arrives at the web server. Inputs to the ASP page come from the client browsers through HTTP POST and GET methods. ASP provides an object model to simplify developers' tasks. Besides objects from the ASP object model like Application, Server, Request, Response, and Session, developers can use any COM components on the server.

If you've already been developing web applications using ASP, you probably agree that it is very easy to end up with intertwined, possibly conflicting HTML markups and server-side scripts. The poor encapsulation model of ASP pages makes them difficult to manage and reuse. Attempts have been made to improve upon this model, including server-side include files and parameterized functions in scripts; however, these attempts come with trade-offs such as time, the management of a network of include files, the performance impact of having nested includes, as well as object ID and variable-scope management.

Developers that deal with cross-browser web applications also run into problems generating HTML according to the client's browser capability. Most of the time, we end up generating only the simplest HTML tags and client-side scripts, which can be understood by many browsers, and foregoing the features of the more advanced browsers. The resulting web application can be only as good as the worst browser it supports. Sometimes, we also attempt to generate different HTML markups for different browsers to take advantage of browser-specific features, resulting in a much better client-side experience; however, this involves much more development time and effort.

Since scripting in ASP is available only to late-bound languages such as VBScript and JavaScript, type-safety is not an option. In addition, server-side scripts in ASP pages get reinterpreted each time the page is accessed, which is not ideal for performance.

Form-state maintenance in an ASP-based application is also labor-intensive-developers must do everything manually, including reposting data, using hidden fields, and session variables. At times, web applications are configured to run in web farm environments where there is more than one web server available to the client. Maintaining session states becomes much harder in these scenarios because it is not guaranteed that the client would return to the same server for the next request. Basically, the developers have to save states manually to SQL Server or other external storage.

Although ASP is a great technology to build dynamic web pages, it has room for improvement. ASP.NET evolved from ASP and overcomes most, if not all, of its shortfalls.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

7.2 ASP.NET

Visual Basic developers have long enjoyed the ease of programming with forms and controls. Writing a VB form-based application is as simple as dragging some controls onto a form and writing some event-handling functions. This is one of the reasons VB has attracted lots of programmers interested in speed of development. Microsoft wisely built this feature into ASP.NET.

ASP.NET simplifies web page development with form-based programming. In ASP.NET, these forms are called *Web Forms* and are analogous to VB forms, replacing ASP pages. Similar to VB, Web Forms programming is also event based. We don't have to write in-line ASP scripts and rely on the top-down parsing interpretation as in ASP programming. To match the rich set of ActiveX controls that VB programmers love in their toolset, ASP.NET equips ASP programmers with *server controls*. To further enhance the productivity of developers, ASP.NET's Web Forms also allow for the separation of the application logic and the presentation layer.

ASP.NET evolves from the ASP programming model with the following additional benefits:

- Clean separation between the application logic (server-side code) and the presentation layer (HTML markup)-no more spaghetti code
- A rich set of server controls that automatically render HTML suitable for any clients and that additionally manage their states
- Enhanced session-state management
- An event-based programming model on the server side, which is simpler and more intuitive than an older script-based model like ASP
- Application logic that can be written in any Microsoft .NET language (VB,C#, Managed C++, etc.); application server-side code is compiled for better performance
- Visual Studio .NET as a RAD tool, which simplifies the development process of Web Forms

[\[Team LiB \]](#)

[\[Team LiB \]](#)

7.3 The System.Web.UI Namespace

Before getting into developing your favorite *Hello, World!* application in ASP.NET, it's important that you become familiar with the underlying structure of ASP.NET. This section describes some of the most important classes packaged in the System.Web.UI namespace in the ASP.NET framework.

The System.Web.UI namespace defines classes and interfaces used in constructing and rendering elements on a Web Form. The most important class in the System.Web.UI is the Control class, which defines properties, methods, and events that are common in all server controls in the Web Forms framework. Another important class in this namespace is Page, which is a derivative of the Control class. All ASP.NET web pages are instances of derivatives of the Page class. To have an extensible framework, the System.Web.UI namespace also includes the UserControl class, which is similar to the Page class except that it is used as the base class for user controls. We will make use of the UserControl and Page classes in [Section 7.5](#) and [Section 7.5.4](#) sections later in this chapter.

7.3.1 Control Class

The Control class is the root of all controls. For example, a text box is a control; a button or a combo box is also a control. The Control class basically encapsulates common functionalities and properties of all user-interface widgets. As you get deeper into ASP.NET development, everything you see is a Control derivative of some sort.

7.3.1.1 Control's properties

The Control class has the following important properties: Controls, ID, ClientID, Parent, EnableViewState, Visible, Context, and ViewState. We will go over each of these properties briefly to show you what the Control class is made up of and how deriving from Control class would create a model that is consistent and easy to work with.

The Controls property represents the children of the control instance; the Parent property defines the parent of the control. These properties enable a hierarchy of controls on a web page. The ID property allows the control to be accessed programmatically by just using the ID and the dot notation to get to the object's properties and methods (i.e., MyObjectId.propertyname). While the ID property allows us to program the control on the server side, ClientID allows us to setup client-side script to access the control on the client side. More information on using ClientID will be shown later in this chapter.

The EnableViewState flag indicates whether the control will maintain its view state, as well as all view states of its child controls. If this flag is set to true, the control will remember its previous view state when the page posts back to itself.^[1] For example, if EnableViewState is set to true, the user's previous selection or form-field data are preserved automatically when the user performs some operation that requires a postback. When the page is sent back to the browser, the user can just continue filling in the form as if he never left it. This is how all derivatives of the Control class maintain their states between requests and free ASP developers from having to simulate view-state behavior with hidden form fields.

[1] Postback is the condition when an ASP page posts the data back to itself for processing. In conventional ASP programming, the states of the fields in the form have to be managed manually. In ASP.NET, we can have these field states managed automatically with a simple `EnableViewState` flag.

The `Context` property enables us to get to information about the current HTTP request, such as the `Application`, `Server`, `Session`, `Request`, and `Response` objects. ASP developers should be familiar with these intrinsic objects. You will likely use the `Context` property when you are processing the web page's `Load` event to get to application- or session-level variables and request parameters to set up your page. Through the `Context` property, you can also get other information, such as cached resources, including database connection for performance improvement; the `trace` property for debugging purposes; and the `user` property for security validation.

The `ViewState` property is an instance of the `StateBag` class, which is used to store name/value pairs of information that can be made accessible across multiple requests for the same web page. These name/value pairs are instances of the `StateItem` class. `ViewState` allows ASP.NET controls to maintain their own state across multiple client roundtrips; it is implemented as a hidden form field on the page. If you've attempted to maintain state for your form in ASP development, you will appreciate this because it is now automatically done for you.

7.3.1.2 Control class methods

The list of methods for the `Control` class is much longer than what we've covered in this section; however, this short list is probably all you need to know to get started with the `Control` class:

DataBind method

Binds the control to a data source. This method is used in conjunction with the data-binding expression syntax on the Web Form. When this method is called, all data-binding tags, `<%# %>`, are re-evaluated so that the new data is bound to the appropriate tag location. Also, any controls that have their `DataSource` property set, retrieve the data from the `DataSource` and fill themselves.

CreateChildControls method

Called before any compositional custom control is rendered. A compositional custom control is similar to a user control. Both of them are composed of other controls to create more complex controls. You would not employ this method simply to use the control. When developing custom controls, this method can be overridden so that custom control developers can create and layout child controls prior to rendering the controls, whether for the first time or for postbacks.

Render method

Similar to the `CreateChildControls`, primarily used to develop custom controls. Control developers override this method to render the control content through the provided `HtmlTextWriter` parameter.

We will revisit the `Render` and `CreateChildControls` methods when we show you how to create custom controls in "Customer Server Controls" later in this chapter.

SaveViewState and LoadViewState methods

Save and reload the state for the control. Server controls maintain their state between requests via these methods.

7.3.2 Page Class

As mentioned earlier, the Page class is actually a derivative^[2] of the Control class. This means it inherits all properties, methods, and events exposed by the Control class. In addition to the inherited things, the Page class defines more specific properties, methods, and events for a web page in the ASP.NET framework.

[2] The Page class derives from TemplateControl, which derives from the Control class.

If you've done ASP development, you already know that Application, Request, Response, Server, and Session are intrinsic objects that you can access while scripting your ASP page. With ASP.NET, these objects are actually properties of the Page class. In addition to these familiar objects, the Page class also exposes other properties such as Cache, ErrorPage, IsPostBack, IsValid, Trace, and Validators.

7.3.2.1 Page class properties and methods

This list is not complete; however, it includes some of the more important features that we want to introduce:

Cache property

Points to a Cache object of the Context for the current page. Here, resources such as DataSet with information retrieved from a database are stored for reuse while the cache item is not yet expired.

ErrorPage property

Specifies the page to display when an error occurs. You can also specify the error page by using the @Page directive, as shown in [Section 7.4](#).

IsPostBack property

Indicates whether the page request is an original request or a postback, since the interaction between the user and the server controls requires a postback to the current page. If IsPostBack is `true`, you should not redo all your page initialization to improve performance.

Validators property

Groups together server controls that can validate themselves inside the Validators property of the Page. (In ASP.NET, a web page usually consists of a number of server controls.) This is so that when the Page needs to validate itself, it can delegate the validation to all of these control and then set the IsValid property to the appropriate value.

Trace property

References a TraceContext object, through which you can issue warning or error messages. Tracing can be switched on or off at any time from the `web.config` setting. `web.config` is an XML-based text file that stores the runtime configuration for an ASP.NET application. Changes to this file take effect immediately. The main configuration file is at the root of your web application; however, you can have a configuration file for each subdirectory in your web application. The closest configuration file overrides the settings of distant configuration files. Being able to switch off tracing in a configuration file like this is much better than doing so manually in ASP development, where you must go through all ASP files to remove all instances of Response. Write debugging messages when you are ready to deploy your application.

LoadControl method

Loads server controls from a file into the page programmatically. This is especially for user controls in ascx files. For ordinary server controls, they can be instantiated directly and added to the page's Controls collection. You can also have static server control declared on the page using the server-side object syntax as described in [Section 7.4](#) later in this chapter.

MapPath method

Maps a virtual path to a physical path for file I/O. This should be familiar to ASP developers.

Validate method

Works with the Server Validation Controls on the page to validate data on the page. If any of the server controls fail to validate, this method returns `false`, and the failed server-validation control renders the error message to the user.

CreateHtmlTextWriter method

Produces an `HtmlTextWriter` object to write HTML to the response stream. This is similar to ASP's `Response.Write` method; however, the `HtmlTextWriter` object is much smarter than the raw `Write` method. It helps you write well-formed HTML.

LoadPageStateFromPersistenceMedium, SavePageStateToPersistenceMedium methods

By default, save and load view state for all controls as hidden fields on the page. If you don't want this setting, you can override the `SavePageStateFromPersistenceMedium` method to save the view state anywhere other than hidden fields. You will also have to override the `LoadPageStateFromPersistenceMedium` method to have the saved view states loaded back onto the page prior to rendering.

7.3.3 UserControl Class

The `UserControl` class is similar to the `Page` class (see the previous section) with the omission of page-specific properties or methods such as `ErrorMessage`, `IsValid`, `User`, `Validators`, `MapPath`, `Validate`, and `CreateHtmlTextWriter`.

The `UserControl` class is typically used as the base class for custom controls. We can also build custom controls by inheriting directly from the `Control` class; however, it's better to start from `UserControl` because it is not as raw as the `Control` class. If you find that `UserControl` supports a number of properties and methods that you don't really want in your custom control, you might choose to inherit the raw `Control` class instead. We show you how to create custom controls in [Section 7.5.4](#) later in this chapter.

7.3.4 System.Web.UI.HtmlControls Namespace

If you've done any client-side DHTML scripting, you know how all HTML tags are mapped to scriptable objects. ASP.NET brings this mapping to the server side. Before the web page is rendered and sent back the client, you can access and manipulate each of the objects on the page.

ASP.NET maps HTML tags with objects in the hierarchy of server-side classes defined in the `System.Web.UI.HtmlControls` namespace. These server objects are called *HtmlControls* because they closely map to standard HTML elements.

For example, here is a simple HTML page that relies on client-side scripting to change the output

page dynamically. (This page won't run on browsers that do not support VBScript client-side scripting or browsers that have client-side scripting turned off.)

```
<html>
<head>
  <script language=vbscript>
    sub cmd1_onclick( )
      txtMessage.InnerHtml = _
        "(Client-side) Your name is: " & frm1.txtName.value
    end sub
  </script>
</head>
<body>
  <form id=frm1>
    Enter Name: <input id="txtName" type="text" size="40">
    <input type=button id="cmd1" value="Click Me">
    <span id="txtMessage"></span>
  </form>
</body>
</html>
```

We will convert this page so that it relies on server control instead of the IE Document Object Model. Since the output of the page is controlled from the server side, the page works regardless of what kind of browser you are using. One drawback to this is that all interaction with the page requires a postback to the server.

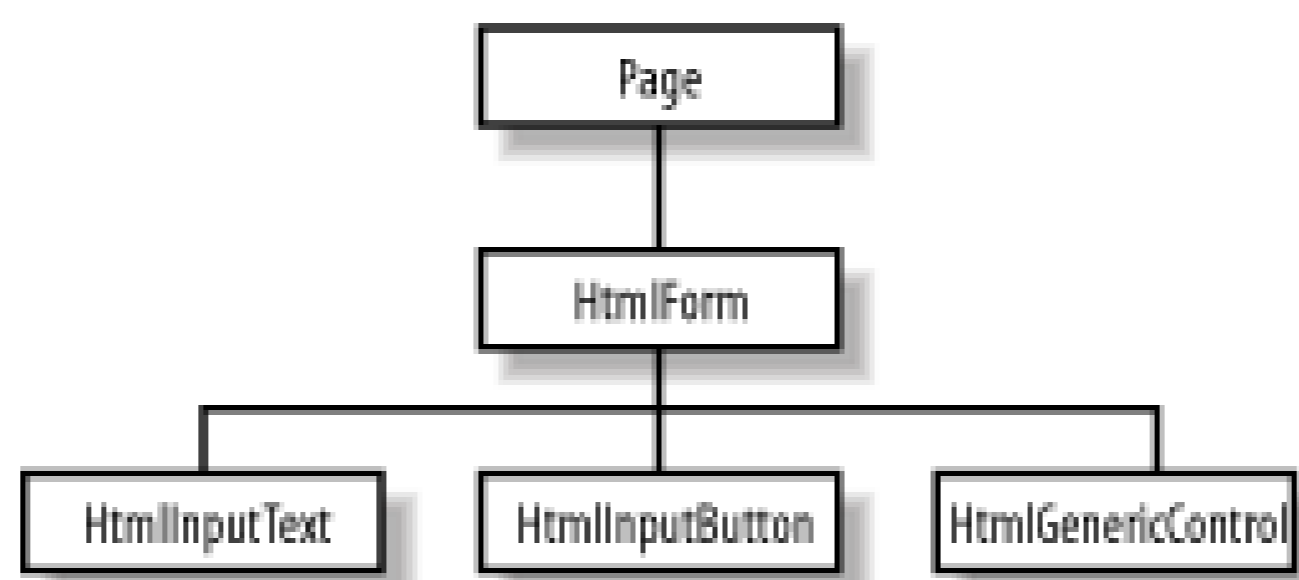
To take advantage of server controls mapping, all you have to do is to add the `id` and `runat` attributes, and your server-side script will be able to access and manipulate the server controls:

```
<html>
<head>
  <script id="scr1" language="c#" runat="server">
    void svr_cmd1_onclick(Object o, EventArgs e)
    {
      txtMessage.InnerHtml =
        "(Server-side) Your name is: " + txtName.Value;
    }
  </script>
</head>
<body>
  <form id="frm1" runat="server">
    Enter Name: <input id="txtName" type="text" size="40" runat="server">
    <input type="button" id="cmd1" value="Click Me"
      onserverclick="svr_cmd1_onclick " runat="server">
    <span id="txtMessage" runat="server"></span>
  </form>
</body>
</html>
```

By adding the `runat="server"` attribute to the HTML form and its controls, you have exposed an `HtmlForm` object, an `HtmlInputText` object, an `HtmlInputButton` object, and an `HtmlGenericControl` object (the span) to your server-side script, as depicted in [Figure 7-1](#). As you can see in the previous

script, you can manipulate the `HtmlGenericControl` object's `txtMessage` to set its `InnerHTML` property.

Figure 7-1. Server-side scriptable objects for the code example



Even though the results of the two simple examples appear to be the same, they are drastically different from the technical point of view. Client-side scripting, as the name implies, runs in the client browser process. On the other hand, when we have controls tagged to run on the server, we can have access to other server resources.

Most classes in the `System.Web.UI.HtmlControls` namespace are derivatives of the `HtmlControl` class which in turn derives from the `Control` class of the `System.Web.UI` namespace. See [Figure 7-2](#) for a graphical presentation of the hierarchy. The `HtmlControl` class serves as the base class for these `HtmlControls` because most HTML elements share common characteristics that are defined in this `HtmlControl` base class. They share properties such as `ID`, `Disabled`, `Value`, `Style`, and `TagName`. Because these `HtmlControls` ultimately are derivatives of the `Control` class, they also have methods and events that the `Control` class exposes.

Figure 7-2. `HtmlControls` object hierarchy

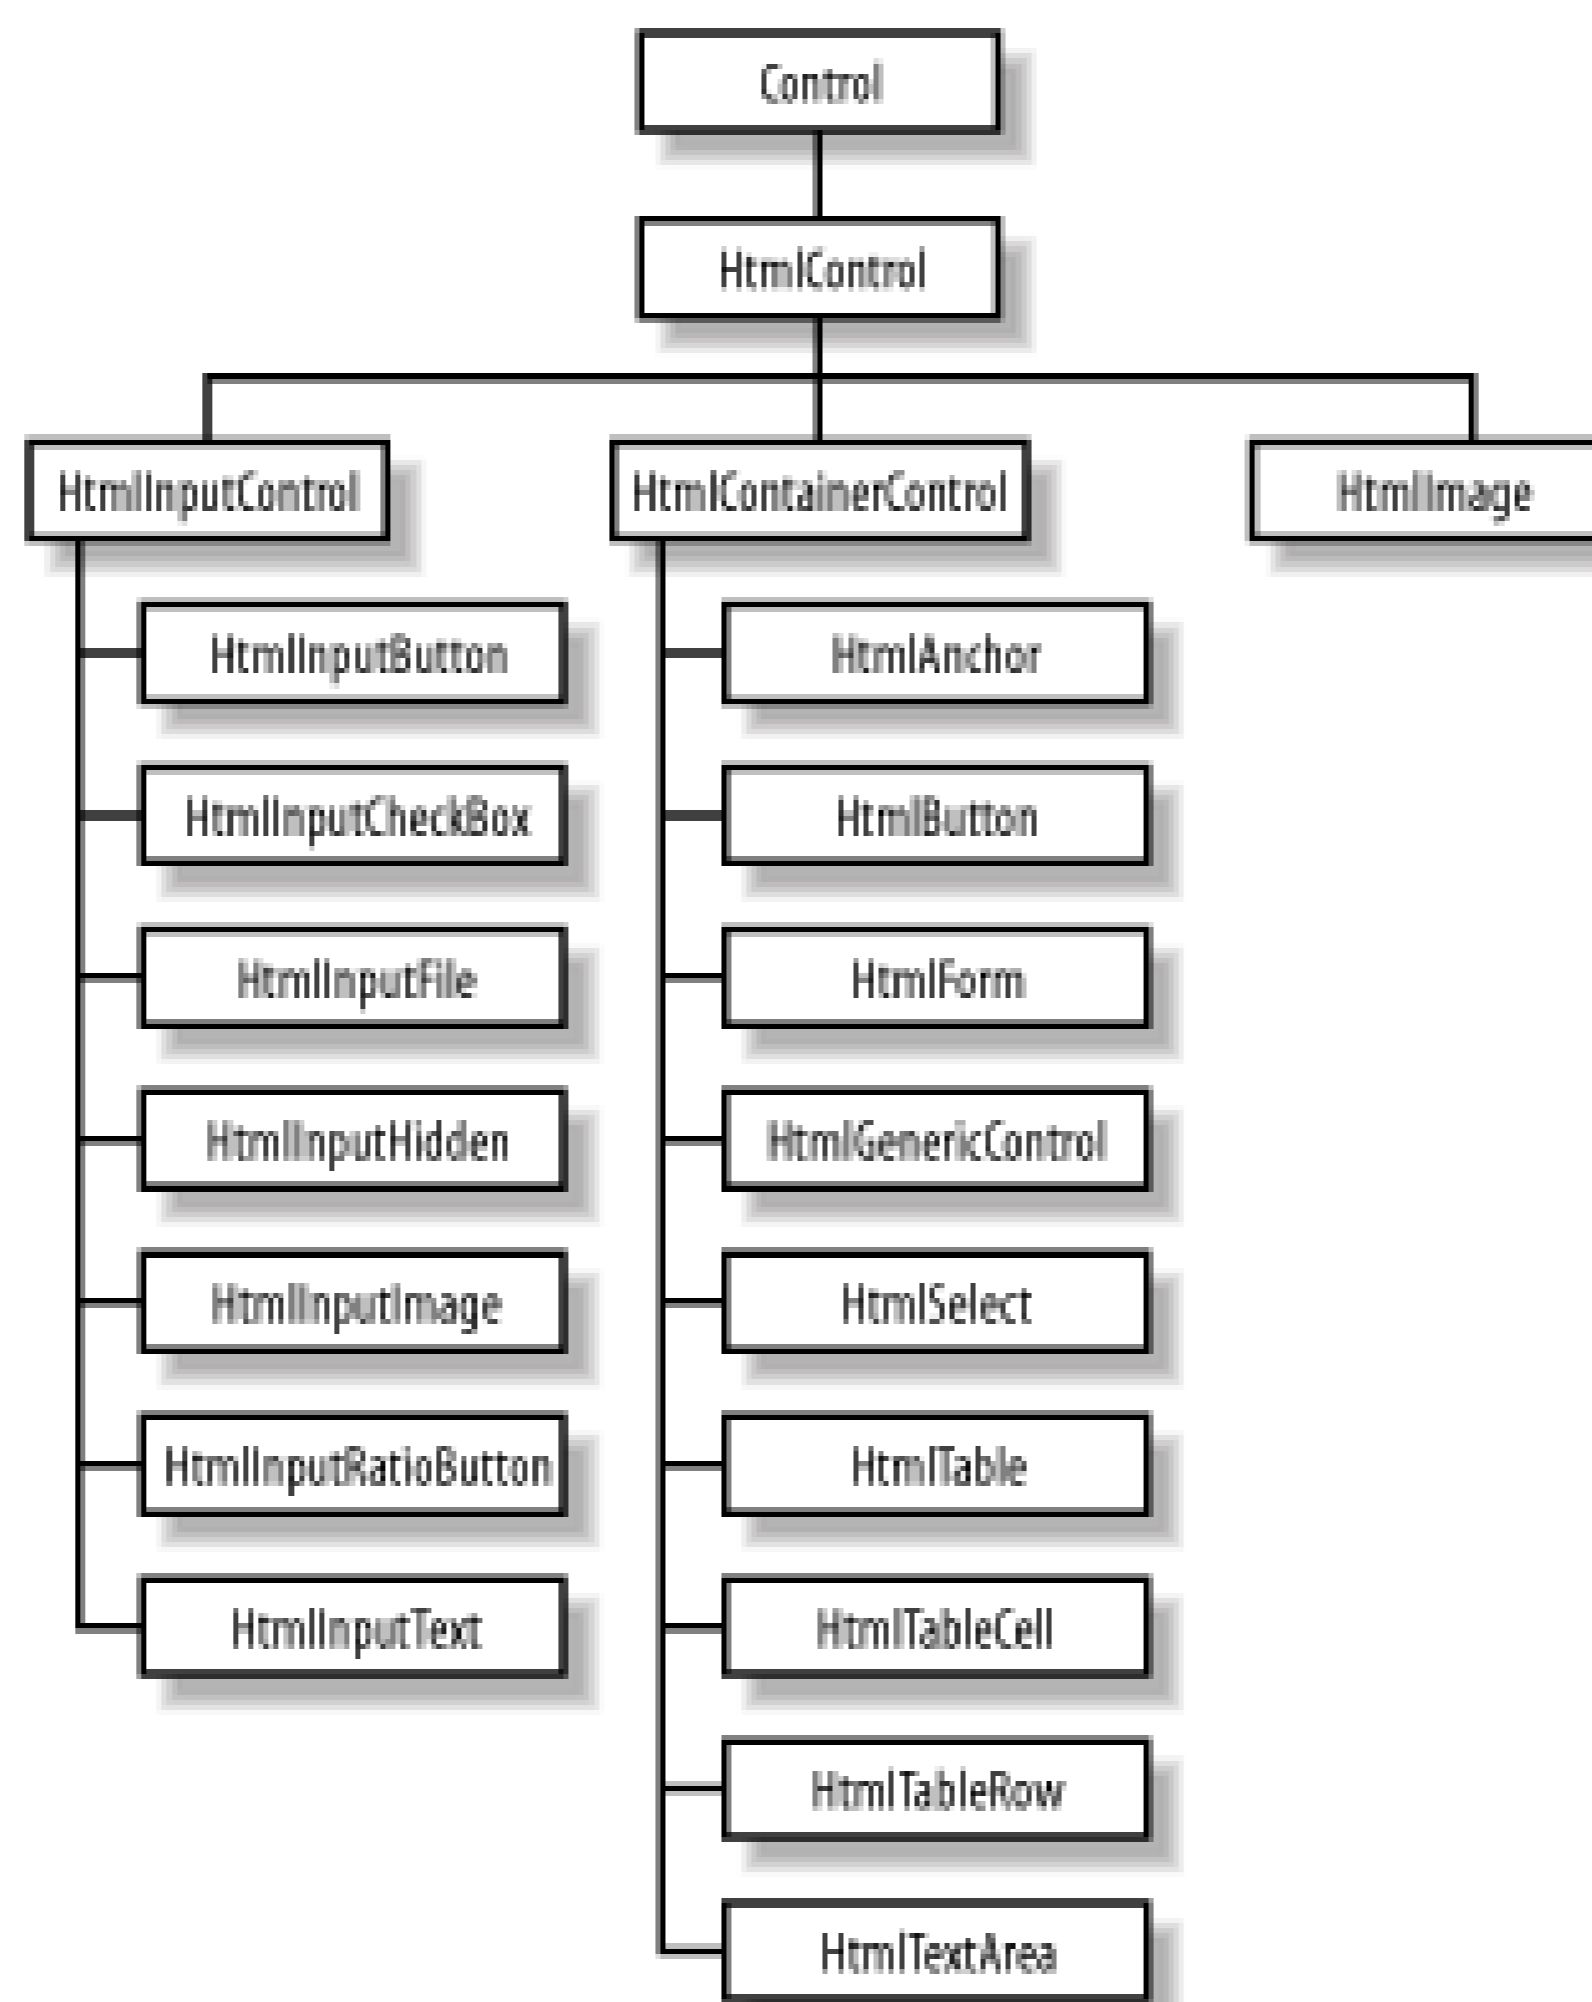


Table 7-1 maps the HtmlControls to standard HTML tags. This means when you have an HTML tag that is flagged to run on the server side with `runat="server"`, ASP.NET creates an appropriate HtmlControl that you can program against.

Table 7-1. HtmlControls mapping to HTML tags

HTMLControl	Description	HTML tag
HtmlImage	Image tag	<code></code>
HtmlInputFile	File selector	<code><input type="file"></code>
HtmlInputHidden	Used for hidden form fields	<code><input type="hidden"></code>
HtmlInputImage	Image input	<code><input type="image"></code>
HtmlInputRadioButton	Radio button	<code><input type="radio"></code>
HtmlInputText	Standard text input	<code><input type="text"></code>
HtmlInputButton	Standard HTML button	<code><input type="button"></code>
HtmlInputCheckBox	Standard HTML checkbox	<code><input type="checkbox"></code>
HtmlForm	Form tag	<code><form></code>
HtmlGenericControl	Miscellaneous generic HTML tags	<code><span, div, etc.></code>

HTMLControl	Description	HTML tag
HtmlSelect	Standard HTML drop-down control	<code><select></code>
HtmlTable	Standard HTML table	<code><table></code>
HtmlTableCell	A cell in a table row	<code><td></code>
HtmlTableRow	A row in a table	<code><tr></code>
HtmlTextArea	Multiline text area	<code><textarea rows=n cols=n></code>
HtmlAnchor	Standard HTML hyperlink control	<code></code> Or <code></code>
HtmlButton	HTML button	<code><button></code>

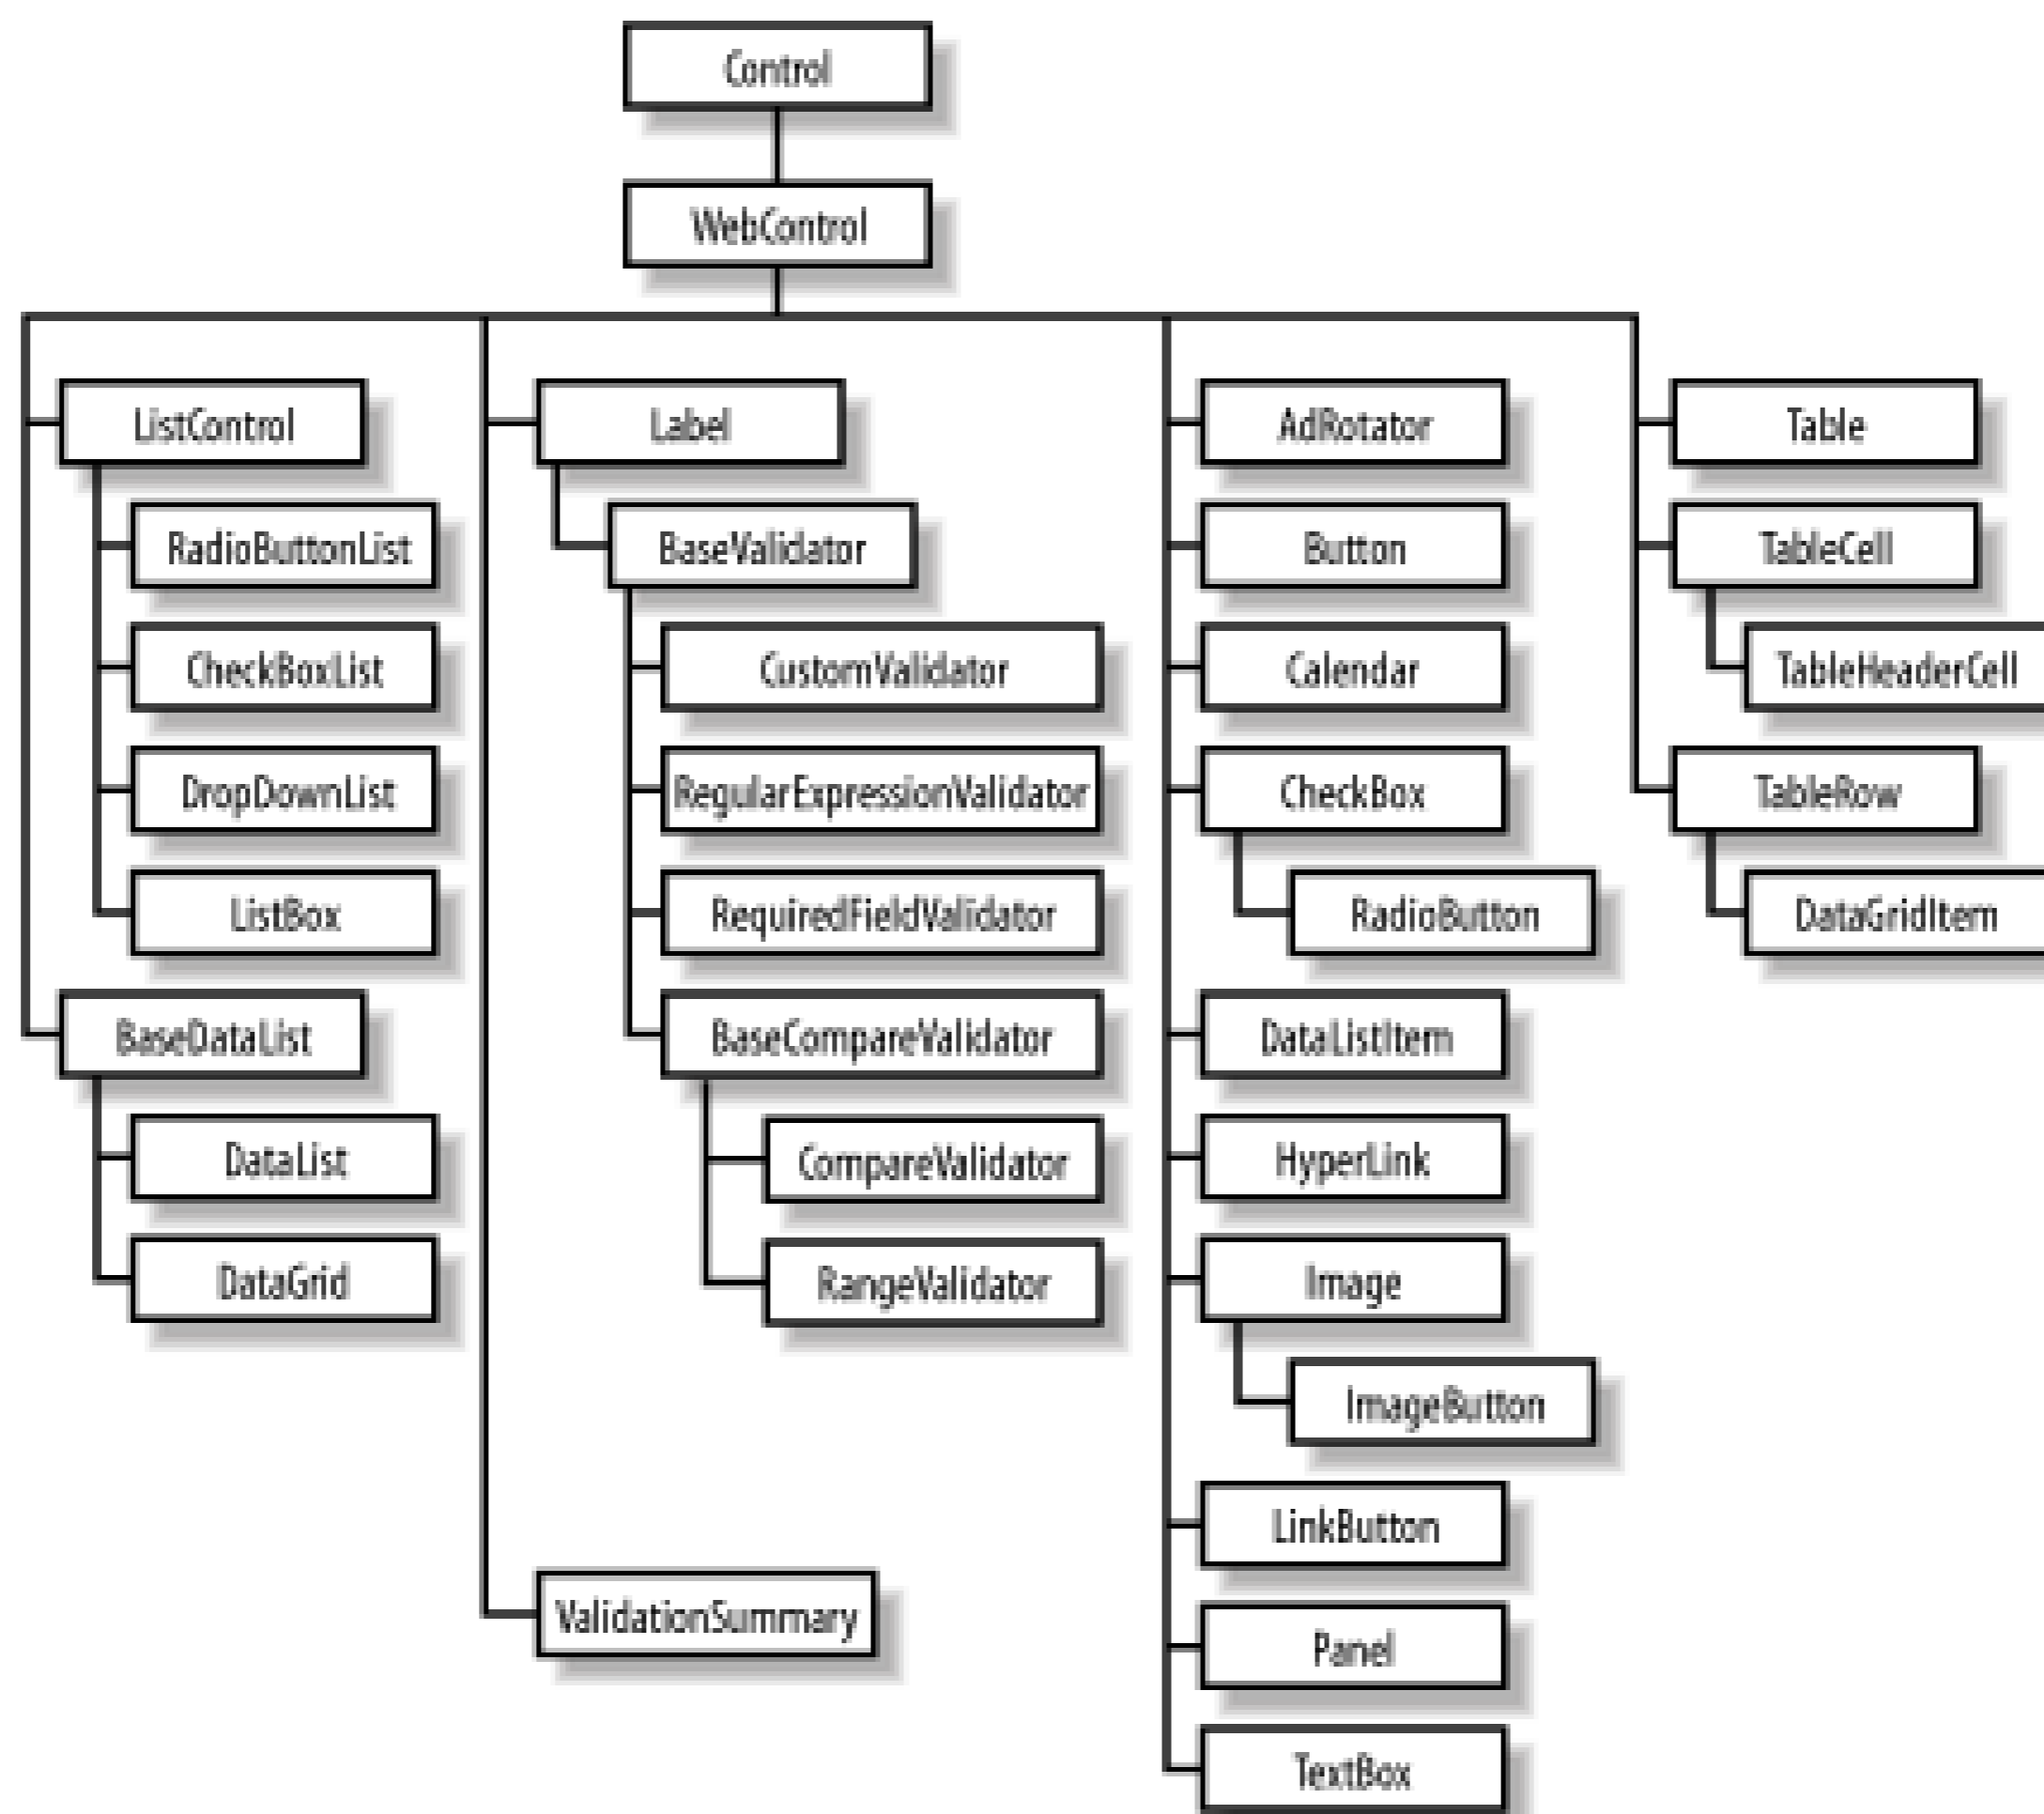
7.3.5 System.Web.UI.WebControls Namespace

While providing HtmlControls, which map to standard HTML elements, ASP.NET also provides another group of UI controls, the WebControl class (see [Figure 7-3](#)). In addition to providing all traditional controls similar to HtmlControls, WebControls also provide much richer controls such as calendars, grids, and validators.

WebControls are richer, more powerful, and more flexible than HtmlControls. It seems that it is the natural choice for new ASP.NET applications; however, HtmlControls are better if you are migrating ASP applications. Another thing that might make you consider using HtmlControls is that with it, your client-side scripts can still access and manipulate the objects.

Most classes in this namespace are based on WebControl, which is again a derivative of the Control class. The WebControl class provides the common properties and methods inherited by all of its descendants, including access key, tab index, tool tip, color, font, and border setting.

Figure 7-3. WebControls object hierarchy



[[Team LiB](#)]

[\[Team LiB \]](#)

7.4 Web Form Syntax

Similar to Active Server Pages, Web Forms are text files consisting of HTML tags and other controlling tags such as directives and script blocks. The default extension for web forms is *aspx*; however, you can use the IIS administration tool to map other file extensions explicitly with *aspnet_isapi.dll* to have them parsed and compiled when accessed, as if they were ASP.NET resources.

There are 10 different syntax elements in ASP.NET; because most of them are carried over from ASP, we list here the familiar ones and discuss only those that are important in ASP.NET:

- Directives
- Code declaration blocks
- Code-rendering blocks
- HTML control syntax
- Custom control syntax
- Data-binding expressions
- Server-side object tags
- Server-side include directives
- Server-side comments
- Literal text

7.4.1 Directives

Previously, all ASP directives were formatted as `<%@ [attribute=value]+ %>` because there was only one type of directive.^[3]

^[3] As noted in the [Preface](#), the plus sign indicates one or more instances of the preceding term in brackets.

ASP.NET adds a number of directives to ASP.NET files. With the new list of directives-Page, Control, Import, Implements, Register, Assembly, OutputCache, and Reference-the syntax for directive is now `<%@ directive [attribute=value]+ %>`. All of the old ASP directives are attributes under the Page directive. If you use the old syntax by omitting the directive name, the attribute/value pairs will be applied to the default directive, which is Page.

7.4.1.1 @ Page

In addition to containing all previous ASP directives (CodePage, EnableSessionState, Language, LCID and Transaction), the ASP.NET Page directive also supports the important attributes ErrorPage, Inherits, Src, and EnableViewState, which we will make use of in this chapter. The complete list of all attributes for the Page directive can be found in the *.NET Framework Developers' Guide*.

```
<@ Page Language="VB" ErrorPage="URL" EnableViewState="true">
```

7.4.1.2 @ Control

Similar to the way the Page directive is used for an ASP.NET page (an *.aspx* file), the Control directive is used for an ASP.NET control (an *.ascx* file). (We get into developing ASP.NET controls in [Section 7.5.4](#) later in this chapter.)

7.4.1.3 @ Import

We can use the Import directive to add namespace references to the current page. Your code can access all classes and interfaces of imported namespaces. For example, if you want to use ADO.NET, you would include the following code:

```
<%@ Import Namespace="System.Data" %>
<%@ Import Namespace="System.Data.OleDb" %>
```

A number of namespaces are automatically imported into all ASP.NET pages to simplify the developers' task:

- System
- System.Collections
- System.Collections.Specialized
- System.Configuration
- System.IO
- System.Text
- System.Text.RegularExpressions
- System.Web
- System.Web.Caching
- System.Web.Security
- System.Web.SessionState
- System.Web.UI
- System.Web.UI.HtmlControls
- System.Web.UI.WebControls

7.4.1.4 @ Implements

Because an *.aspx* file is basically a Page class derivative, it too can also implement an interface. The @ Implements directive is used to declare that the *.aspx* implements the specified interface. For example, the following line declares that the page implements the IPostBackEventHandler interface:

```
<%@ Implements Interface="System.Web.UI.IPostBackEventHandler" %>
```

7.4.1.5 @ Register

This directive registers custom server controls for use in the current page by specifying the aliases to be used as prefixes for class names. It is used in conjunction with the custom server-control elements to provide a concise way of specifying server-control names. The following line of code registers a custom control to be used in this page:

```
<%@ Register Tagprefix="Ch07"
            TagName="MyCustomControl
            Src="MyCustomControl.ascx" %>
```

The name of the control is `MyCustomControl`; the prefix used when declaring the control is `Ch07`; the source for the control is in `MyCustomControl.ascx`. (We demonstrate this when we develop our ASP.NET controls later in this chapter.)

If the server custom control was not done through `.ascx` but instead, through inheriting `UserControl`, the syntax for the `Register` directive is the following:

```
<%@ Register Tagprefix="MyCustomCtrl"
    Namespace="MyNamespace"
    Assembly="MyDll.dll" %>
```

7.4.1.6 @ Assembly

The `Assembly` directive specifies the assembly to which the current page belongs. This effectively makes all the classes and interfaces belonging to the assembly accessible to the current page. For example, the following line of code specifies that the current page belong to the `Ch07` assembly:

```
<%@ Assembly Name="Ch07" %>
```

This means that code in this page can access anything in the `Ch07` assembly as long as the `Ch07` assembly is compiled and placed in the [bin](#) subdirectory of this web application.

7.4.1.7 @ OutputCache

You can use the `OutputCache` directive to control the output-caching duration for the current page. This is similar to setting up the expiration for the response object in ASP programming. The `Duration` attribute of the `OutputCache` directive defines the time in seconds until the page expires.

7.4.1.8 @ Reference

The `@ Reference` directive is used to add a reference to a page or a control to this `aspx` page.

7.4.2 Code Declaration Blocks

As in ASP, *code declaration blocks* define the code to be parsed and run for the page. In these blocks, the `runat` attribute specifies whether the code block is client-side or server-side code. For server-side programming, set this attribute to `server`. If you ignore the `runat` attribute, IIS will interpret the code block as client-side code, which is used for Dynamic HTML (DHTML).

```
<script runat="server" [language="code language"]>
    Code
</script>
```

For both client- and server-side code declaration blocks, you can also use the `src` attribute to point to an external source file containing the code. This helps separate the code from the HTML content of the page. The value for `src` can be a relative path or a URL to a source file. The URL can be on the same or a different web server:

```
<script runat="server"
  [language="code language" ]
  [src="externalfilename" ] />
```

7.4.3 Code-Rendering Blocks

There are no changes to this syntax versus that in ASP. Inline code or inline expressions specified in these code-rendering blocks are executed when the page is rendered. All these blocks are enclosed between the tags `<%` and `%>`. The language used in these tags is specified in the language attribute of the Page directive.

7.4.4 HTML-Control Syntax

HTML controls are very similar to standard HTML elements, with the exception of the `id` and the `runat` attributes. If you've developed web applications with DHTML, you should be familiar with the `id` attribute of an HTML element and how to programmatically reference the client-side control representing the HTML element. The difference in this case is that the control is not on the client side but on the server side. For example, the following code represents an HTML server button control:

```
<input id="cmd1" runat="server"
  type="button" value="Click Me" />
```

All HTML server controls must be inside a `<form runat="server">` control because Web Forms use the POST method to maintain the controls' states.

When encountering an HTML element tagged with `id` and the `runat` attribute set to `server`, ASP.NET creates the appropriate scriptable server `HtmlControl` object. For example, the previous HTML snippet generates a server `HtmlControl` of type `HtmlInputButton` that has an `id` of `cmd1`.

You can bind an event handler to this control's event to handle notification from this control, such as the `onclick` event. There are two ways to bind an event handler to a control's event, the declarative way and the programmatic way. The declarative is done inside the HTML element tag as an attribute/value pair. The attribute is the name of the event, and the value is the name of the event-handling function. For example, to handle the `onclick` event, add this to the previous HTML tag:

```
onserverclick="handleServerClick"
```

The programmatic way to bind an event to the handler involves a line of code that assigns a delegate to the event property of the control. In C#, the code to bind the `ServerClick` event of the button to the event handler `handleServerClick` is:

```
cmd1.ServerClick += new System.EventHandler(handleServerClick);
```

If you've used client-side DHTML in your web applications, event binding should be nothing new to you, except for some subtle differences. The first difference is obvious: the event handler runs on the server before the page is sent back to the browser, instead of running on the client side. The other difference is that all event-handler functions for server-side must have two parameters: `Sender` and `Event`. The `Sender` parameter is of type `object`, indicating the source element that caused the event to happen; the `Event` parameter is of type `EventArgs`, which is the actual event fired. In DHTML scripting, we would inspect the `window.event` object to find out which element was the source of the

event and other event information.

7.4.5 Custom-Control Syntax

Similar to HTML Server Controls, *custom controls* also have `id` and `runat` attributes; however, custom controls are not standard HTML elements. To insert a custom control into a page, use the following syntax:

```
<tagprefix:tagname id="controlID" runat="server" eventname="eventHandler" />
```

Notice that all custom controls' tags have a tag prefix, which is an alias to the namespace in which the control is defined. See the Register directive earlier in this chapter for information on registering namespaces' aliases. Binding events to their handlers for custom controls is the same as for HTML controls. Even though we show the two ways of binding events, it is preferable to bind events using the second method because it cleanly separates the HTML tags from the code behind the screen.

All web controls mentioned in the WebControls namespace can be inserted in the same manner (these controls have the prefix `asp`). For example, you can have the following tags in your *aspx* page:

```
<asp:TextBox id=txt1 runat=server></asp:TextBox>
<asp:Button id=cmd1 runat=server Text="Web Button"></asp:Button>
<asp:Label id=label1 runat=server></asp:Label>
```

These tags result in three objects generated from the three classes: `TextBox`, `Button`, and `Label`, from the `System.Web.UI.WebControls` namespace. In your server script, you can access and manipulate these objects to render your page appropriately.

7.4.6 Data-Binding Expressions

Data-binding expressions bind the server controls with some data sources. The syntax to bind data is:

```
<%# data-binding-expression %>
```

Examine the following block of code to see the simplest data binding:

```
<asp:Label text='<%# TestData %>' runat=server/>
```

The data-binding expression here indicates that the label's text content is bound to a publicly defined property, `TestData`, of the Web Form. This means that when data binding occurs for the form, `<%# TestData %>` will be replaced by the content of the `TestData` property. Let's define this property for the Web Form:

```
public string TestData = "Hello World";
```

The Web Forms page framework does not perform data binding automatically. The developers must explicitly call the `DataBind()` method to activate the evaluation of the data-binding expression and perform the substitution. We can call the page's `DataBind` method upon the page-load event or

whenever we change the `TestData` property and want it reflected on the page. This example calls `DataBind()` upon page load to bind the Label's text to the `TestData` variable:

```
<html>
  <head><title>Data Binding Sample</title></head>
  <body>

    <script language="C#" runat=server>
      /* Declare the variable we want to bind to. */
      public string TestData;
      void Page_Load(Object oSender, EventArgs oEvent) {
        TestData = "Hello World!\n";
        Page.DataBind( );
      }
    </script>

    <asp:Label text='<%=# TestData %>' runat=server/>

  </body>
</html>
```

Let's try something a little more complicated. In the next block of tags, we have three labels bound to three different properties of an object called `currStudent`:

```
Name: <asp:Label text='<%=# currStudent.FirstName %>' runat=server/>
<asp:Label text='<%=# currStudent.LastName %>' runat=server/> <br/>
SSN: <asp:Label text='<%=# currStudent.SSN %>' runat=server/>
```

The `currStudent` object is a publicly accessible property of the current page:

```
<script language="C#" runat=server>
  public class CStudent {
    /* Declare the variable we want to bind to. */
    public string FirstName;
    public string LastName;
    public string SSN;
  }
  public CStudent currStudent;

  void Page_Load(Object oSender, EventArgs oEvent) {
    currStudent = new CStudent( );
    currStudent.FirstName = "Jack";
    currStudent.LastName = "Daniel";
    currStudent.SSN = "123-45-6789";
    Page.DataBind( );
  }
</script>
```

You can have this `currStudent` object filled with data coming from any source then perform a `DataBind` call to update the page with the current student's information. The assumption here, of course, is that the `Student` class provides the previously mentioned properties.

7.4.7 Server-Side Object Tags

Server-side object tags statically declare and instantiate COM and .NET objects. The syntax to declare server-side objects in *global.asax* is:

```
<object id="id" runat="server" scope="scope" class=".NET class name" >  
<object id="id" runat="server" scope="scope" progid="COM ProgID" >  
<object id="id" runat="server" scope="scope" classid="COM classID" >
```

Scope can be pipeline, application, or session, which means the object is available on the page, as an application variable, or as a session variable, respectively. To dynamically add a server-side object to the page, you would use the `Page.LoadControl()` method or just instantiate the control directly.

7.4.8 Other Elements

Server-side includes server-side comments and literal text which are exactly the same as in ASP. Therefore, we will not go over them here.

[\[Team LiB \]](#)

[[Team LiB](#)]

7.5 ASP.NET Application Development

In conventional ASP programming, developers typically access the Request object to get the parameters needed to render the page and render the content of the page through either the Response object or code-rendering blocks. We also use other ASP objects such as the Application, Session, and Server objects to manage application variables, session variables, server settings, and so on.

As mentioned earlier, ASP.NET is intended to change all this spaghetti madness by introducing a much cleaner approach to server-side scripting framework: Web Forms, or programmable pages, and server controls.

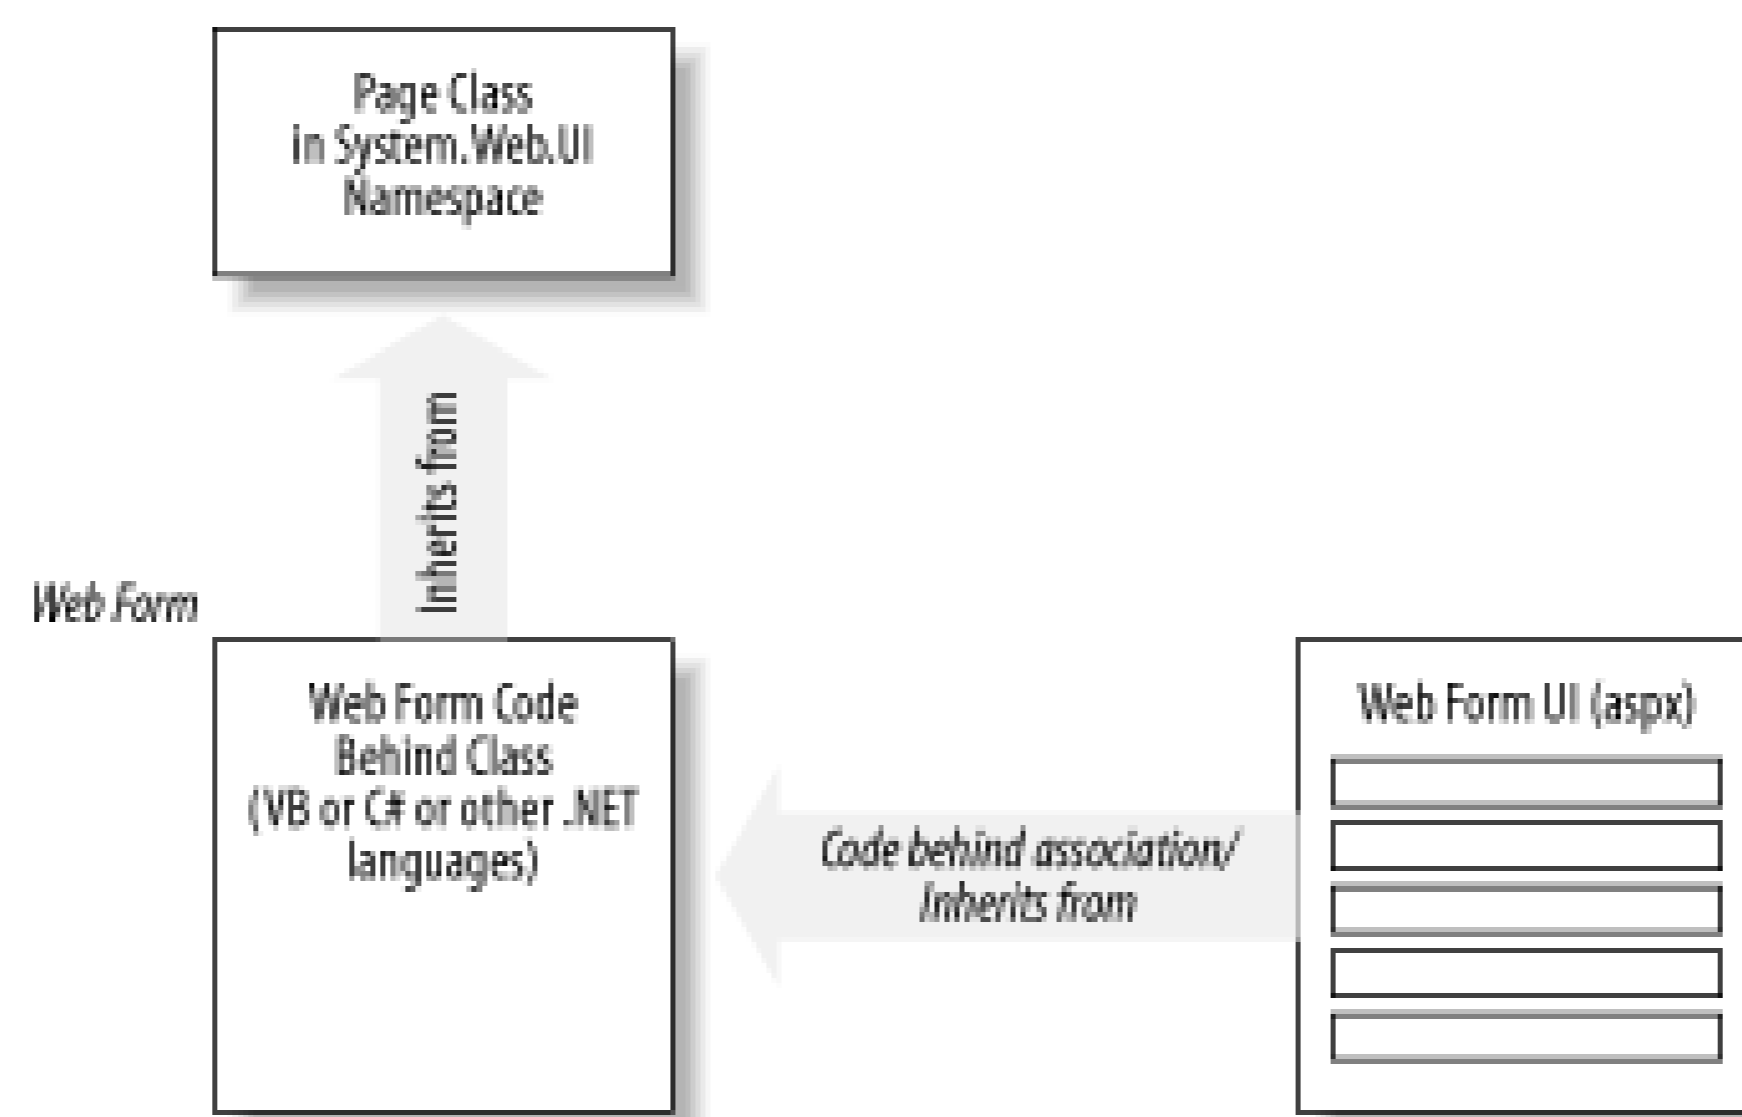
In the following sections, we cover the components of a Web Form, its life cycles, the server controls that the Web Form contains, event-handling for these server controls, as well as how to create your own server controls.

7.5.1 Web Form Components

Similar to VB Forms, a Web Form consists of two components: the form with its controls, and the code behind it that handles events associated with the form's controls. A Web Form has the file extension *.aspx* and contains HTML elements, as well as server controls. The code behind the form is usually located in a separate class file. Note that while it is possible to have both the form and the code in one file, it is better to have separate files. This separation of user interface and application code helps improve the spaghetti-code symptom that most ASP-based applications are plagued with.

ASP.NET provides the Page class in the System.Web.UI namespace. This class encapsulates all common properties and methods associated with web pages. The code behind the class derives from this Page class to provide extensions specific to the page we're implementing. The *aspx* file provides the form layout and control declarations. [Figure 7-4](#) illustrates the relationship between the Page base class, the Web Form code behind the class, and the Web Form user interface (UI).

Figure 7-4. Web Form components



As a Web Form developer, you will have to provide the latter two. The Web Form UI is where you declare server controls with appropriate IDs. The code behind the class is where you programmatically access server controls declared in the Web Form UI, as well as handle events from these controls. The following simple example shows the *aspx* page, the code behind source file, and how they work together. The *aspx* file (*TestEvent.aspx*) contains only HTML tags and a directive that links to the code behind:

```
<%@ Page language="c#" codebehind="TestEvents.cs" inherits="CTestEvents"%>
<html>
  <head><title>Testing Page Events with codebehind</title></head>
  <body>
    <form runat=server>
      Init Time: <asp:Label id=labelInit runat=server/><br/>
      Load Time: <asp:Label id=labelLoad runat=server/><br/>
      <input type=submit />
    </form>
  </body>
</html>
```

The code-behind, *TestEvents.cs*, contains the class *CTestEvents* to which the *aspx* page is referring:

```
using System;

public class CTestEvents : System.Web.UI.Page {
  protected System.Web.UI.WebControls.Label labelInit;
  protected System.Web.UI.WebControls.Label labelLoad;

  public CTestEvents( ) {
    labelInit = new System.Web.UI.WebControls.Label( );
    labelLoad = new System.Web.UI.WebControls.Label( );
  }

  public void Page_Init(Object oSender, EventArgs oEvent) {
    labelInit.Text = DateTime.Now.ToString( );
  }

  public void Page_Load(Object oSender, EventArgs oEvent) {
```

```

        labelLoad.Text = DateTime.Now.ToString(  );
        if(IsPostBack) {
            labelLoad.Text += "(PostBack)";
        }
    }
}

```

You must compile *TestEvents.cs* and place the DLL in the */bin* directory under your web application's virtual directory before trying to access the *aspx* page.^[4]

[4] The Web Application directory is the root virtual directory where your web application resides. To set up the virtual directory, use the IIS Administration Tool.

The command to compile this C# file is:

```
csc /t:library TestEvents.cs
```

ASP.NET parses the Web Form files to generate a tree of programmable objects, where the root is the Page-derived object representing the current Web Form. This is similar to how the IE browser parses the HTML file and generates a tree of scriptable objects to be used in DHTML; however, the tree of objects for the Web Form files resides on the server side.

As you are already aware from our survey of the System.Web.UI namespace, the Page class actually derives from the Control class. In a sense, a Web Form is a hierarchy of Control-derived objects. These objects establish the parent-child relationship through the Parent and Controls properties.

Besides the Controls and Parent properties, the Page class also provides other useful properties, which are familiar to ASP developers—such as the Request, Response, Application, Session, and Server properties.

Because the Web Form is nothing but a programmable page object, using this object-oriented model is much more intuitive and cleaner than the conventional ASP development. As opposed to the linear execution of server-side scripts on an ASP page, ASP.NET enables an event-based object-oriented programming model.

Let's take an example of a web page that contains a form with numerous fields. One or more of these fields display list information from a database. Naturally, we have code in the ASP page to populate these fields so that when a user requests this ASP page, the generated page would have the content ready. As soon as the last line of data is written to the browser, the ASP page is done. This means that if there were errors when the user submits the form, we will have to repopulate all the database-driven form fields, as well as programmatically reselect values that the user chose prior to submitting the form. In ASP.NET, we don't have to repopulate the database-driven fields if we know that the page has already been populated. Furthermore, selected values stay selected with no manual handling. The next couple of sections describe the concept in more detail.

7.5.1.1 Web Form events

The Page class exposes events such as Init, Load, PreRender, and Unload. Your job as a developer is to handle these events and perform the appropriate task for each of these stages. This is much better than the linear execution model in ASP programming because you don't have to worry about the location of your initialization scripts.

The first event that happens in the life of a Web Form is the Init event. This is raised so that we can have initialization code for the page. Please note that because the controls on the page are not yet created at this point, this initialization should only contain code that does not rely on any controls. This event is raised once for each user of the page.

Most developers are more familiar with the Load event that follows the Init event. Subsequently, it is raised each time the page is requested. When this event is raised, all child controls of the Web Form are loaded and accessible. You should be able to retrieve data and populate the controls so that they can render themselves on the page when sent back to the client.

The following example shows the how the Init and Load events can be handled in ASP.NET. In this example, we show both the HTML and its code together in one file to make it simpler:

```
<html>
  <head><title>Testing Page Events</title></head>
  <body>

    <script language="C#" runat="server">
      void Page_Init(Object oSender, EventArgs oEvent) {
        labelInit.Text = DateTime.Now.ToString( );
      }

      void Page_Load(Object oSender, EventArgs oEvent) {
        labelLoad.Text = DateTime.Now.ToString( );
        if(IsPostBack) {
          labelLoad.Text += "(PostBack)";
        }
      }
    </script>

    <form runat="server">
      Init Time: <asp:Label id="labelInit" runat="server"/><br />
      Load Time: <asp:Label id="labelLoad" runat="server"/><br />
      <input type="submit" />
    </form>
  </body>
</html>
```

The first time you access this page, the Init event happens, followed by the Load event. Because these events happen quickly, both the Init Time and Load Time will probably show the same time. When you click on the submit button to cause the page to reload, you can see that the Init Time stays what it was, but the Load Time changes each time the page is reloaded.

The PreRender event happens just before the page is rendered and sent back to the client. We don't often handle this event; however, it depends on the situation. You might want to alter the state of some of the "server-side" objects before rendering the page content.

The last event in the life of a Web Form is the Unload event. This happens when the page is unloaded from memory. Final cleanup should be done here. For example, while unloading you should free the unmanaged resources that you've allocated at the Init event.

Beside these page-level events, controls on the page can raise events such as ServerClick and ServerChange for HtmlControls, as well as Click, Command, CheckedChanged,

SelectedIndexChanged, and TextChanged events for WebControls. It is the handling of these events that makes ASP.NET truly dynamic and interactive.

7.5.2 The Life Cycle of a Web Form

In ASP, the web page starts its life when a client requests a particular page. IIS parses and runs the scripts on the ASP page to render HTML content. As soon as the page rendering is complete, the page's life ceases. If you have forms that pass data back to the ASP page to be processed, the ASP page runs as a new request, not knowing anything about its previous states. Passing data back to the original page for processing is also referred to as postback.

In ASP.NET, things are a little different. The page still starts at the client's request; however, it appears to stay around for as long as the client is still interacting with the page. For simplicity's sake we say that the page stays around, but in fact, only the view states of the page persist between requests to the page. These view states allow the controls on the server to appear as if they are still present to handle server events. We can detect this postback state of the page via the `IsPostBack` property of the Page object and forego certain costly reinitialization. The handling of events during these postbacks is what makes ASP.NET so much different than conventional ASP development.

In the following example, we extend the previous example to handle the postback. When the Load event is handled for the first time, we populate the drop-down list box with data. Subsequently, we indicate only the time the event is raised without reloading the data. This example also demonstrates the server event handler `handleButtonClick` that was bound to the `ServerClick` event of the button:

```
<html>
  <head><title>Testing Page Events</title></head>
  <body>

    <script language="C#" runat="server">
      void Page_Init(Object oSender, EventArgs oEvent) {
        labelInit.Text = DateTime.Now.ToString( );
      }

      void Page_Load(Object oSender, EventArgs oEvent) {
        labelLoad.Text = DateTime.Now.ToString( );
        if(!IsPostBack) {
          selectCtrl.Items.Add("Acura");
          selectCtrl.Items.Add("BMW");
          selectCtrl.Items.Add("Cadillac");
          selectCtrl.Items.Add("Mercedes");
          selectCtrl.Items.Add("Porche");
        } else {
          labelLoad.Text += " (Postback)";
        }
      }

      void handleButtonClick(Object oSender, EventArgs oEvent) {
        labelOutput.Text = "You've selected: " + selectCtrl.Value;
        labelEvent.Text = DateTime.Now.ToString( );
      }
    </script>
  </body>
</html>
```



```

</script>

<form runat="server">
  Init Time: <asp:Label id="labelInit" runat="server"/><br/>
  Load Time: <asp:Label id="labelLoad" runat="server"/><br/>
  Event Time: <asp:Label id="labelEvent" runat="server"/><br/>
  Choice: <select id="selectCtrl" runat="server"></select><br/>
  <asp:Label id="labelOutput" runat="server"/><br/>
  <input type="button" value="update"
        OnServerClick="handleButtonClick" runat="server" />
</form>

</body>
</html>

```

The life cycle of a Web Form consists of three main stages: Configuration, Event Handling, and Termination. As mentioned earlier, these stages span across many requests to the same page, as opposed to the serving-one-page-at-a-time policy found in ASP.

Init and Load events map to the configuration stage (i.e., when the page is first requested and each time the page is reloaded via postbacks). Events such as Button.Click and ListControl.SelectedIndexChanged map to the Event Handling stage, where user interactions with the page are handled via postbacks. The Termination stage involves the Dispose method and the Unload event. The postbacks allow the ASP.NET web page to appear like a continuous thread of execution while spanning multiple, full round-trips. In the old ASP world, each round-trip is a brand new request to the page unless the developer has built in a custom and elaborated framework similar to that of ASP.NET.

7.5.2.1 Configuration

In the Configuration stage, the page's Load event is raised. It is your job to handle this event to set up your page. Because the Load event is raised when all the controls are already up and ready, your job is now to read and update control properties as part of setting up the page. In the previous code example, we handled the Load event to populate the drop-down list with some data. We also updated the labelLoad control's Text to display the time the Load event happens. In your application, you will probably load the data from a database and initialize form fields with default values.

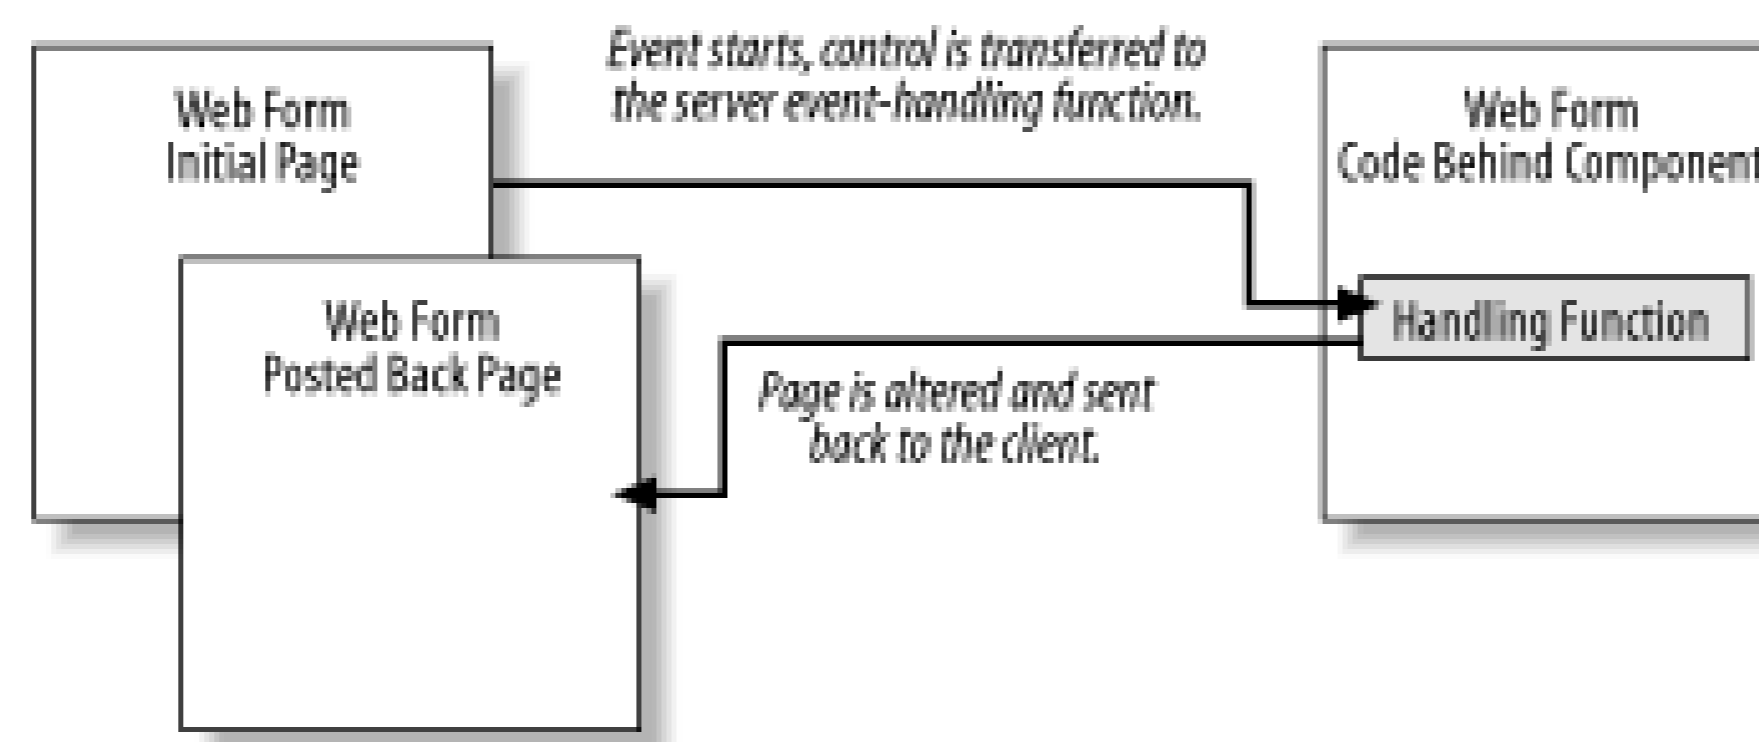
The page's IsPostBack property indicates whether this is the first time the page is loaded or if it is a postback. For example, if you have a control that contains a list of information, you will only want to load this control the first time the page is loaded by checking the IsPostBack property of the page. When IsPostBack is `true`, you know that the list control object is already loaded with information. There is no need to repopulate the list. In the previous code example, we skipped over the population of the drop-down and just displayed a string "(Postback)".

You might need to perform data binding and re-evaluate data-binding expressions on the first and subsequent round trips to this page.

7.5.2.2 Event handling

In this middle stage, the page's server event-handling functions are being called as the result of some events being triggered from the client side. These events are from the controls you've placed on the Web Form. [Figure 7-5](#) depicts the life cycle of an event.

Figure 7-5. The Web Form event life cycle



7.5.2.3 Termination

At this stage, the page has finished rendering and is ready to be discarded. You are responsible for cleaning up file handles, releasing database connections, and freeing objects. Although you can rely on the CLR to perform garbage collection of managed resources for you, we strongly advise you to clean up after yourself because garbage collection only happens periodically. On heavily loaded systems, if the garbage-collection cycle is not optimal, the unfreed resources can still exhaust memory and bring your system to a halt.^[5]

[5] These so-called logical memory leaks equate to areas where large blocks of memory (managed or unmanaged) were allocated and used for a brief amount of time but won't be free automatically until it is out of scope. Because in ASP.NET the scope from the configuration phase to the termination phase is rather small, this problem does not appear to be that bad. For Windows Form applications or Windows Services, this problem, if undetected, could be catastrophic. The way to expedite the deallocation of objects is to set long-lived objects to null (or nothing in VB). If the object implements the `IDisposable` interface, call its `Dispose` method as soon as you can. The garbage collector won't pick up this kind of unintentional memory usage if you have references to objects that are unused, yet remain in scope throughout the life of the application.

We can perform the cleanup for the previous example with the `Unload` event handler shown as follows. Because there is nothing to clean up in this simple example, we just show you the function as a template:

```

void Page_Unload(Object oSender, EventArgs oEvent) {
    // Cleaning up code here
}
  
```

7.5.3 Server Controls

As we saw from the `System.Web.UI.HtmlControls` and `System.Web.UI.WebControls` namespaces, server controls are programmable controls that run on the server before the page is rendered by ASP.NET. They manage their own states between requests to the same page on the server by inserting a hidden field that stores the view state of the form. This eliminates the need to repopulate the value of form fields with the posted value before sending the page back to the client.

Server controls are also browser-independent. Because they are run on the server side, they can rely on the `Request.Browser` property to get the client's capability and render appropriate HTML.

Since the server controls are just instantiations of .NET classes, programming the server controls yields easy-to-maintain code. Especially when you have custom server controls that encapsulate other controls, web application programming becomes nothing more than gluing these blocks together.

All HTML controls and web controls mentioned in `System.Web.UI.HtmlControls` and `System.Web.UI.WebControls` are server controls shipped with ASP.NET.

7.5.4 Custom Server Controls

As you become more familiar with the ASP.NET framework and the use of server controls on your Web Form, you will eventually need to know how to develop these server controls yourself. In ASP.NET, there are two ways of creating custom server controls: the *pagelet* approach, which is easy to do but rather limited in functionality, and the Control base class (or UserControl) *derivative* approach, which is more complicated but also more powerful and flexible.

7.5.4.1 Pagelets

Until recently, code reuse in ASP development has been in the form of server-side includes. If you have common UI blocks or scripts, you can factor them into an include file. Use the syntax `<!-- #include file="url" -->` to include the common file into the main page to return to the browser. This approach is fine, but it has serious limitations. The main thing is to make sure the HTML tag IDs and script variable names are unique. This is because IIS does nothing more than merge the include file when it parses server-side includes. The include file ends up being in the same scope with the container file. You cannot include the same file twice because there will be tag ID and script conflicts.

With ASP.NET, you can factor out common HTML and scripts into what is currently called a pagelet and reuse it without worrying about the ID conflicts. A pagelet is a Web Form without a body or a form tag, accompanied by scripts. The HTML portion of the pagelet is responsible for the layout and the user interface, while the scripts provide the pagelet with programmability by exposing properties and methods. Because the pagelet is considered a user control, it provides an independent scope. You can insert more than one instance of the user control without any problem.

The container Web Form must register the pagelet as a user control with the `@Register` directive and include it on the page with the `<prefix:tagname>` syntax. If more than one copy of the pagelet is used in a container page, each of them should be given different IDs for the container page's script to work correctly. The script on the container Web Form can access and manipulate the pagelet the same way it does any other server controls. The next example shows how an address form is reused as a pagelet. You might display this address form to allow the web user to register with your application or to display the shipping and billing addresses when the web user checks out:

```
<table>
  <tr>
    <td><asp:Label id="labelName" runat="server">Name</asp:Label></td>
    <td><asp:TextBox id="txtUserName" runat="server"
      Width="332" Height="24"></asp:TextBox></td>
  </tr>
</table>
```

```

        <td><asp:Label id="labelAddr1" runat="server">Address</asp:Label></td>
        <td><asp:TextBox id="txtAddr1" runat="server"
            Width="332" Height="24"></asp:TextBox></td>
    </tr>
    <tr>
        <td><asp:Label id="labelAddr2" runat="server"></asp:Label></td>
        <td><asp:TextBox id="txtAddr2" runat="server"
            Width="332" Height="24"></asp:TextBox></td>
    </tr>
    <tr>
        <td><asp:Label id="labelCity" runat="server">City</asp:Label></td>
        <td>
            <asp:TextBox id="txtCity" runat="server"></asp:TextBox>
            <asp:Label id="labelState" runat="server">State</asp:Label>
            <asp:TextBox id="txtState" runat="server" Width="34" Height="24">
                </asp:TextBox>
            <asp:Label id="labelZIP" runat="server">ZIP</asp:Label>
            <asp:TextBox id="txtZIP" runat="server" Width="60" Height="24">
                </asp:TextBox>
        </td>
    </tr>
    <tr>
        <td><asp:Label id="labelEmail" runat="server">Email</asp:Label></td>
        <td><asp:TextBox id="txtEmail" runat="server"
            Width="332" Height="24"></asp:TextBox></td>
    </tr>
</table>

```

```

<script language="C#" runat="server" ID="Script1">
    public String UserName {
        get { return txtUserName.Text; }
        set { txtUserName.Text = value; }
    }
    public String Address1 {
        get { return txtAddr1.Text; }
        set { txtAddr1.Text = value; }
    }
    public String Address2 {
        get { return txtAddr2.Text; }
        set { txtAddr2.Text = value; }
    }
    public String City {
        get { return txtCity.Text; }
        set { txtCity.Text = value; }
    }
    public String State {
        get { return txtState.Text; }
        set { txtState.Text = value; }
    }
    public String ZIP {
        get { return txtZIP.Text; }
        set { txtZIP.Text = value; }
    }

```

```

    }
</script>

```

To use your pagelet, register it as a server control via the `@Register` directive, as shown in the next block of code. After registering, include the tag for the pagelet as if it was a normal server control. Specify the prefix, the tag name, the server control's ID, and set the `runat` property to `server`:

```

<%@ Register TagPrefix="Acme" TagName="Address" Src="Address.ascx" %>
<%@ Page language="c#" %>
<html>
<head>
    <script language="C#" runat="server">
        void Page_Load(Object oSender, EventArgs evt) {
            addr.UserName = "Jack Daniel";
        }
    </script>
</head>
<body>
    Welcome to the E-Shop.
    Registering with E-Shop will allow for monthly updates of bargains . . .
    <form method="post" runat="server">
        <p><Acme:Address id="addr" runat="server"></Acme:Address></p>
        <p><asp:Button id="cmdClear" runat="server" Text="Clear"></asp:Button>
            <asp:Button id="cmdSubmit" runat="server" Text="Submit">
                </asp:Button></p>
    </form>
</body>
</html>

```

You should be able to programmatically access the properties of the pagelet through the server control's ID (`addr` in this case). In the previous example, we accessed the `UserName` property of the `Address` pagelet via its ID:

```
addr.UserName = "Jack Daniel";
```

For an e-commerce checkout page, you could have two instances of `<Acme:Address>` on the same page: one for the billing and the other for the shipping address. Your script should access these instances of the pagelet via the ID you assign to each address control.

You can also programmatically instantiate instances of the pagelet through the use of the Page's `LoadControl` method. The first thing is to declare a variable of type `Control` in your script to host your pagelet. This is because the `Control` is the root of all objects, including your pagelet. Then instantiate the variable with a call to the `LoadControl`, passing in the filename of the control page. To make the control visible on the page, add the control to the Page's collection of controls. Because you currently have an instance of the `Control` object, you won't be able to call the pagelet's properties and methods until you *cast* the variable from `Control` type to your pagelet type. This is similar to having an `Object` variable in Visual Basic to hold a COM component. To access the COM-component methods and properties, you would cast the `Object` variable to the component type. Pagelets when loaded are automatically typed as `pagename_extension`. For example, if your pagelet were named `myControl.ascx`, the type generated for it would be `myControl_ascx`. The boldface line in the following example shows you how to cast `addr1` from `Control` to type `Address_ascx` in order to access the `UserName` property of the pagelet:

```

<%@ Register TagPrefix="Acme" TagName="Address" Src="Address.ascx" %>
<%@ Page language="C#" %>
<html>
<head>
  <script language="C#" runat="server">
    void Page_Load(Object oSender, EventArgs evt) {
      addr.UserName = "Jack Daniel";
      Control addr1;
      addr1 = LoadControl("Address.ascx");
      ((Address_ascx)addr1).UserName = addr.UserName;
      this.frm.Controls.AddAt(3, addr1);
    }
  </script>
</head>
<body>
  <form id="frm" method="post" runat="server">
    Billing Address:<br/>
    <Acme:Address id="addr" runat="server"></Acme:Address>
    Shipping Address:<br/>
    <p><asp:Button id="cmdClear" runat="server" Text="Clear">
      </asp:Button>
      <asp:Button id="cmdSubmit" runat="server" Text="Submit">
      </asp:Button>
    </p>
  </form>
</body>
</html>

```

This example, the checkout page, shows you how to declare a pagelet statically in your page with the `<Acme:Address>` tag, as well as how to dynamically create an instance of the custom control `Address` with the Page's `LoadControl()` method. Once you've created the control dynamically, you must cast the object to the control type before manipulating it.

The `AddAt()` method is used to insert the `Address` pagelet at a particular location in the checkout page. Instead of declaring the dynamic pagelet as a `Control`, you can also declare it as its type, which is `Address_ascx`. This way, you have to cast it only once when loading the dynamic control:

```

Address_ascx addr2 = (Address_ascx)LoadControl("Address.ascx");
addr2.UserName = "ABC";

```

7.5.4.2 Control derivatives

Although it is easy to create custom controls using the pagelet approach, this technique is not flexible enough to create more powerful custom controls, such as ones that expose events or hierarchy of controls. With ASP.NET, you can also create custom controls by inheriting from the `Control` base class and overriding a couple of methods.

The following example shows you how to create the simplest custom control as a `Control` derivative:

```

namespace MyWebControls
{
  using System;

```

```

using System.Web.UI;
using System.Web.UI.WebControls;
using System.ComponentModel;

public class MyWebControl : System.Web.UI.WebControls.WebControl
{
    //protected override void Render(HtmlTextWriter output)
    //{
    //    output.Write("custom control testing via Render( )");
    //}

    protected override void CreateChildControls( )
    {
        Table tbl = new Table( );
        TableRow row = new TableRow( );
        TableCell cell = new TableCell( );
        HyperLink a = new HyperLink( );
        a.NavigateUrl = "http://msdn.microsoft.com";
        a.ImageUrl = "image url";
        cell.Controls.Add (a);
        row.Cells.Add(cell);
        tbl.Rows.Add(row);

        row = new TableRow( );
        cell = new TableCell( );
        cell.Controls.Add (new LiteralControl("custom control testing"));
        row.Cells.Add(cell);
        tbl.Rows.Add(row);

        tbl.BorderWidth = 1;
        tbl.BorderStyle = BorderStyle.Ridge;

        Controls.Add(tbl);
    }
}

```

As you can see, the `MyWebControl` object derives from the `WebControl` class. We have seen that `WebControl` ultimately derives from the base `Control` class. All we really do here is override either the `Render` or the `CreateChildControls` methods to construct the custom web control. If you choose to override the `Render` method, you will have to generate the HTML for your custom control through the `HtmlTextWriter` object, `output`. You can use methods such as `Write`, `WriteBeginTag`, `WriteAttribute`, and `WriteEndTag`.

In our example, we override the `CreateChildControls` method. Instead of worrying about the actual HTML tag and attribute names, we then create ASP.NET objects directly by their class names, such as: `Table`, `TableRow`, `TableCell`, `HyperLink`, and `LiteralControl`, to construct a hierarchy of objects under a *table*. We can also manipulate attributes for the objects via their properties. At the end of the method, we add the *table* to the custom control's collection of controls.

You will have to compile the previous control code to generate a DLL assembly (i.e., `csc /t:library MyWebControls.cs`). To use the control, deploy the assembly by copying it to the `/bin` directory of

your web application. Then you should be able to register the control with the `@Register` directive and use it as if it was a server control provided by ASP.NET. If you are using Visual Studio .NET, you can add a reference to the control assembly file or the control project for the test web project that uses the control.

Your custom-control test page should now look like the following:

```
<%@ Page language="c#" %>
<%@ Register TagPrefix="WC" Namespace="MyWebControls"
           Assembly="MyWebControls" %>
<html>
<head>
  <script language="C#" runat=server>
    void Page_Load(object sender, EventArgs e) {
      MyWebControls.MyWebControl myCtrl;
      myCtrl = new MyWebControls.MyWebControl( );
      this.Controls.Add(myCtrl);
    }
  </script>
</head>
<body>
  <form method="post" runat="server">
    This is the main page
    <WC:MyWebControl id="myControl1" runat="server" />
  </form>
</body>
</html>
```

As you can see, we register the custom control with the `@Register` directive and alias the namespace `MyWebControls` with the `WC` prefix. In the body of the Web Form, we can add the custom-control tag as `<WC:MyWebControl>`.

In addition to inserting the custom control onto the page declaratively, as shown earlier, we can also programmatically create the custom control at runtime. The `Page_Load` code demonstrates this point:

```
MyWebControls.MyWebControl myCtrl;
myCtrl = new MyWebControls.MyWebControl( );
this.Controls.Add(myCtrl);
```

The output page is shown in [Figure 7-6](#).

Figure 7-6. Custom control test output, statically and dynamically

This is the main page



7.5.5 Event-Driven Programming

There are two ways to associate event handlers-functions that handle the event-to the UI controls.

Refer to [Section 7.4](#) earlier in this chapter, particularly where we describe the syntax for server controls. All we do to bind an event from a control to an event handler is use the *eventname=eventhandlername* attribute/value pair for the control. For example, if we want to handle the *onclick* event for the HTML control *input*, all we do is the following. Note that for the HTML controls, the server-side click event is named *onserverclick*, as opposed to the client-side click event, *onclick*, which can still be used in DHTML scripting:

```
<input id="cmd1" runat="server"
  onserverclick="OnClickHandler"
  type="button" value="click me">
```

For an ASP.NET web control, the syntax is the same:

```
<asp:Button id="cmd2" runat="server"
  onclick="OnClickHandler2"
  Text="click me too"></asp:Button>
```

After binding the event to the event-handling function name, we have to provide the actual event handler:

```
void OnClickHandler(object sender, EventArgs e)
{
  // Code to retrieve and process the posted data
}
```

The second way of binding events is delegation. You don't have to have any notion of code in the *aspx* file, not even the event-handling function name. All you have to do is register the event handler with the control's event-handler property. For web controls, the event-handler property for button click is *Click*. For HTML controls, it's *ServerClick*:

```
ControlID.Click += new System.EventHandler (this.EventHandlerName);
```

```
ControlID.ServerClick += new System.EventHandler (this.EventHandlerName);
```

7.5.6 Custom Server Controls and Client Script

Although ASP allows for the generating of dynamic pages from the server, the HTML generated usually needs the help of client-side scripting to provide a richer user experience. Again, in ASP.NET, the Server Controls are not meant to just generate static HTML. Client-side scripts can still be incorporated to further enhance the user interaction. There are two general way of doing this. The first is to include the client-side scripts within the body of the ASPX page, or somewhere the custom control can get to. The second way is to make the custom control emit its related client-side script while rendering. Either way, the client-side events will need to be tied to the script either through declarative means (i.e., *attributeEvent= eventHandler*) or programmatically through adding attributes to the rendering control dynamically. This is also where we can use the ClientId property of the Control object to allow client-side script to refer to the control on the client side.

We show two examples in this section to describe how you can do both. The pros and cons for each of these should be weighed technically as by other factors, such as time and usage of the control. For example, if you write the custom control for internal use in your company and the time-to-market is extremely short, you might opt to just include the client-side script manually because it might be easier to debug and change client or server-side code independently. On the other hand, if the custom control you are writing will be distributed widely, you might want to package it so that the control users do not have to worry about having to set up any external scripts. Enough said, let's see some examples. In this first example, we change the simple custom control shown earlier to include some invocation of client-side script and add the client script to the test page that uses the control:

```
namespace MyWebControls
{
    . . .

    public class MyWebControlWithClientScriptV1 :
        System.Web.UI.WebControls.WebControl
    {
        protected override void CreateChildControls( )
        {
            . . .
            cell.Controls.Add (new LiteralControl("custom control testing"));
            cell.Attributes.Add("onmouseover", "MouseOverHandler(this);");
            cell.Attributes.Add("onmouseout", "MouseOutHandler(this);");
            row.Cells.Add(cell);
            . . .
        }
    }
}
```

The changes from the custom control are highlighted. Basically, we just add two attributes "onmouseover" and "onmouseout" to the cell that contains the text "custom control testing" to call two client-side functions: MouseOverHandler and MouseOutHandler, respectively. What we need to do now is add these two client-side script functions to the custom control test page:

```
<%@ Page language="c#" Trace="true"%>
<%@ Register TagPrefix="WC" Namespace="MyWebControls"
           Assembly="MyWebControlWithClientScriptV1 "%>
<html>
<head>
```

```

<script language="javascript">
    function MouseOverHandler(ctl) {
        ctl.style.color="red";
    }
    function MouseOutHandler(ctl) {
        ctl.style.color="black";
    }
</script>
</head>
<body>
    <form id="myForm1" method="post" runat="server">
        <WC:MyWebControlWithClientScriptV1 id="myControl1" runat="server" />
    </form>
</body>
</html>

```

The two client-side script functions change the color of the text within the custom control when the mouse moves over and out of it.

As you can see, in order to use this version of the custom control, you have to know to include the client-side script functions to avoid errors. This is not ideal if you are writing custom control for a living. We move on to the second example where you don't have to do anything special to use the custom control.

Let's start with the test ASPX page:

```

<%@ Page language="c#" Trace="true"%>
<%@ Register TagPrefix="WC" Namespace="MyWebControls"
        Assembly="MyWebControlWithClientScriptV2 "%>
<html>
<head>
</head>
<body>
    <form id="myForm1" method="post" runat="server">
        <WC:MyWebControlWithClientScriptV2 id="myControl1" runat="server" />
    </form>
</body>
</html>

```

Notice that there is no client-side script block and the class name for the control is V2.

Now on the control side, we add the following block of code in the overridden CreateChildControls method:

```

        if(!Page.IsClientScriptBlockRegistered("myScript")) {
            string sScript =
@"
<!-- ctrl generated -->
<script language=""javascript"">
    function MouseOverHandler(ctl) {
        ctl.style.color=""red"";
    }
    function MouseOutHandler(ctl) {

```

```

        ctl.style.color="black";
    }
</script>
<!-- ctrl generated -->
";
        Page.RegisterClientScriptBlock("myScript", sScript);

```

Here, we basically ask the Page (that hosts the control) to see if "myScript" has been registered. If not, we register the included script with the name "myScript". When the control is rendered to the browser, the script will be rendered only once for this version of the control.

This example is just to demonstrate how client-side scripts can be inserted into the stream rendering to the browser. In practice, you might only render the script header block that points to an external script file that you've installed on a specific location on the web server. This improves performance in the long run by having the client-side script file cached by the browser. In case you really don't want anybody to change the client-side script that you have installed on the web server, you can include the client-side script file in your assembly as a resource. You can then load and register the script block from the DLL file.

7.5.7 ASP.NET and Caching

When talking about caching in ASP.NET, there are two kind of caching you should know. "Output caching" involves caching the generated html for part or the whole ASPX page so that it can be resent to the client without regenerating the cache content. The second type of caching in ASP.NET is application specific "Application data caching". Here you write code to cache data generated via expensive operations so that you don't have to regenerate the data for every request. Output caching is done through a declarative mean and it's automatic while Application data caching is programmable and flexible but not automatic.

Let's look at the following simple ASPX file:

```

<%@ Page language="c#" %>
<html>
<body>
    <form id="frml" runat="server">
        Page generated: <% DateTime t = DateTime.Now; Response.Write(t); %><br/>
    </form>
</body>
</html>

```

If you browse to this ASPX file, you will see the time the page was generated on the server. Every time you refresh your browser, the page is regenerated.

You can cache the content of this page by just adding the output cache directive:

```

<%@ OutputCache duration="30" varybyparam="none" %>

```

With this change, every time you refresh the browser within 30 seconds from the first request, the page does not get regenerated.

There are attributes associating with the OutputCache directive dealing with:

Duration

How long the item stays in the cache (in seconds).

Location

Where the cache is stored (for an ASPX file).

Shared

Whether the cached item can be shared with multiple pages (for an ASCX file).

The rest of the attributes specify how to uniquely cache the item based on differences by control IDs, by HTTP header variables, by QueryString variables or Form variables, and by custom management (*VaryByControl*, *VaryByCustom*, *VaryByHeader*, *VaryByParam*).

The example shows that we cache the page with no distinction on param. Of course, you can set it up so that the system will cache multiple versions of the page based on params.

Now that we know how to cache pages or controls on a page, let's take a look at how application content caching has changed from ASP. In ASP, developers have used Application variables or even Session variables as the cache mechanism. This works fine, but there are a number of considerations the developers have to worry about around issues such as memory resource and freshness of data. ASP.NET introduces the Cache object that is associated with each instance of ASP.NET application. You can add items into this Cache collection with priority, duration (absolute duration or relative to latest access), file or key dependencies (so that when the file, directory, or dependent cache item changes, the cache item is cleared. In addition, you can also have call-back function so you will know when the cache item is removed from the cache.

```
<%@ Page language="c#" %>
<script runat="server" language="c#">
    void Page_Load( ) {
        if(Cache["myItem"] == null) {
            Cache.Insert("myItem", // Cache name
                DateTime.Now, // Cache data
                null, // Cache dependency
                DateTime.Now.AddSeconds(30), // Absolute
                TimeSpan.Zero // Relative
            );
        }
        myLabel.Text = Cache["myItem"].ToString( );
    }
</script>
<html>
<body>
    <form id="frm1" runat="server">
        <asp:Label id="myLabel" runat="server"></asp:Label>
    </form>
</body>
</html>
```

The above code behaves similar to the output cache example we had earlier. The cache item, which is the current date time value is stored for 30 seconds. In the real world, you would probably cache something that would cost a little more than *DateTime.Now*, such as a *DataSet* that contains a number of rows or tables from a database.

For the second example, let's see how we can cause the cache to refresh the data using the dependency. We continue with the previous example and add the following code to Page_Load to load *myDoc.xml* into the cache and display this *xml* content in myLabel2. The cache dependency specifies that when *myDoc.xml* is changed, this cache item should be invalidated:

```
if(Cache["myItem2"] == null) {
    System.Xml.XmlDocument oDoc = new System.Xml.XmlDocument( );
    oDoc.Load(Server.MapPath("myDoc.xml"));
    Cache.Insert("myItem2",
                oDoc,
                new CacheDependency(Server.MapPath("myDoc.xml")));
    myLabel2.Text = "<br/>Refresh time: " + DateTime.Now.ToString( );
}
myLabel2.Text += "<xmp>"
                + ((System.Xml.XmlDocument)Cache["myItem2"]).InnerXml
                + "</xmp>";
```

We also have to add another label object inside the form tag:

```
<br/>
<asp:Label id="myLabel2" runat="server"></asp:Label>
```

Now if you navigate to this test page for the first time, the first label will be the date time on the server and the second label will contain "Refresh time" and the xml in *myDoc.xml*. Refreshing the page won't reload the xml file until this file is modified.^[6]

[6] This is merely an example to demonstrate the CacheDependency. Reading a file in a web application is dangerous, especially when there is no error handling code.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

7.6 ASP.NET and Web Services

The ASP.NET framework simplifies development of web services.^[7] All the low-level work, such as packaging and unpacking data in XML format and utilizing HTTP protocol to transport the web messages between distributed components, are done by the framework. This allows the developers to focus on the application logic.

^[7] We've seen web services in detail in [Chapter 6](#). This section reviews web services briefly to remind you that the underlying support for web services is ASP.NET.

The .NET Framework uses *asmx* as the default file extension for web services, as opposed to *aspx* for Web Forms and *ascx* for web controls.

7.6.1 The WebService Directive

All *asmx* files start with the `@WebService` directive that instructs ASP.NET on how to compile the code, as well as the main class name. The `WebService` directive has the following attributes:

Language

Specifies the language in which the code was written. This instructs the ASP.NET framework to use the appropriate compiler to build your web service. Use `VB` for Visual Basic, `C#` for C#, and `JS` for JScript .NET. As other languages emerge, obviously you can specify other languages.

Class

Specifies the main class, which exposes web methods. The ASP.NET framework instantiates this class in order to serve the web methods to the clients.

Codebehind

Specifies the source file for your code, which allows for complete code/ASP separation.

You can easily create a simple Web Service similar to the following *asmx* file:

```
<%@ WebService Language="VB" class="MathClass" %>
imports System.Web.Services
Public Class MathClass
  <WebMethod> _
  public function Add(a as integer, b as integer) as integer
    return(a + b)
  end function
end class
```

Note the line continuation symbol right after `<WebMethod>`. If you prefer to separate your code completely from any ASP.NET elements, you could have the code for your Web Service saved in a separate file and specify the `Codebehind` attribute of the `@WebService` directive to point to the code file:

```
<%@ WebService Language="VB" Codebehind="MathClass.vb" Class="MathClass" %>
```

The source for *MathClass.vb* looks exactly like the *asmx* shown earlier minus the first line. You can use the following command line to compile *MathClass.dll*.

```
vbc /t:library /r:System.Web.Services.dll MathClass.vb
```

As with all code-behind, the binary has to be deployed in the */bin* directory under the application.

7.6.2 The WebMethod Attribute

Public methods of any classes can be tagged with the WebMethod attribute to be made accessible from the Web. The syntax for tagging attributes to methods is different for each .NET language. For example, in C# the tag takes the following form:

```
[WebMethod(attribute="value" attribute="value" . . . )]
    public returnType FunctionName(paramsList)
```

In VB, angle brackets are used instead of square brackets and the assignment symbol is ":= " instead of just "=". Also note that the whole web method declaration is on a single line. If you want to separate them for readability, use the line continuation symbol "_":

```
<WebMethod(attribute:="value" attribute="value" . . . )> Public Function
    FunctionName(paramsList) as returnType
```

```
<WebMethod(attribute:="value" attribute="value" . . . )> Public Sub
    SubName(paramsList)
```

7.6.3 Using Web Services

If you are using Visual Studio .NET, you can choose Project/Add Web Reference and then type in the URL where the Web Service resides.^[8] For our purpose, we'll point to the Web Service we created in the last chapter, PubsWS. The URL to this Web Service on our server is *http://localhost/PubsWS/PubsWS.asmx*. The default web reference name is the server name where the Web Service was found. After adding the web reference, you can access the proxy object to the Web Service you are calling via the type *servername.proxyObjectName*. For your case, it is *localhost.PubsWS*.^[9]

[8] In VS.NET 1.1, you can also browse the web services on your local machine in addition to browsing web services on UDDI servers.

[9] You can rename the web reference when adding it to your project. This way the Web Service will be *<yourwebservicename>.proxyObjectName* instead of *servername.proxyObjectName*.

The following code excerpt demonstrates how to use the Web Service through the proxy. We create an instance of the proxy object and then ask it to relay the message to the real Web Service to get the list of authors. The result will be streamed back in XML format, which is reconstructed into a DataSet object. We then bind *DataGrid1*, which is just a DataGrid object that we have on the Web Form, to the default view of the first table of the DataSet. Finally, we ask for the actual binding to take place. The resulting page is the grid populated with rows from the Authors table of the Pubs sample database:


```
localhost.PubsWS ws = new localhost.PubsWS( );  
DataSet ds = ws.GetAuthors( );  
DataGrid1.DataSource = ds.Tables[0].DefaultView;  
DataGrid1.DataBind( );
```

Instead of using Visual Studio .NET to locate and automatically generate the proxy class, you can also use the information from the previous chapter to generate the source for the proxy class yourself. You can then include or compile the source into a DLL and add the DLL to the project as a reference. In any case, the end result is the same. Here is an example that links against the proxy we created in the previous chapter and fills a grid with data:

```
<%@ Page Language="C#" %>  
<%@ Import Namespace="System.Data" %>  
  
<!-- Link to the proxy generated by wsdl.exe -->  
<%@ Assembly Src="PubsWS.cs" %>  
  
<html>  
  <head>  
    <title>SOAP Client</title>  
  </head>  
  <body>  
  
    <!-- Make the SOAP call and fill the data grid. -->  
    <%  
      PubsWS ws = new PubsWS( );  
      DataSet ds = ws.GetAuthors( );  
      dg.DataSource = ds.Tables[0].DefaultView;  
      dg.DataBind( );  
    %>  
  
    <!-- Create a data grid. -->  
    <asp:DataGrid id="dg" runat="server"/>  
  
  </body>  
</html>
```

[[Team LiB](#)]

[\[Team LiB \]](#)

7.7 Data Binding and the Use of Templates

Although all web controls can be data bound, only DataGrid, DataList, and Repeater use a template to control the display of data items. In this section, we show you how to perform simple data binding with some common web controls and how to use an HTML template to provide fully customized data-bound controls.

In ASP.NET, data binding between a control and its data source simply means that the data is copied from the source and placed onto the control upon the call to the `DataBind()` method. This is different than data binding in Windows Forms, where a link is maintained between the control and the data source. We cover Windows Forms and data binding in the next chapter.

In its simplest form, data binding is the act of binding a control to a data source. Previously, data binding required that an ADO recordset be a data source, which is not too flexible. There was no way to bind, for example, an array or a collection of objects to a control. With ASP.NET, the whole data-binding business is revamped. The only requirement to make your data source bindable to web controls is that your data source implements the `System.Collections.ICollection` interface. In other words, a bindable data source is a collection of homogeneous objects that the web controls can obtain data items from.

Although it is possible to write your own data classes that implement the `ICollection` interface and bind them to web controls, numerous classes exist that do this for you, such as `Array`, `ArrayList`, `DataView`, `HashTable`, `Queue`, `SortedList`, and `Stack`. All you have to do is put your data in these forms, and you can bind your data to web controls.

Here is the simplest form of data binding. In the form file, all we have are two list boxes with `ids list0` and `list1`:

```
<asp:listbox id="list0" runat="server"></asp:listbox>
<asp:listbox id="list1" runat="server"></asp:listbox>
```

In the `Page_Load` event handler in the code-behind source file, we construct the data sources of type `Array`, which implement the `ICollection` interface we mentioned earlier, and then bind the list controls with the data sources:

```
int[] myArray0 = new int[7] { 1, 2, 3, 5, 7, 11, 13 };
string[] myArray1 = new string[7] {
    "Monday",
    "Tuesday",
    "Wednesday",
    "Thursday",
    "Friday",
    "Saturday",
    "Sunday"
};

list0.DataSource = myArray0;
list0.DataBind();
list1.DataSource = myArray1;
```

```
list1.DataBind( );
```

[Figure 7-7](#) shows the output of this page.

Figure 7-7. Data binding with data from arrays

Most of the time, we tend to bind data from data sources that come from a database. The next example pulls data from the Categories table of the familiar NorthWind database. We can still use the `asp:listbox` control, but this time, we specify the item's property we want for the text, as well as the value property of the list box. We did not have to do this for the previous example because the items that the list box binds to are of simple types (int and string). If we were to have an array of objects, we would have to specify the property we want to bind to `datavaluefield` and `datatextfield` the way we do in the following example:

```
<asp:listbox id=ListBox1 runat="server"
    datavaluefield="CategoryID"
    datatextfield="CategoryName">
</asp:listbox>
```

Again, in the code-behind source file, we have the code to construct the data source and to bind the data source to the list control. Note that because we are using ADO.NET to get the data from the database, we must have references to `System.Data` and `System.Data.OleDb` namespaces. The `DefaultView` property of class `Table` is of type `DataView`, which implements the `ICollection` interface:

```
System.Data.DataSet m_ds = new System.Data.DataSet( );
String sConn =
    "provider=SQLOLEDB;server=(local);database=NorthWind;Integrated Security=SSPI";
String sSQL =
    "select * from Categories";

System.Data.OleDb.OleDbDataAdapter da =
    new System.Data.OleDb.OleDbDataAdapter(sSQL, sConn);
da.Fill(m_ds, "Categories");

ListBox1.DataSource = m_ds.Tables["Categories"].DefaultView;
ListBox1.DataBind( );
```

[Figure 7-8](#) shows the output for this example.

Figure 7-8. Data binding with data from a database

7.7.1 DataGrid

The DataGrid control takes data binding a step further by allowing more than one property of the bound item to be displayed. This section's example shows you how to control the binding of data columns to the grid, as well as how to customize the look and feel of the DataGrid using style.

By default, the DataGrid automatically binds all columns of the data source in the order that they come from the database. Sometimes this is not the behavior you would want. To fully control what columns bind and in which order you want the binding to happen, switch off the `autogeneratecolumns` attribute of the DataGrid, and provide the `columns` property, as shown in the following sample:

```
<asp:DataGrid id=DataGrid1 runat="server"
    ForeColor="Black"
    autogeneratecolumns=false>

    <columns>
        <asp:boundcolumn datafield=CategoryID
            headertext="ID" readonly=True/>
        <asp:boundcolumn datafield=CategoryName
            headertext="Category" />
        <asp:boundcolumn datafield=Description
            headertext="Description" />
    </columns>

    <SelectedItemStyle bgcolor="#ffcc99" font-bold=True/>

    <AlternatingItemStyle BackColor="Gainsboro"/>

    <FooterStyle BackColor="Silver" ForeColor="White"/>

    <ItemStyle BackColor="White"/>

    <HeaderStyle BackColor="Navy" Font-Bold="True" ForeColor="White"/>

</asp:DataGrid>
```

[Figure 7-9](#) shows the result of this example.

Figure 7-9. DataGrid data binding

In addition to using `asp:boundcolumn` to bind a column of the DataGrid to a column of the data source, you can also use `asp:buttoncolumn` to insert a column with buttons that generate notifications. You can handle these notifications to perform predefined tasks such as selecting the item, removing it, and adding it to the shopping basket. You can also have `asp:hyperlinkcolumn` insert links to other pages in a column, `asp:editcommandcolumn` control editing of the selected row of data, or `asp:templatecolumn` customize the display of your column of data.

There are a number of styles that you use to control the visual formatting of your DataGrid control. The HeaderStyle and FooterStyle, as the names imply, control the style for the header and the footer of the DataGrid. The ItemStyle, AlternatingItemStyle, SelectedItemStyle, and EditItemStyle are used for each type of items in the list. The PagerStyle controls the visual appearance and layout of the paging interface.

The code-behind source file for binding data to the DataGrid is similar to that of the previous example. Basically, we bind the set of the DataSource property of the DataGrid to the DefaultView of the Categories table and perform the binding with the DataBind method:

```
DataGrid1.DataSource = m_ds.Tables["Categories"].DefaultView;
DataGrid1.DataBind( );
```

7.7.2 DataList

Unlike the DataGrid control, where the data binding is still in a tabular form, the DataList control allows to you lay out the list in any way^[10] through the use of HTML templates.

[10] For both the DataList and DataGrid controls, the table html tag is used to lay out the output. The difference is that the DataGrid actually uses the table to display tabular data while the DataList uses the table for the lay out only and flows the data into the rows in the table.

Within a DataList tag, you can customize a number of templates. The templates that can be customized include:

- AlternatingItemTemplate
- EditItemTemplate
- FooterTemplate
- HeaderTemplate
- ItemTemplate
- SelectedItemTemplate
- SeparatorTemplate

Specific tags are used to set up the style for each type of items you want to display. Similar to the previous list, you also have ItemStyle, SelectedItemStyle, and so on.

In the following example, we only show you one template, the ItemTemplate, which is applied to all items in the list. In this template, we use Web Form data-binding syntax to bind two properties of the data item, the CategoryID and CategoryName fields. In this simple template, the CategoryID will always be shown with Verdana font in size 10.

You can also control the flow of the DataList by setting attributes such as `repeatcolumns`, `repeatdirection` (vertical, horizontal), or `repeatlayout` (flow, table):

```
<asp:DataList id=DataList1 runat="server"
    repeatcolumns=3
```

```

repeatdirection=Horizontal>

<AlternatingItemStyle BackColor="Gainsboro"/>
<ItemTemplate>
  <font face=Verdana size=10>
    <%=# DataBinder.Eval(Container.DataItem, "CategoryID") %>
  </font>
  <%=# DataBinder.Eval(Container.DataItem, "CategoryName") %>
</ItemTemplate>

</asp:DataList>

```

The code behind the data binding is shown here:

```

DataList1.DataSource = m_ds.Tables["Categories"].DefaultView;
DataList1.DataBind( );

```

[Figure 7-10](#) shows the output of this DataList data-binding example.

Figure 7-10. DataList data binding with template

1 Beverages	2 Condiments	3 Confections
4 Dairy Products	5 Grains/Cereals	6 Meat/Poultry
7 Produce	8 Seafood	

7.7.3 Repeater

The ASP.NET Repeater control is completely driven by an HTML template to repeatedly display each of the data items bound to it. When the page renders, the Repeater control loops through all the records in the data source and generates HTML to display the record according to the HTML template. This is as free-form as you can get for data binding. You can have templates to generate bulleted lists, numbered lists, comma-separated lists, and tabs. Unlike the DataGrid or DataList that use an HTML table to control the layout of the data, the Repeater control does not add any HTML of its own. You can specify the layout however you want using the templates described below.

There are only five templates in the Repeater control:

- AlternatingItemTemplate
- FooterTemplate
- HeaderTemplate
- ItemTemplate
- SeparatorTemplate

We will use two of these templates to control the display of the item and its separator.

Again, we bind two fields of each item to the template. The end result is a comma-separated list of URLs that link another Web Form to display more detailed information about the clicked category. As you can see, we also use Web Form data-binding tags, `<%#` and `%>`, to perform the binding. The CategoryID fills the `cat` parameter of the query string to the fictitious *DisplayCategory.aspx* Web Form, and the CategoryName is the display text for the anchor tag.

You could also replace the anchor tag and comma with graphical images to make your page more visually appealing:

```
<asp:Repeater id=Repeater1 runat="server">

    <ItemTemplate>
        <A HREF="http://YourURL/DisplayCategory.aspx?cat=
        <%# DataBinder.Eval(Container.DataItem, "CategoryID") %>"
        ><%# DataBinder.Eval(Container.DataItem, "CategoryName") %>
        </A>
    </ItemTemplate>
    <SeparatorTemplate>, </SeparatorTemplate>

</asp:Repeater>
```

[Figure 7-11](#) shows the result of using the data repeater to bind data.

Figure 7-11. Data binding using repeater and template

As with the other controls, the Repeater needs to be bound to a data source:

```
Repeater1.DataSource = m_ds.Tables["Categories"].DefaultView;
Repeater1.DataBind();
```

As you can see, using a template to bind data to these list-bound controls can be very simple, yet powerful. However, you should be aware of how the generated HTML will look. You should not have complicated, bloated templates that will result in unappealing, large files. In web application development, the page size is directly proportional to the response time the customer experiences.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

7.8 State Management and Scalability

ASP.NET overcomes all major limitations of ASP when it comes to managing session states. As you are aware from ASP development, a session state is nothing but a named variable that is cached at the server for the duration of the web user's session. As the user navigates through the web application, the session state retains its value as long as the session is not expired.

ASP session-state management can be summarized as follows:

- The session starts, and the web application assigns a unique key to the user.
- This key is stored in an HTTP cookie. Along each subsequent request, the client browser sends the unique key back to the server.
- The server looks up the states stored for this particular key and processes the request accordingly.

Although this has worked fine for all these years, we've found out that there were a number of limitations to live with or work around. The biggest limitation is that the session state is process-dependent, which is impossible to implement in a web farm environment without custom session management.

7.8.1 ASP.NET Session-State Management

ASP.NET improves on ASP session-state management by giving you the option to move to an out-of-process model. By having all web servers in the farm pointing to a common server that hosts the out-of-process state manager, the web client can be redirected around the farm without losing the session states.

By using an out-of-process model, we no longer have the problem of losing session states when the IIS process is cycled. This means that if the web server application crashed for whatever reason and restarted within the session timeout duration, the web clients could still have all their session states intact. Of course, if the out-of-process state manager crashed, that is a whole different issue. This leads to the next improvement of ASP.NET—the ability to persist session state to a database.

The idea of persisting session state to a database is not new. Many of us have implemented this as the workaround for dealing with web farm configuration. However, ASP.NET makes it easier.

Similar to all other configurations in ASP.NET, session management is done through the use of the *web.config* files. There are two levels of configuration: machine and application. Machine-level configuration associates with the *machine.config* file stored in *WinNT\Microsoft.NET\Framework\<version>\CONFIG\machine.config*, while the application-level configuration uses the *web.config* file in the application root directory. The application-level configuration overrides the machine-level configuration.

The following code is a portion of the *web.config* file dealing with session-state management:[\[11\]](#)

[11] The content of this file is case-sensitive.

```
<configuration>
  <system.web>
    <sessionState
      mode="InProc"
      cookieless="false"
      timeout="20" />
    </system.web>
  </configuration>
```

[Table 7-2](#) lists the properties of the SessionState class.

Table 7-2. Properties of the SessionState class

Property	Description
mode	<code>Off</code> indicates that session state is disabled; <code>InProc</code> stores session data locally; <code>StateServer</code> stores session state on a remote server; and <code>SQLServer</code> stores it on a SQL Server.
Cookieless	Specifies whether to rely on the client acceptance of cookie. If this property is set to <code>true</code> , ASP.NET inserts the unique key to the URL for navigation between pages within the application instead of setting it in the client's cookie.
Timeout	Specifies session timeout in minutes. This is a sliding window of time: it starts counting down for each request. The default is 20 minutes.
stateConnectionString	Specifies the server and port of the remote session-state server (not a SQL Server). The format is <code>tcpip=HOST:PORT</code> , as in <code>tcpip=192.168.254.1:42424</code> . Use this only when <code>mode=StateServer</code> .
sqlConnectionString	Represents a SQL Server connection string, such as <code>user id=sa;password=;database=ASPState;server=(local)</code> . This is required when <code>mode=SQLServer</code> .

7.8.1.1 Out-of-process session-state management

When you set the session-state mode to run on a remote server (`mode=StateServer`), you must prepare the remote server to run the state management service automatically.

ASP.NET SDK includes an NT service call *ASP.NET State Service* to be used for out-of-process session-state management. Before setting your *web.config* files to use the out-of-process mode, you will have to start the ASP State service by going to the NT Services Management Console and start the service. You might want to change the *startup type* to `automatic` so that this service will start automatically at subsequent reboots.

7.8.1.2 SQL Server session-state management

To use this mode, the SQL Server machine has to be prepared. ASP.NET SDK includes a SQL script to create the ASP State database, which is where all session states are stored. Find this SQL script (*InstallSqlState.sql*) at `%SystemRoot%\Microsoft.NET\Framework\BUILDNUMBER\`. To apply the script to your SQL Server, use the SQL Server command-line tool *osql.exe* or SQL Query Analyzer. We use the latter because it allows us to inspect the script to get a better understanding of how this mode of session management is implemented. You will have to stop and restart SQL Server because the script alters the `master` to run the `ASPState_Startup` helper procedure at SQL startup time.

7.8.1.3 Cookieless session-state management

In ASP development, it is a usual practice to impose the requirement that the clients' web browsers be set up to accept cookies so that we can use session state the way it is meant to be used. However, when this requirement is not in place, especially for business-to-consumer (B2C) kinds of applications, the developers have to package the session ID along with the URL as a variable in the query string or as a form field and manage the session states manually.

With ASP.NET, as you can see from the `sessionstate` section of the configuration file, all you do is flip the setting of `cookieless` to `true`, and everything is automatically done for you. Session state can be used as if nothing has changed.

To setup and experiment with these session-state configuration, we've created two fictitious *asp.net* pages: `login.aspx` and `main.aspx`. The main page redirects the user to the login page if the user has not logged in. The login page redirects the user to the main page when the user is authenticated. When the user logs in, session variable `UserName` will be populated.

The following is the source for the simplified login page:

```
<HTML>

<script language="VB" runat="server">
Sub cmdLogin_Click(ByVal sender As System.Object, _
                  ByVal e As System.EventArgs)
    ' more processing here
    Session("UserName") = txtUID.Text
    Response.Redirect("Main.aspx")
End Sub
</script>

<body>
<form id="Form1" method="post" runat="server">
<table>
  <tr>
    <td>User ID</td>
    <td><asp:TextBox id="txtUID"
                    runat="server"></asp:TextBox></td>
  </tr>
  <tr>
    <td>Password</td>
```

```

        <td><asp:TextBox id="txtPWD"
                    textmode="password"
                    runat="server">
                    </asp:TextBox></td>

    </tr>
<tr>
    <td></td>
    <td><asp:Button id="cmdLogin"
                    runat="server"
                    Text="Login"
                    onclick="cmdLogin_Click">
                    </asp:Button></td>

    </tr>
</table>
</form>
</body>
</HTML>

```

The skeleton for the main page is as follows:

```

<HTML>

<script language="VB" runat="server">
Sub Page_Load(ByVal sender As System.Object, ByVal e As System.EventArgs)
    If (Session("UserName") <> "") Then
        labelData.Text = "Welcome back, " + Session("UserName")
    Else
        Response.Redirect("Login.aspx")
    End If
End Sub
</script>

<body>
<form id="Form1" method="post" runat="server">
    <asp:Label id="labelData" runat="server"></asp:Label>
</form>
</body>
</HTML>

```

In the first scenario, we will use session-state mode `InProc`. Because the IIS process handles the session state, if we simulate a web server restart by issuing the command `iisreset` and trying to refresh the main page, it will redirect us to the login page.

In the second scenario, we change the session-state mode to `StateServer` and start the ASP.NET Session State Service (i.e., the command line `net start aspnet_state`). Note that here we are running the Session State Service on the same machine as the web server even though we can have this service running on a separate server for more reliability. This time around, the session state persists through the resetting of the web server. Of course, if we restart the ASP.NET Session State Service itself, the main page will still redirect us to the login page.

Now that we've seen in-process and out-of-process session-state management, the last scenario we try will be to have session state persisted to a database. This is as simple as setting the `mode` and the

`sqlConnectionString` attributes of the `sessionState` node in the `web.config` file. Of course, we ran `InstallSqlState.sql` on the SQL server to generate the schema and supporting stored procedures needed by ASP.NET to persist state into the database. The result is similar to the previous trials, however. Because the session data are stored in `tempdb`, they are cleared when the SQL server is restarted. As a side note, remember to have SQL Server Agent start automatically so that the cleanup session-state job can be run correctly.

7.8.2 Performance Versus Scalability and Reliability

As we've said, ASP.NET introduces an out-of-process model of session-state management, which enables more scalable solutions, but not without a cost. Out-of-process communication performs much worse than in-process communication, not to mention persisting the session states to a database. You should weigh the benefits of each of the different modes of state management to find the one that is most suitable for your application. [Table 7-3](#) summarizes the different modes and their trade-offs.

Table 7-3. Session-state management communication modes

Mode	Description
In-process	This mode gives you the best performance. It is not reliable, because it is memory-based. It is not scalable, because it is process-based. If you are setting up a web farm, you will have to make sure that subsequent requests are going to the same server.
Out-of-process	The reliability factor is still in question because this mode is still memory based. However, because a separate process manages the session state, it is more reliable than the in-process mode. Because of the out-of-process communication overhead, it is much slower than in-process mode. It is scalable for use in web farms.
SQL Server	This mode gives you the highest level of reliability at the cost of performance. It is scalable for use in web farms.

[[Team LiB](#)]

[Team LiB]

7.9 Forms Authentication in ASP.NET

The previous example in the session management section only demonstrates how session state can be managed in ASP.NET. If you want to expand the example to handle your application authentication, every single *aspx* file in addition to *main.aspx* should check for the session variable "UserName" and redirect to the *Login.aspx* file if this session variable is not found. This is too much work, at least in the .NET world. We take this opportunity to show you how to do forms authentication in ASP.NET. By definition, forms authentication is basically a setup where unauthenticated requests are automatically redirected to a designated login form. Once the user provides the login information and the login form processes it successfully, the user is then redirected back to the original page along with an "authenticated cookie." Subsequent requests do not get redirected to the login form until the cookie expires.

The first thing you will have to do is edit the `web.config` file to set the authentication mode to "Forms" and setup the URL for the login page and the name of the authentication cookie:

```
<configuration>
  <system.web>
    <authentication mode="Forms">
      <forms loginUrl="login.aspx" name=".authToken"/>
    </authentication>
    <authorization>
      <deny users="?" />
    </authorization>
  </system.web>
</configuration>
```

In this `web.config` file, we've specified that the authentication mode is "Forms" for form-based authentication. Other settings are "Windows," "Passport," and "None" (case-sensitive), which we will not cover in this book because of its size. The `loginUrl` is where the system should redirect the request if the user is not yet authenticated and the `name` attribute is the name of the cookie to store the authentication info. We also have to set up the authorization so that this web application will deny all unauthenticated users.

Since we specify that the login URL is *login.aspx*, let's see the content of this file:

```
<HTML>
<script language="VB" runat="server">
Sub cmdLogin_Click(ByVal sender As System.Object, _
                  ByVal e As System.EventArgs)
  ' more processing here
  FormsAuthentication.RedirectFromLoginPage(txtUID.Text, false)
End Sub
</script>
<body>
  . . .
</body>
</HTML>
```

Once we have authenticated the credentials, we call a helper method of FormsAuthentication object to redirect to whatever page the client was from. The first parameter is the user name and the second Boolean variable tells the function not to persist the cookie across browser sessions. Note the difference between this and the home-grown authentication via the session example we had earlier. Here, we don't have to remember what URL to return to.

The *main.aspx* page now looks like this:

```
<HTML>
<body>
<script language="VB" runat="server">
Sub Page_Load(ByVal sender As System.Object, ByVal e As System.EventArgs)
    labelData.Text = "Welcome back, " + Context.User.Identity.Name
End Sub
Sub Logout(ByVal sender As System.Object, ByVal e As System.EventArgs)
    FormsAuthentication.Signout( )
    Response.Redirect("Login.aspx")
End Sub
</script>
<form id="Form1" method="post" runat="server">
    <asp:Label id="labelData" runat="server">Label</asp:Label>
    <asp:Button id="cmdLogout" runat="server" onclick="Logout" Text="Logout"></asp:Button>
</form>
</body>
</HTML>
```

[Team LiB]

[\[Team LiB \]](#)

7.10 Summary

Throughout this chapter, we've introduced you to ASP.NET and the benefits that it brings to web application development. These benefits include a new and extended web page life cycle that involve events driven from the client browsers, server controls that manage their own states, the separation of user interface and the code behind, the replacement of late-bound scripting languages with strong typed compiled languages, and the new and improved session-state management that improves scalability.

If you are trying to embrace the web paradigm by using ASP.NET for your web application, along with web services for integration between sites, you are right on target. However, not all applications are suitable for the Web. There is a huge market of standard applications where the development is done in traditional VB, C, C++, Java, and so on, and there is no need for it to be web-based. In the next chapter, we cover Windows Forms, which map to traditional Windows applications.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Chapter 8. Windows Forms

If the goal of Microsoft .NET is to embrace the Web, what will happen to conventional Windows applications? It turns out the .NET Framework not only benefits the development of web applications, but improves the way standard Windows applications are built. In this chapter, we provide you with an understanding of what Windows Forms are, how to use Windows Forms .NET classes to create Windows Forms-based applications, and how you can still "embrace the Web" while creating Windows applications. [\[1\]](#)

[1] For further information, see O'Reilly's *.NET Windows Forms in a Nutshell*, by Ian Griffiths and Matthew Adams, and *Programming .NET Windows Applications*, by Jesse Liberty.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

8.1 Introducing Windows Forms

If you have developed Windows applications since the early 1990s, chances are you have used raw Windows APIs such as RegisterClass, CreateWindow, ShowWindow, GetMessage, TranslateMessage, and DispatchMessage. You certainly had a WinMain entry point in your application. Inside this function, you registered your application with Windows, created and showed the window, and handled messages from the system. Every Windows application has to have a message loop that collects Windows messages and dispatches them to the message-handler function that you've registered through RegisterClass function. As a developer, much of your job is handling Windows messages, such as WM_CREATE, WM_SIZE, or WM_CLOSE, that you create and pump into the system with PostMessage or SendMessage.

Classic Windows development is tedious and error-prone. The result is that application frameworks were built as an abstraction on top of all these Windows APIs. Frameworks such as the Microsoft Foundation Class Library (MFC) and Active Template Library (ATL) were created to help Windows application developers focus more on the task of solving business problems than on how to handle certain Windows messages. These frameworks provide the plumbing, or the template, of a Windows application. The developer's responsibility is to deal with business logic.

Although it is much easier to develop Windows applications using these frameworks, it is again sometimes necessary to go down to the Windows API level when the Framework does not give you the controls you need. This situation causes inconsistency in the code. Moreover, there exist numerous frameworks similar to MFC and ATL, such as the Object Windows Library (OWL) from Borland, zApp from Rogue Wave, Windows add-on scripts for Python such as the Win32 Extensions or PythonWin GUI Extensions, Visual Basic, and other homegrown frameworks, causing developers much grief when switching from one to another.

Windows Forms provides a unified programming model for standard Windows application development. It is similar to the native Windows API with regard to level of abstraction; however, it is much richer and more powerful. Instead of depending on functions such as the native Windows API, Windows Forms provides a hierarchy of classes. Instead of calling CreateWindow for any type of user-interface widgets, you create the particular type of user-interface control using the appropriate class. You might think that MFC and other frameworks already provide hierarchy of classes. What other benefits can Windows Forms bring that make it stand out from the crowd? The answer is the language-independent aspect of this new framework. Any .NET language can use this collection of classes that make up the Windows Forms object model.

If you've developed Windows applications in C++ and Visual Basic, you might think that it would be nice to have the power of C++ to work in an integrated development environment like that of VB. It is now possible with Visual Studio .NET and Windows Forms. Windows Forms brings a VB-like integrated development environment to C#, Managed C++, and other languages.

In current Windows application development, if you use COM, DCOM, or ActiveX components, deployment of your application requires extensive configuration. You would probably at least use the *regsvr32* utility to register and unregister components from the Windows Registry on the client machine. All these setup-related deployment tasks are eliminated by Microsoft .NET-by Windows Forms in particular. Now, all you have to do to install an application is copy the executable onto the

client machine.

Because Windows Forms is part of the Microsoft .NET grand scheme, it fully supports and integrates with web services, ADO.NET, and the .NET classes. You can have Windows Forms as the frontend to your web application by using .NET classes such as `HttpRequest` and `HttpResponse`. These classes allow your Windows Forms application to communicate with web servers. You can also use web services from your Windows Forms application. Remember that Windows Forms applications are no longer just stand-alone applications.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

8.2 The System.Windows.Forms Namespace

In this section, we describe the architecture of Windows Forms and introduce the classes that make up the Windows Forms namespace.

Windows Forms architecture is rather simple. It takes the form of *controls* and *containers*. This is similar to Java JFC model, where container types of classes are Panel, Window, JComponent, and so on, and control types of classes are Button, Checkbox, Label, and so on. Most user-interface classes in the Windows.Forms namespace derive from the Control class. In a sense, everything that you see in a Windows Forms application is a control. If a control can contain other controls, it is a container. The application user interface consists of a form object acting as the main container, as well as the controls and other containers that reside on the form.

Similar to the native Windows API common functions, the System.Windows.Forms namespace provides a common set of classes you can use and derive from to build Windows Forms applications. The classes and interfaces in this namespace allow you to construct and render the user-interface elements on a Windows Form.

As we have seen from the last chapter, the System.Web.UI namespace provides the classes for building web applications. Similarly, the System.Windows.Forms namespace provides the classes for building standard applications. The System.Windows.Forms namespace is analogous to the System.Web.UI namespace, as described in the previous chapter.

Similar to the Control and Page classes in the System.Web.UI namespace, Control and Form are the two most important classes in the System.Windows.Forms namespace.

8.2.1 Control Class

Control is the base class of all UI controls in Windows Forms applications. It provides common properties for all controls, as well as common user-interface control behaviors, such as accepting user input through the keyboard or mouse and raising appropriate events.

[Table 8-1](#) is a list of some representative properties, methods, and events that you will most likely encounter. For the complete list, check out the Microsoft .NET SDK.

Table 8-1. Common Control properties, methods, and events

Properties	Description
Controls, Parent	These properties allow for constructing hierarchy of controls. The Controls property lists all child controls, while the Parent property points to the parent of the current control.

Properties	Description
Enabled, Focused, Visible	These properties control the visual states of the control.
Left, Right, Top, Bottom, Width, Height, Size, Location	These properties control the location and size of the control.
Methods	Description
Show, Hide, Focus, Select	These methods manipulate the control's visual state.
Refresh, Invalidate, Update	These methods control when and what portion of the screen needs redrawing. The Refresh method immediately forces the control to redraw itself and all of its children. The Invalidate and Update methods selectively control the portion of the screen that needs to be redrawn.
ProcessCmdKey, WndProc	If you develop your own controls, override these methods to intercept the Windows messages. This is similar to how Windows developers handled Windows messages when they developed Win32 applications using the native Win32 API.
Events	Description
Click, MouseDown, MouseUp, MouseMove, MouseWheel	You have to write event handlers to handle these events for your applications.
KeyDown, KeyUp, KeyPress	Similar to the mouse events, these keyboard-events can be handled via custom event handlers.


The Control class also provides behaviors, such as data binding, context menu, drag and drop, anchoring and docking; and properties, such as font, color, background, cursor, and so on.

8.2.2 Form Class

A *form* in Windows Forms is similar in concept to a *page* in Web Forms. It is a container type of control that hosts other UI controls. You manipulate the properties of the Form object to control the appearance, size, and color of the displayed form. A Windows Form is basically a representation of any window displayed in your application.

A standard form contains a titlebar, which contains an icon, title text, and control box for the Minimize, Maximize, and Close buttons (see [Figure 8-1](#)). Most of the time, a form also contains a menu right under the titlebar. The working area of the form is where child controls are rendered. A border around the whole form shows you the boundary of the form and allows for resizing of the form. Sometimes, the form also contains scrollbars so that it can display more controls or larger controls than the size of the working area of the form.

Figure 8-1. An empty application



You can manipulate the form's standard visual elements with properties such as `Icon`, `Text`, `ControlBox`, `MinimizeBox`, `MaximizeBox`, and `FormBorderStyle`. For example, if you want the title text of the form to read `Hello World`, you include the assignment `formName.Text = "Hello World";`. To have a form without the control box in the top right corner, set the `ControlBox` property to `false`. If you want to selectively hide the Maximize or the Minimize button in the control box, set the `MaximizeBox` or `MinimizeBox` property to `false`.

You can assign a menu to your form by setting the `Menu` property of the form with an instance of the `MainMenu` class. We will show you how to do this in [Section 8.3](#) later in this chapter.

Similar to Submit and Reset buttons in a web page's form, a form will frequently include OK and Cancel buttons to submit or to reset the form. In Windows Forms, you can assign any button to the `AcceptButton` property of the form to make it the default button when the user hits the Enter key. Similarly, you can set up the `CancelButton` property to handle the Escape key.

The Form class supports a number of methods itself, along with the methods it inherits from the base class. `Activate`, `Show`, `Hide`, `ShowDialog`, and `Close` are a few of the imperative methods used in any form to control the window-management functionality of a form. As we get into [Section 8.3](#) later in this chapter, you will see these methods in action.

8.2.2.1 Extending existing controls

Because Windows Forms API is object oriented, extending controls is as easy as deriving from the control you want to extend and adding methods, properties, and events, or overriding the default behavior of the control:

```
class MyCustomTextBox : TextBox
{
    // Customization goes here.
}
```

8.2.2.2 Creating composite controls

Composite controls are controls that contain other controls. By definition, it ought to be derived from the `ContainerControl` class; however, the Windows Forms object model provides the `UserControl` class, which is a better starting point for your custom composite controls (`UserControl` actually derives from `ContainerControl`):

```
class MyCustomComposite : UserControl
```

```
{
    // Composite controls go here.
}
```

While deriving from `UserControl` class to create your custom composite controls is not a hard task, Microsoft Visual Studio .NET is an excellent tool for making this task even easier. It truly is an effort to raise the bar on RAD tools. Developers' productivity benefits greatly from support tools like these.

8.2.3 Application Class

The `Application` class provides static methods to start, stop, or filter Windows messages in an application. All Windows Forms applications contain a reference to this `Application` class. More specifically, all Windows Forms applications start with something like the following:

```
System.Windows.Forms.Application.Run(new MyForm( ));
```

While this class provides other methods and properties beside the `Run` method, this method is really the only essential one. The rest of the methods (listed in the rest of this section) are low-level and not frequently used.

The `Run` method starts the application thread's message loop. This method has two signatures. The first signature involves no parameters, which are normally used for non-GUI applications:

```
System.Windows.Forms.Application.Run( );
```

The second signature takes a form as a parameter, as you can see from the first example. The form `MyForm` is the entry point to a GUI Windows Forms application.

[Table 8-2](#) summarizes the `Application` class.

Table 8-2. Common `Application` properties and methods

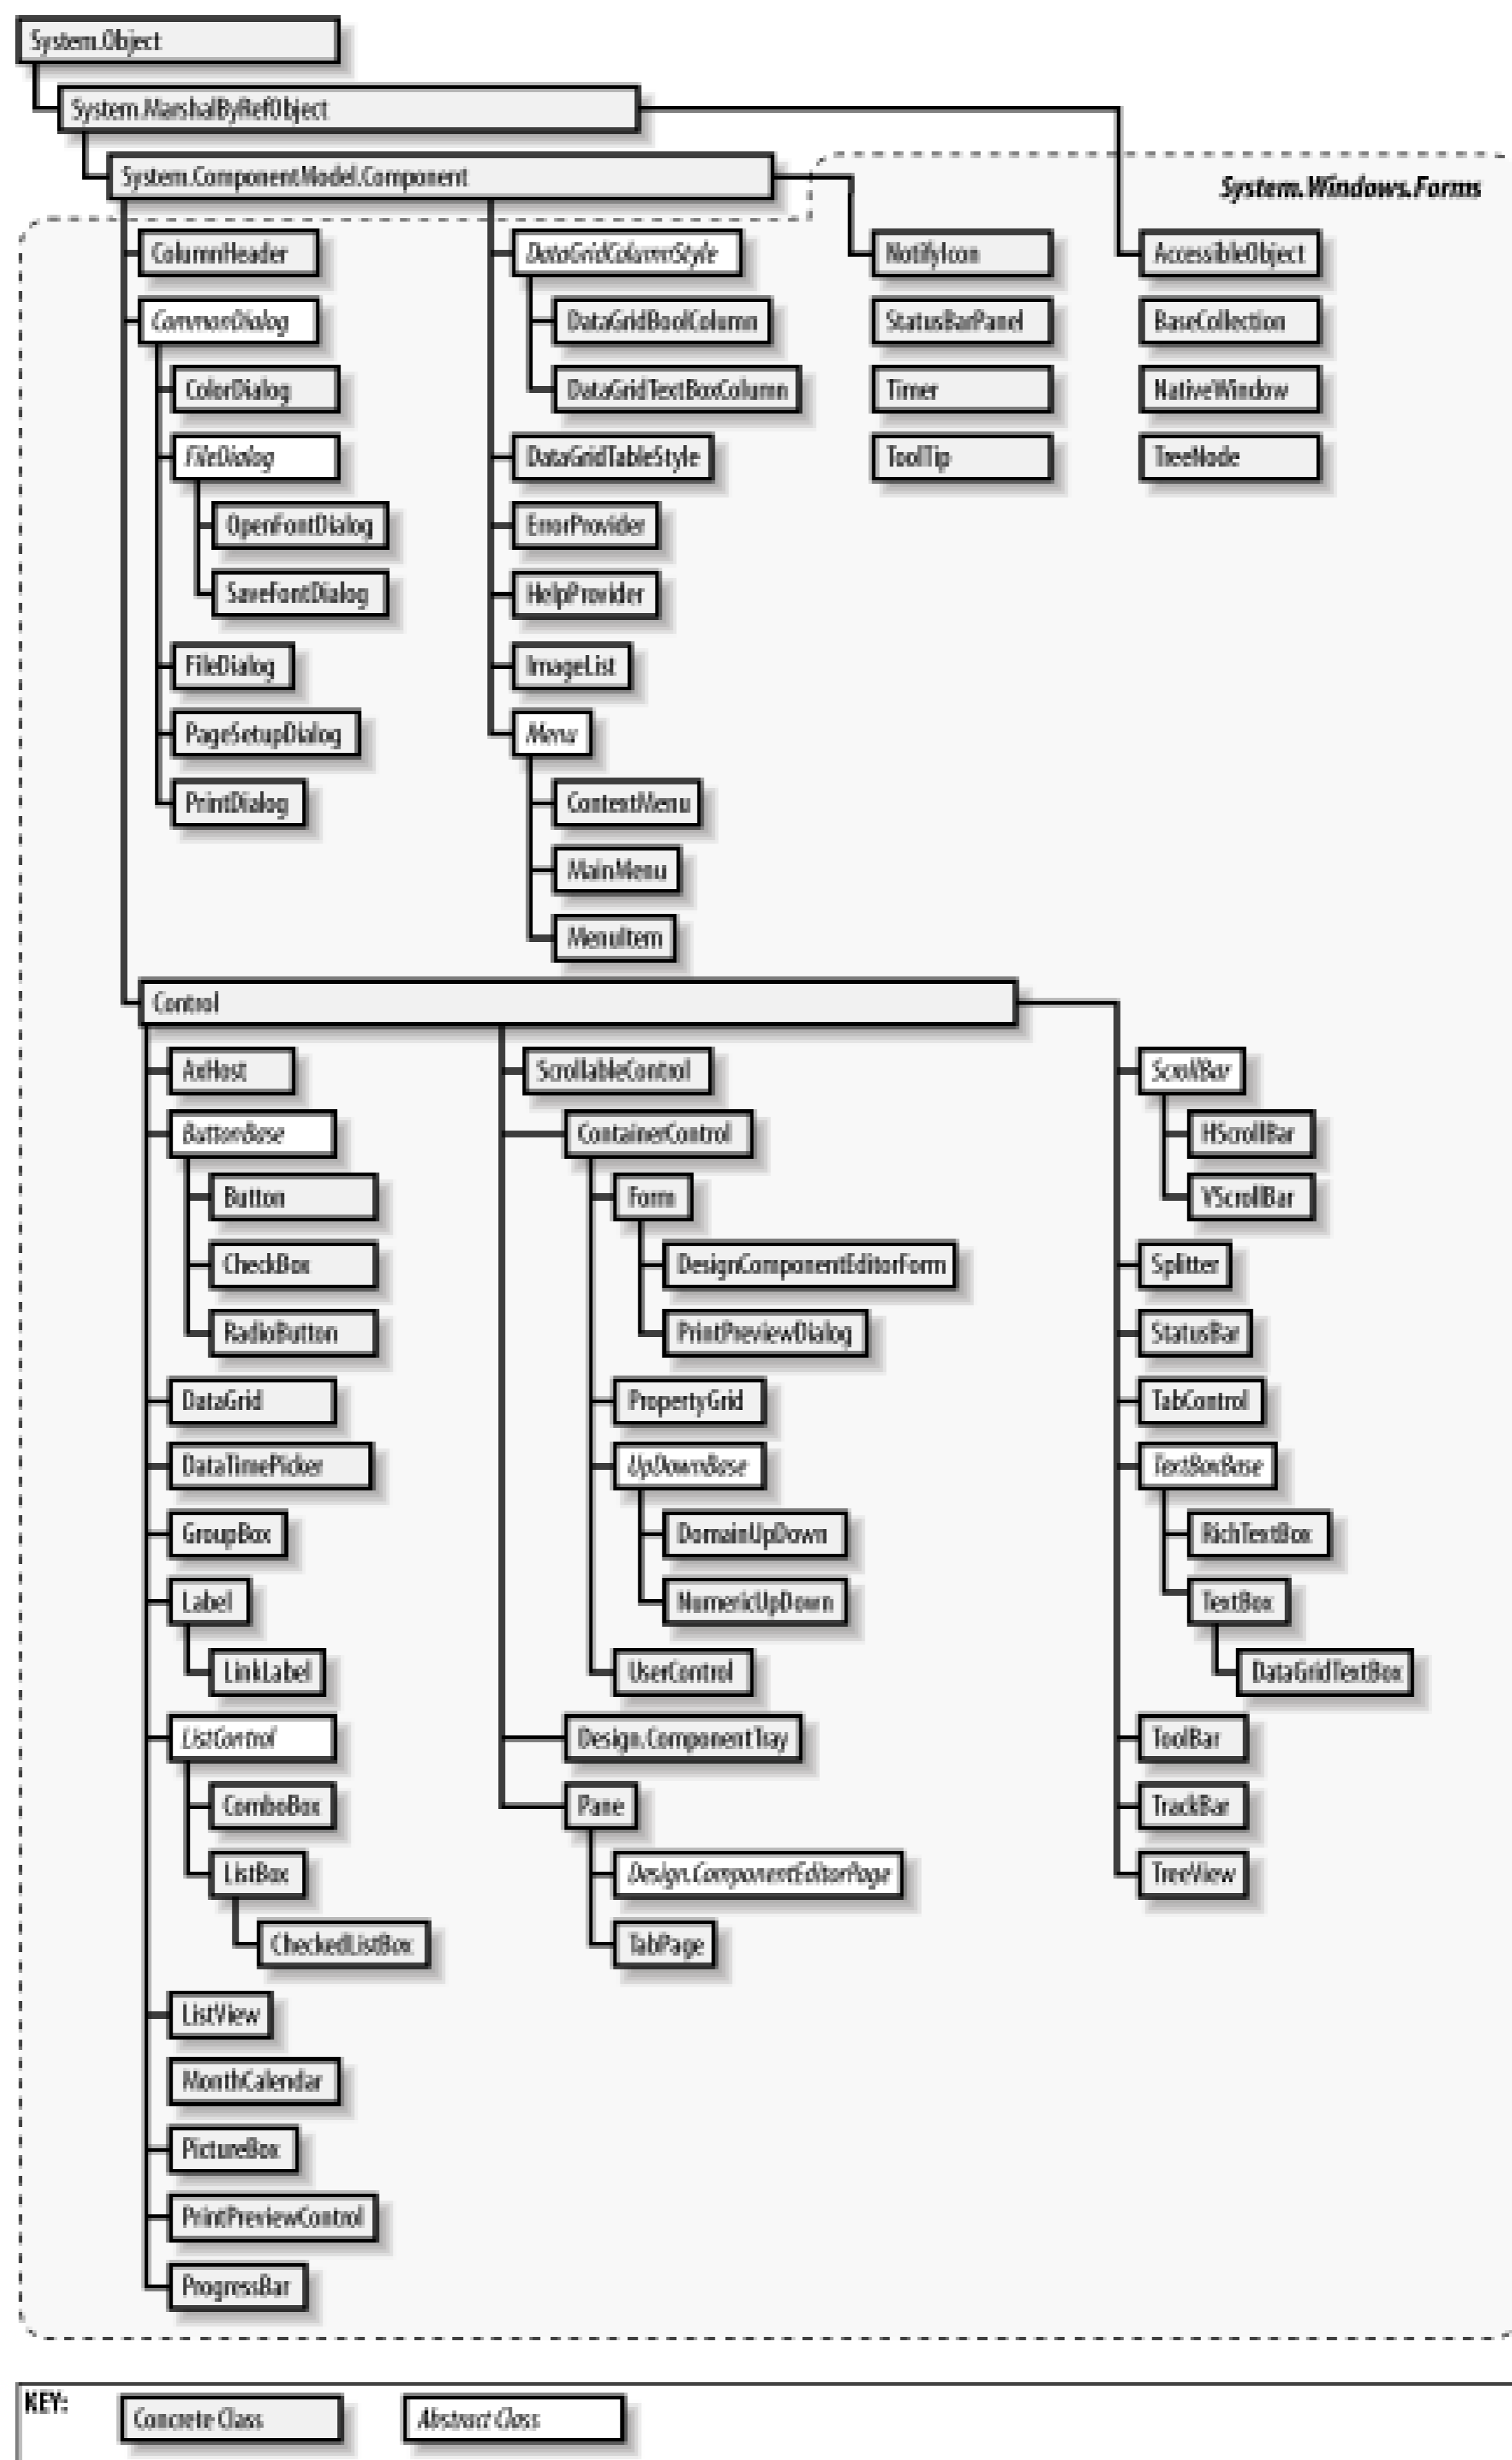
Properties	Description
<code>CommonAppDataRegistry</code>	This is the common application registry key under which common data is stored and shared among all users. It is application specific. If a path does not exist, one is created in the following format: <i>Base Path\CompanyName\ ProductName\ ProductVersion</i> .
<code>StartupPath</code>	This property is the path in which the executable started.
<code>UserAppDataRegistry</code>	This is the registry key where roaming user's data are kept.
Methods	Description
<code>Run</code>	This method starts the application whether it is GUI-based or not.
<code>Exit</code>	This method stops the application by sending the stop message to all message loops in all threads in the application.

Properties	Description
ExitThread	Similarly, this method stops the current message loop in the current thread.
AddMessageFilter	You can also add a message filter to the application to intercept and filter Windows messages. ^[2]
RemoveMessageFilter	You can also remove the message filter.
DoEvents	This method processes all Windows messages currently in the message queue.

[2] The only parameter you need to provide to this method is an object that implements the `IMessageFilter` interface. Currently, the only method in this interface is `PreFilterMessage`, which you have to override to intercept and filter any message. If your `PreFilterMessage` method returns `true`, the Windows message is consumed and not dispatched to its destination. You can let the message pass through by returning `false` in this method.

[Figure 8-2](#) illustrates the hierarchy of Windows Controls in the `System.Windows.Forms` namespace. These controls are placed on the form to create Windows Forms applications and on a `UserControl` container to create UI Controls (similar to current ActiveX controls). This figure does not include the `Application` class.

Figure 8-2. `System.Windows.Forms` Windows Controls class hierarchy



[Team LiB]

[\[Team LiB \]](#)

8.3 Windows Forms Development

The Form class in the System.Windows.Forms namespace represents a standard window that contains Windows controls. In this section, we walk you through the development of a Windows Forms application and introduce you to the rich set of Windows controls that can be used on a Windows Form.

8.3.1 Windows Forms Application

All Windows Forms applications start out with a derived class from the System.Windows.Forms.Form class. A simple Windows Forms application looks like the following:

```
public class MyForm : System.Windows.Forms.Form
{
    public MyForm( )
    {
        Text = "Hello World";
    }
    public static void Main( )
    {
        System.Windows.Forms.Application.Run(new MyForm( ));
    }
}
```

Basically, you define a class MyForm, which derives from the System.Windows.Forms.Form class. In the constructor of MyForm class, you set the Text property of the Form to `Hello World`. That's all there is to it. The static Main function is the entry point to all applications. In the entry-point function you call the static method `Application.Run`, which starts the message loop for the application. Because you also pass a form-derived object MyForm to the Run method, what we have is a Windows Forms application.

You can also include references to the namespaces to avoid typing the fully qualified name of classes such as System.Windows.Forms.Form or System.Windows.Forms.Application. To do this, include the following line at the beginning of the source file and omit the System.Windows.Forms prefix to your class names:

```
using System.Windows.Forms;
```

To build the previously listed application, we use the command-line C# compiler. Notice that the target type is an executable, not a DLL, as when we compiled our web service PubsWS (type this command all on one line):^[3]

^[3] You can also compile the simple file with `csc MyForm.cs` but it's better to know how to specify the target type and the references that your source relies on.

```
csc /t:winexe
```

```

/r:system.dll
/r:System.Windows.Forms.dll
MyForm.cs

```

The standard Form object that is shown on the screen doesn't do much; however, it demonstrates the simplicity of creating a Windows Forms application. You can exit the application by clicking on the Close button of the Control Box on the titlebar of the form. When you do this, a quit message is injected into the message loop, and, by default, it is processed and the Application instance will stop.

8.3.2 Windows Controls

Windows Forms applications can be much more involved than the application shown earlier; however the underlying concepts are the same. In this section, we introduce you to the rich set of Windows controls that you can use on your form, as well as data binding to some of these controls. We also show how event handling works in Windows Forms applications.

8.3.2.1 Adding controls onto the form

First of all, we create and add the control to the Controls collection of the form:

```

Button btn1 = new Button( );
btn1.Text = "Click Me";
this.Controls.Add(btn1);

```

Adding other types of controls follows the same convention. There are three basic steps:

1. Create the control.
2. Set up the control's properties.
3. Add the control to the Controls collection of the Form object.

8.3.2.2 Binding the event handler

This is swell, but what does the application do when you click on the button? Nothing. We have not yet bound the event handler to the button's event. To do that, we first have to create the event handler. An event handler is nothing more than a normal function, but it always has two parameters: object and EventArgs. The object parameter is filled with event originator. For example, if you clicked on a button on a form, causing the `Click` event to fire, the object parameter to the event handler will point to the button object that you actually clicked on. The EventArgs object represents the event itself. Using the same example, the EventArgs parameter will be the `Click` event with event arguments, such as the coordinates of the mouse, which button got clicked and so on. The following code excerpt shows the event handler for the `Click` event on a button:

```

void btn1_onclick(Object sender, EventArgs e)
{
    Text = "Sender: " + sender.ToString( ) + " - Event: " + e.ToString( );
}

```

```
}
```

That event handler changes the title of the form each time the button is clicked. Now that we have created the event handler, we assign it to the event `Click` of the button:

```
btn1.Click += new EventHandler(btn1_onclick);
```

That line of code constructs an `EventHandler` object from the method we passed in and passes the newly created object to the `Click` event of the button. We basically register a callback function when `Click` happens. (You may want to review [Chapter 2](#) where we discuss delegates.) Here is the complete example:

```
using System;
using System.Windows.Forms;

public class MyForm : Form
{
    void btn1_onclick(object sender, EventArgs e)
    {
        Text = "Sender: " + sender.ToString( ) +
            " - Event: " + e.ToString( );
    }

    public MyForm( )
    {
        Text = "Hello World";

        Button btn1 = new Button( );
        btn1.Text = "Click Me";
        this.Controls.Add(btn1);

        btn1.Click += new EventHandler(btn1_onclick);
    }

    public static void Main( )
    {
        Application.Run(new MyForm( ));
    }
}
```

When the user clicks on the button, our event handler is called because we've already registered for the click event. It is possible to add more than one event handler to a single event by repeating the assignment line for other event handlers. All handlers that are registered to handle the event are executed in the order in which they're registered. For example, we add the following function to the code:

```
void btn1_onclick2(object sender, EventArgs e)
{
    MessageBox.Show(String.Format("Sender: {0} - Event: {1}",
        sender.ToString(), e.ToString( )));
}
```

and one more line to associate this function as an event handler in the case button `btn1` is clicked:

```
btn1.Click += new EventHandler(btn1_onclick);
btn1.Click += new EventHandler(btn1_onclick2);
```

The result is as expected. Both event handlers get called.

You can also easily remove the event handler. Replace `+=` with `-=`:

```
btn1.Click -= new EventHandler(btn1_onclick);
```

Binding event handlers to events at runtime provides the developer with unlimited flexibility. You can programmatically bind different event handlers to a control based on the state of the application. For example, a button click can be bound to the *update function* when the data row exists or to the *insert function* when it's a new row.

As you can see, the process of binding event handlers to events is the same in Windows Forms as in Web Forms. This consistency of programming model is possibly due their shared substrate, the CLR in both environments.

8.3.2.3 Data binding

There are two kinds of data binding in Windows Forms. The first involves simple Windows controls such as Label, TextBox, and Button. These simple controls can be bound to a single value only. The second involves Windows controls that can manage lists of data such as ListBox, ComboBox, and DataGrid. These list controls are bound to lists of values.

Let's look at the first type of data binding. In the following example, we bind text boxes to fields in a table from the Pubs database. We extend the simple *Hello, World* Windows Form application to include data access and data binding.

The first thing is to obtain the data from the database. (It's a good time to review ADO.NET in [Chapter 5](#) if you did not read the book in the order presented.) Let's take a look at the following example of a C# file:

```
using System;
using System.Windows.Forms;
using System.Data;
using System.Data.OleDb;

public class MyForm : Form
{
    public static void Main( )
    {
        Application.Run(new MyForm( ));
    }

    private TextBox m_txtFirstName, m_txtLastName, m_txtPhone;
    private Button m_btnPrev, m_btnNext;
    private CurrencyManager m_lm;
    private DataSet m_ds;
```

```
public MyForm( )
{
    Text = "Simple Controls Data Binding";

    // Create the first name text box.
    m_txtFirstName = new TextBox( );
    m_txtFirstName.Dock = DockStyle.Top;

    // Create the last name text box.
    m_txtLastName = new TextBox( );
    m_txtLastName.Dock = DockStyle.Top;

    // Create the phone text box.
    m_txtPhone = new TextBox( );
    m_txtPhone.Dock = DockStyle.Top;

    // Add both first name and last name to the panel1.
    Panel panel1 = new Panel( );
    panel1.Dock = DockStyle.Left;
    panel1.Controls.Add(m_txtFirstName);
    panel1.Controls.Add(m_txtLastName);
    panel1.Controls.Add(m_txtPhone);
    // Add panel1 to the left of the form.
    this.Controls.Add(panel1);

    // Create the up button and bind click to event handler.
    m_btnPrev = new Button( );
    m_btnPrev.Text = "Up";
    m_btnPrev.Dock = DockStyle.Top;
    m_btnPrev.Click += new EventHandler(btnPrev_onclick);

    // Create the down button and bind click to event handler.
    m_btnNext = new Button( );
    m_btnNext.Text = "Down";
    m_btnNext.Dock = DockStyle.Top;
    m_btnNext.Click += new EventHandler(btnNext_onclick);

    // Add both the up and down buttons to panel2.
    Panel panel2 = new Panel( );
    panel2.Dock = DockStyle.Right;
    panel2.Width = 50;
    panel2.Controls.Add(m_btnNext);
    panel2.Controls.Add(m_btnPrev);
    // Add panel2 to the right of the form.
    this.Controls.Add(panel2);

    // Fill the dataset with the authors table from Pubs database.
    m_ds = new DataSet( );
    string oSQL = "select au_fname, au_lname, phone from authors";
    string oConnStr =
        "provider=sqloledb;server=(local);database=pubs;Integrated Security=SSPI";
```

```

OleDbDataAdapter oDA = new OleDbDataAdapter(oSQL, oConnStr);
oDA.Fill(m_ds, "tbl");

// Bind the Text property of last name text box to field au_lname.
m_txtLastName.DataBindings.Add("Text",
                                m_ds.Tables["tbl"],
                                "au_lname");

// Bind the Text property of first name text box to field au_fname.
m_txtFirstName.DataBindings.Add("Text",
                                 m_ds.Tables["tbl"],
                                 "au_fname");

// Bind the Text property of phone text box to field phone.
m_txtPhone.DataBindings.Add("Text",
                             m_ds.Tables["tbl"],
                             "phone");

// Obtain the list manager from the binding context.

m_lm = (CurrencyManager)this.BindingContext[m_ds.Tables["tbl"]];
}

protected void btnNext_onclick(object sender, EventArgs e)
{
    // Move the position of the list manager.
    m_lm.Position += 1;
}
protected void btnPrev_onclick(object sender, EventArgs e)
{
    // Move the position of the list manager.
    m_lm.Position -= 1;
}
}

```

UI controls derive from the Control class, and inherit the DataBindings property (which is of type ControlsBindingCollection). This DataBindings property contains a collection of Binding objects that is used to bind any property of the control to a field in the list data source.

To bind a simple control to a record in the data source, we can add a Binding object to the DataBindings collection for the control using the following syntax:

```
controlName.DataBindings.Add("Property", datasource, "columnname");
```

where `controlName` is name of the simple control that you want to perform the data binding. The `Property` item specifies the property of the simple control you want to be bound to the data in column `columnname`.

The C# source file shows how to bind the Text property of the TextBox control `m_txtLastName` to the `au_lname` column of Authors table of the DataSet `m_ds`, as well as `m_txtFirstName` and `m_txtPhone` to columns `au_fname` and `phone`.

To traverse the list in the data source, we will use the `BindingManagerBase` object. The following excerpt of code shows you how to get to the binding manager for the data source bound to the controls on the form. In this case, because the data is of list type, the binding manager returned from the `BindingContext` is a `CurrencyManager`:^[4]

[4] If the data source returns only one data value, the `BindingManagerBase` actually points to an object of type `PropertyManager`. When the data source returns a list of data value, the type is `CurrencyManager`.

```
// Obtain the list manager from the binding context.
m_lm = (CurrencyManager)this.BindingContext[m_ds.Tables["tbl"]];
```

To demonstrate the use of `BindingManagerBase` to traverse the data source, we add two buttons onto the form, `btnNext` and `btnPrev`. We then bind the two buttons' click events to `btnNext_onclick` and `btnPrev_onclick`, respectively:

```
protected void btnNext_onclick(object sender, EventArgs e)
{
    m_lm.Position += 1;
}

protected void btnPrev_onclick(object sender, EventArgs e)
{
    m_lm.Position -= 1;
}
```

As you use `BindingManagerBase` to manage the position of the list-in this case, the current record in the Authors table-the `TextBox` controls will be updated with new values. [Figure 8-3](#) illustrates the user interface for the simple controls data-binding example.

Figure 8-3. Simple controls data binding

Now let's take a look at the other type of data binding. In this example, we will bind the whole authors table to a `DataGrid`:

```
using System;
using System.Windows.Forms;
using System.Data;
using System.Data.OleDb;

public class MyForm : Form
{
    public static void Main( )
    {
        Application.Run(new MyForm( ));
    }
}
```

```
private Button m_btn1;
private TextBox m_txt1;
private DataGridView m_dataGrid1;

public MyForm( )
{
    Text = "Hello World";

    m_txt1 = new TextBox( );
    m_txt1.Text = "select * from authors";
    m_txt1.Dock = DockStyle.Top;
    this.Controls.Add(m_txt1);

    m_btn1 = new Button( );
    m_btn1.Text = "Retrieve Data";
    m_btn1.Dock = DockStyle.Top;
    m_btn1.Click += new EventHandler(btn1_onclick);
    this.Controls.Add(m_btn1);

    m_dataGrid1 = new DataGridView( );
    m_dataGrid1.Dock = DockStyle.Fill;
    this.Controls.Add(m_dataGrid1);

    this.AcceptButton = m_btn1;
}

protected void btn1_onclick(object sender, EventArgs e)
{
    try {
        DataSet ds = new DataSet( );
        string oConnStr =
            "provider=sqloledb;server=(local);database=pubs;Integrated Security=SSPI";
        OleDbDataAdapter oDA =
            new OleDbDataAdapter(m_txt1.Text, oConnStr);
        oDA.Fill(ds, "tbl");

        /* You can specify the table directly like this
        *
        *     m_dataGrid1.DataSource = ds.Tables["tbl"];
        *
        * or specify the datasource and the table separately
        * like this:
        */
        m_dataGrid1.DataSource = ds;
        m_dataGrid1.DataMember = "tbl";

    } catch(Exception ex) {
        MessageBox.Show("An error has occurred. " + ex.ToString( ));
    }
}
}
```


Data binding for controls of type List in Windows Forms is similar to that of Web Forms. However, you don't have to call the `DataBind` method of the control. All you have to do is set the `DataSource` property of the UI control to the data source. The data source then has to implement the `IEnumerable` or `IListSource` (or `IList`, which implements `IEnumerable`) interfaces. As it turns out, there are hundreds of classes that can be used as data source, including `DataTable`, `DataView`, `DataSet`, and all array or collection type of classes.

The process for `DataGrid` data binding is also simple: just set the `DataSource` property of the `DataGrid` object to the data source, and you're all set. We name the table `tbl` when we add it to `DataSet` with the data adapter's `Fill()` method; therefore, the following line of code just indexes into the collection of tables in the `DataSet` using the table name:

```
m_dataGrid1.DataSource = ds.Tables["tbl"];
```

If the data source contains more than one table, you will also have to set the `DataMember` property of the `DataGrid` to the name of the table you want the control to bind to:

```
m_dataGrid1.DataSource = ds;  
m_dataGrid1.DataMember = "tbl";
```

The results of binding the two tables to the `DataGrid` are shown in [Figure 8-4](#) and [Figure 8-5](#).

Figure 8-4. Binding the authors table to the `DataGrid`

Figure 8-5. Binding the titles table to the `DataGrid`

title_id	title	type	pub_id	price	advance	royalty	ytd_sales	notes	pubdate
BU1032	The Buzz Eye	business	1389	19.99	5000	10	4895	An overview o	6/12/1991
BU1111	Cooking with	business	1389	11.95	5000	10	3876	Helpful hints o	6/9/1991
BU2075	You Can Com	business	0736	2.99	10025	24	18722	The latest me	6/30/1991
BU7832	Straight Talk	business	1389	19.99	5000	10	4895	Annotated an	6/22/1991
MC2222	Silicon Valley	med_cook	0877	19.99	0	12	2032	Favorite recip	6/9/1991
MC3021	The Gourmet	med_cook	0877	2.99	15000	24	22246	Traditional Fr	6/18/1991
MC3026	The Psycholo	UNDECIDED	0877	(null)	(null)	(null)	(null)	(null)	2/8/2001
PC1035	But Is It Usea	popular_comp	1389	22.95	7000	16	8780	A survey of so	6/30/1991
PC8888	Secrets of Sil	popular_comp	1389	20	8000	10	4895	Muckraking re	6/12/1994
PC9999	Net Etiquette	popular_comp	1389	(null)	(null)	(null)	(null)	A must-read i	2/8/2001
PS1372	Computer Ph	psychology	0877	21.99	7000	10	375	A must for the	10/21/1991
PS2091	Is Anger the E	psychology	0736	10.95	2275	12	2045	Carefully rize	6/15/1991
PS2106	Life Without F	psychology	0736	7	6000	10	111	New exercise,	10/5/1991
PS3333	Prolonged Dal	psychology	0736	19.99	2000	10	4072	What happen	6/12/1991
PS7777	Emotional Se	psychology	0736	7.99	4000	10	3336	Protecting you	6/12/1991

8.3.2.4 Arranging controls

After adding controls onto the form and setting the event handlings and data bindings, you are fully functional. However, for the visual aspect of your application, you might want to change the layout of the controls on the form. You can do this by setting up physical locations of controls with respect to the container to which the controls belong,^[5] or you can dock or anchor the controls inside the container.

[5] This is similar to VB programming. Controls initially have absolute positions on the form, but they can be programmatically moved and resized while the application is running.

Docking of a control is very simple. You can dock your control to the top, left, right, or bottom of the container. If you dock your control to the top or the bottom, the width of your control will span the whole container. On the same token, if you dock the control to the left or the right, its height will span the height of the container. You can also set the Dock property to `DockStyle.Fill`, which will adjust the control to fill the container.

The anchoring concept is a bit different. You can anchor your control inside your container by tying it to one or more sides of the container. The distance between the container and the control remains constant at the anchoring side.

You can also use a combination of these techniques by grouping controls into multiple panels and then organizing these panels on the form. With docking and anchoring, there is no need to programmatically calculate and reposition or resize controls on the form.

If you've ever done Java Swing development, you might notice that the current Microsoft .NET Windows Forms framework is similar to JFC with respect to laying out controls; however, it is missing the Layout Manager classes such as GridLayout and FlowLayout to help lay out controls in the containers. We hope that in future releases of the .NET SDK, some sort of layout manager will be included.^[6] Currently, if you are writing your Windows Forms application using Visual Studio .NET, you will have more than enough control over the layout of controls on your form.

[6] Microsoft provides some interesting examples of how you can develop layout managers. The URL is <http://msdn.microsoft.com/library/en-us/dndotnet/html/custlaywinforms.asp>.

8.3.3 Visual Inheritance

Visual inheritance was never before possible on the Windows platform using Microsoft technologies.

Prior to the release of Microsoft .NET (and we are only talking about VB development here), developers used VB templates to reuse a form. This is basically a fancy name for copy-and-paste programming. Each copy of a VB template can be modified to fit the current use. When the template itself is modified, copies or derivatives of the template are not updated. You either have to redo each one using copy-and-paste or just leave them alone.

With the advent of Microsoft .NET, where everything is now object-oriented, you can create derived classes by inheriting any base class. Since a form in Windows Forms application is nothing more than a derived class of the base Form class, you can actually derive from your Form class to create other Form classes.

This is extremely good for something like a wizard-based application, where each of the forms looks similar to the others. You can create the common look-and-feel form as your base class and then create each of the wizard forms by deriving from this base class.

8.3.4 MDI Applications

There are two main styles of user interfaces for Windows-based applications: Single Document Interface (SDI) and Multiple Document Interface (MDI).^[7] For SDI applications, each instance of the application can have only one document. If you would like more than one open document, you must have multiple instances of the application running. MDI, on the other hand, allows multiple documents to be open at one time in one instance of the application. Another good thing about MDI application is that, depending of the type of document currently open, the main menu for the application changes to reflect the operations that you can perform on the document.

^[7] Other styles are Explorer, Wizard, etc., but we are not going discuss all of them in this book.

While it is easy to implement both SDI and MDI applications using the Windows Forms architecture, we only show an example of MDI in this section.

MDI application architecture borrows the same pattern of Windows Forms architecture. Basically, you have one form acting as the container form and other forms acting as child forms.

The Form class provides a number of properties and methods to help in the development of MDI applications, including `IsMdiContainer`, `IsMdiChild`, `MdiParent`, `MdiChildren`, `ActiveMdiChild`, and `LayoutMdi()`.

The first thing we want to show you is the bare minimum main form for our MDI application:

```
using System;
using System.Windows.Forms;

public class MdiMainForm : Form
{
    public MdiMainForm( )
    {
        this.Text = "MDI App for Text and Images";
        // This is the MDI container.
        this.IsMdiContainer = true;
    }
    public static void Main(string[] args)
```

```

    {
        Application.Run(new MdiMainForm( ));
    }
}

```

Believe it or not, this is basically all you have to do for the main form of the MDI application! For each of the child forms that we will be spawning from this main form, we will set its `MdiParent` property to point to this main form.

In the following code excerpt, we load a child form of the main form:

```

. . .
Form a = new Form( );
a.MdiParent = this;
a.Show( );

Form b = new Form( );
b.MdiParent = this;
b.Show( );
. . .

```

Again, all it takes to spawn a child form of the MDI application is a single property, `MdiParent`. In your application, you will replace the type for forms `a` and `b` with your own form classes. (As shown later in this chapter, we have `ImageForm` and `TextForm`.)

One other point that makes MDI applications interesting is the fact that there is one set of main menus and it is possible for child forms to merge their menus with the MDI frame. We also show you how to incorporate menus into our main MDI form, and later in this section, how the child form's menus are merged to this main menu.

The whole menu architecture in Windows Forms application revolves around two classes^[8] `MainMenu` and `MenuItem`. `MainMenu` represents the complete menu for the whole form. A `MenuItem` represents one menu item; however, each menu item contains child menu items in the `MenuItems` property. Again, you start to see the pattern of controls and containers here, too, although the menu classes are not controls.^[9] For example, if we are to have two top-level menus (e.g., `File` and `Window`), then basically, we have to set up the `MainMenu` object so that it contains two menu items in its `MenuItems` property. We can do so using the `Add` method of the `MenuItems` property to insert menu items dynamically into the collection. If we know ahead of time the number of menu items, we can declaratively assign an array of menu items to this property. Recursively, we can have the `File` or the `Window` menu items contain a number of sub-menu items in their `MenuItems` property the same way we set up the main menu.

[8] The third class is `ContextMenu`, but we won't discuss it in the scope of this book.

[9] `MainMenu`, `MenuItem` and `ContextMenu` derive from `Menu`.

Let's take a look at the source code:

```

using System;
using System.Windows.Forms;

public class MdiMainForm : Form
{

```

```
// Menu Items under File Menu
private MenuItem mnuOpen, mnuClose, mnuExit;

// Menu Items under the Window Menu
private MenuItem mnuCascade, mnuTileHorz, mnuTileVert,
                mnuSeparator, mnuCloseAll, mnuListMDI;

// The File and Window Menus
private MenuItem mnuFile, mnuWindow;

// The Main Menu
private MainMenu mnuMain;

public MdiMainForm( )
{
    this.Text = "MDI App for Text and Images";

    // File Menu Item
    mnuFile = new MenuItem( );
    mnuFile.Text = "&File";
    mnuFile.MergeOrder = 0;

    // Window Menu Item
    mnuWindow = new MenuItem( );
    mnuWindow.MergeOrder = 2;
    mnuWindow.Text = "&Window";

    // Main Menu contains File and Window
    mnuMain = new MainMenu( );
    mnuMain.MenuItems.AddRange(
        new MenuItem[2] {mnuFile, mnuWindow});

    // Assign the main menu of the form.
    this.Menu = mnuMain;

    // Menu Items under File menu
    mnuOpen = new MenuItem( );
    mnuOpen.Text = "Open";
    mnuOpen.Click += new EventHandler(this.OpenHandler);
    mnuClose = new MenuItem( );
    mnuClose.Text = "Close";
    mnuClose.Click += new EventHandler(this.CloseHandler);
    mnuExit = new MenuItem( );
    mnuExit.Text = "Exit";
    mnuExit.Click += new EventHandler(this.ExitHandler);
    mnuFile.MenuItems.AddRange(
        new MenuItem[3] {mnuOpen, mnuClose, mnuExit});

    // Menu Items under Window menu
    mnuCascade = new MenuItem( );
    mnuCascade.Text = "Cascade";
    mnuCascade.Click += new EventHandler(this.CascadeHandler);
```

```

        mnuTileHorz = new MenuItem( );
        mnuTileHorz.Text = "Tile Horizontal";
        mnuTileHorz.Click += new EventHandler(this.TileHorzHandler);
        mnuTileVert = new MenuItem( );
        mnuTileVert.Text = "Tile Vertical";
        mnuTileVert.Click += new EventHandler(this.TileVertHandler);
        mnuSeparator = new MenuItem( );
        mnuSeparator.Text = "-";
        mnuCloseAll = new MenuItem( );
        mnuCloseAll.Text = "Close All";
        mnuCloseAll.Click += new EventHandler(this.CloseAllHandler);
        mnuListMDI = new MenuItem( );
        mnuListMDI.Text = "Windows . . . ";
        mnuListMDI.MdiList = true;

        mnuWindow.MenuItems.AddRange(
            new MenuItem[6] {mnuCascade, mnuTileHorz, mnuTileVert ,
                mnuSeparator, mnuCloseAll, mnuListMDI});

        // This is the MDI container.
        this.IsMdiContainer = true;
    }
    public static void Main(string[] args)
    {
        Application.Run(new MdiMainForm( ));
    }
    . . .
}

```

(Note that this source-code listing is completed in the event handlers listing that follows.)

We first declare all the menu items that we would like to have, along with one MainMenu instance in the class scope. In the main-application constructor, we then instantiate the menu items and set their Text properties. For the two top-level menu items, we also set the MergeOrder property so that we can control where the child forms will merge their menu to the main form menu. In this case, we've set up the File menu to be of order 0 and the Window menu to be of order 2. As you will see later, we will have the child menu's MergeOrder set to 1 so that it is between the File and Window menus.

We then add both the File and the Window menus to the main menu's MenuItems collection by using the AddRange() method:

```

mnuMain.MenuItems.AddRange(
    new MenuItem[2] {mnuFile, mnuWindow});

```

Note that at this time, the File and Window menus are still empty. We then assign mnuMain to the MainMenu property of the Form object. At this point, we should be able to see the File and Window menus on the main form; however, there is no drop-down yet.

Similar to how we create menu items and add them to the main menu's MenuItems collection, we add menu items into both the File and Window menu. However, one thing is different. We also bind event handlers to the Click events of the menu items. Let's take one example, the Open menu item:

```
mnuOpen = new MenuItem( );
mnuOpen.Text = "Open";
mnuOpen.Click += new EventHandler(this.OpenHandler);
```

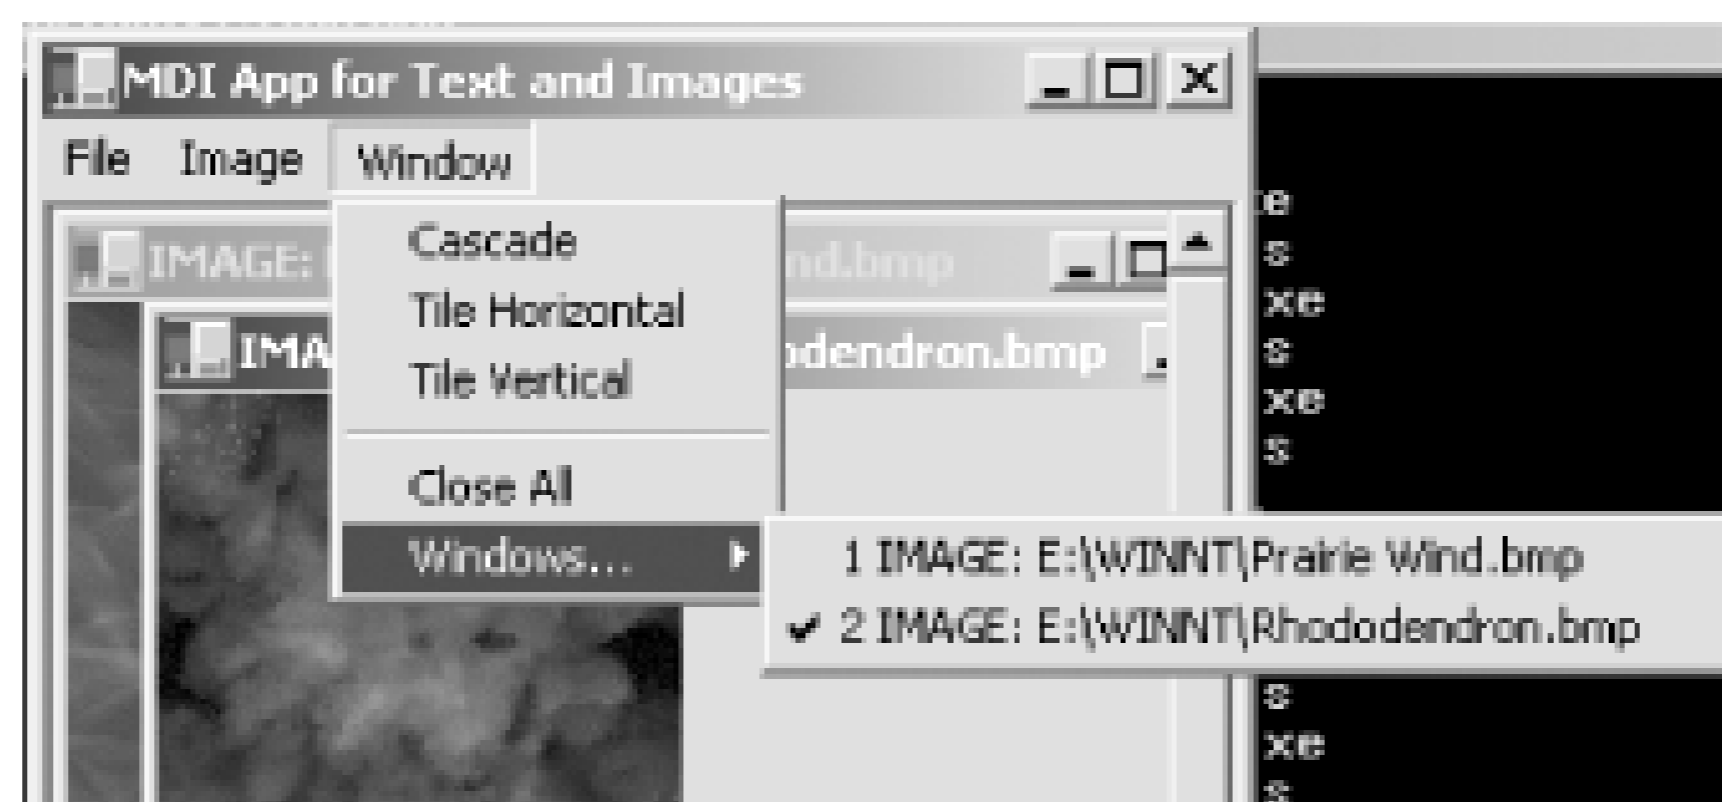
Note that the syntax for binding the event handler `OpenHandler` to the event `Click` of the `MenuItem` class is similar to any other event binding that we've seen so far. Of course, we will have to provide the function body in the MDI main class.

While we are talking about menus, another interesting piece of information is the `mnuListMDI` `MenuItem` at the end of the `Window` menu. We set the `MdiList` property of this `MenuItem` to `true`, as shown in the following code fragment, so that it will automatically show all the opened documents inside the MDI application.

```
mnuListMDI.Text = "Windows . . . ";
mnuListMDI.MdiList = true;
```

See [Figure 8-6](#) for an example of how this feature shows up at runtime.

Figure 8-6. MdiList autogenerated menu entries



The following code is for the event handlers that we've set up for various menu items in this main form (this completes the `MdiMainForm` class listing):

```
protected void OpenHandler(object sender, EventArgs e)
{
    //MessageBox.Show("Open clicked");
    OpenFileDialog openFileDialog = new OpenFileDialog( );
    if(openFileDialog.ShowDialog( ) == DialogResult.OK)
    {
        try
        {
            String sFN = openFileDialog.FileName;
            String sExt = sFN.Substring(sFN.LastIndexOf("."));
            sExt = sExt.ToUpper( );
            //MessageBox.Show(sFN + " " + sExt);
            if(sExt == ".BMP" || sExt == ".JPG" || sExt == ".GIF")
            {
                ImageForm imgForm = new ImageForm( );
                imgForm.SetFileName(sFN);
                imgForm.MdiParent = this;
                imgForm.Show( );
            }
        }
    }
}
```

```
    }
    else if(sExt == ".TXT" || sExt == ".VB" || sExt == ".CS")
    {
        TextForm txtForm = new TextForm( );
        txtForm.SetFileName(sFN);
        txtForm.MdiParent = this;
        txtForm.Show( );
    }
    else
    {
        MessageBox.Show("File not supported.");
    }
}
catch(Exception ex)
{
    MessageBox.Show ("Error: " + ex.ToString( ));
}
}
```

```
protected void CloseHandler(object sender, EventArgs e)
{
    if(this.ActiveMdiChild != null)
    {
        this.ActiveMdiChild.Close( );
    }
}
```

```
protected void ExitHandler(object sender, EventArgs e)
{
    this.Close( );
}
```

```
protected void CascadeHandler(object sender, EventArgs e)
{
    this.LayoutMdi(MdiLayout.Cascade);
}
```

```
protected void TileHorzHandler(object sender, EventArgs e)
{
    this.LayoutMdi(MdiLayout.TileHorizontal);
}
```

```
protected void TileVertHandler(object sender, EventArgs e)
{
    this.LayoutMdi(MdiLayout.TileVertical);
}
```

```
protected void CloseAllHandler(object sender, EventArgs e)
{
    int iLength = MdiChildren.Length;
    for(int i=0; i<iLength; i++)
```



```

    {
        MdiChildren[0].Dispose( );
    }
}

```

The functionality of the OpenHandler event handler is simple. We basically open a common file dialog box to allow the user to pick a file to open. For simplicity's sake, we will support three image formats (BMP, GIF, and JPG) and three text file extensions (TXT, CS, and VB). If the user picks the image-file format, we open the ImageForm as the child form of the MDI application. If a text-file format is selected instead, we use the TextForm class. We will show you the source for both the ImageForm and TextForm shortly.

To arrange the children forms, we use the LayoutMdi method of the Form class. This method accepts an enumeration of type MdiLayout. Possible values are Cascade, ArrangeIcons, TileHorizontal, and TileVertical.

The form also supports the ActiveMdiChild property to indicate the current active MDI child form. We use this piece of information to handle the File Close menu item to close the currently selected MDI child form.

To handle the CloseAll menu click event, we loop through the collection of all MDI child forms and dispose them all.

The following is the source for ImageForm class:

```

using System;
using System.Drawing;
using System.Windows.Forms;

public class ImageForm : System.Windows.Forms.Form
{
    private MenuItem mnuImageItem;
    private MenuItem mnuImage;
    private MainMenu mnuMain;
    private Bitmap m_bmp;

    public ImageForm( )
    {

        mnuImageItem = new MenuItem( );
        mnuImageItem.Text = "Image Manipulation";
        mnuImageItem.Click += new EventHandler(this.HandleImageItem);

        mnuImage = new MenuItem( );
        mnuImage.Text = "&Image";
        mnuImage.MergeOrder = 1; // Merge after File but before Window.
        mnuImage.MenuItems.AddRange(new MenuItem[1] {mnuImageItem});

        mnuMain = new MainMenu( );
        mnuMain.MenuItems.AddRange( new MenuItem[1] {mnuImage});
        this.Menu = mnuMain;
    }
}

```

```

public void SetFileName(String sImageName)
{
    try
    {
        m_bmp = new Bitmap(sImageName);
        Invalidate( );
        this.Text = "IMAGE: " + sImageName;
    }
    catch(Exception ex)
    {
        MessageBox.Show ("Error: " + ex.ToString( ));
    }
}

protected override void OnPaint(PaintEventArgs e)
{
    if(m_bmp != null)
    {
        Graphics g = e.Graphics;
        g.DrawImage(m_bmp, 0, 0, m_bmp.Width, m_bmp.Height);
    }
}

protected void HandleImageItem(object sender, EventArgs e)
{
    MessageBox.Show("Handling the image.");
}
}

```

Because this ImageForm class needs to draw the image file on the form, we include a reference to the System.Drawing namespace. To render the image file onto the form, we rely on the Bitmap and Graphics classes. First of all, we get the input filename and construct the Bitmap object with the content of the input file. Next, we invalidate the screen so that it will be redrawn. In the overridden OnPaint method, we obtained a pointer to the Graphics object and asked it to draw the Bitmap object on the screen.

One other point that we want to show you is the fact that the Image menu item has itsMergeOrder property set to 1. We did this to demonstrate the menu-merging functionality of MDI applications. When this form is displayed, the main menu of the MDI application changes to File, Image, and Window.

To complete the example, following is the source to the TextForm class:

```

using System;
using System.Windows.Forms;
using System.IO;

public class TextForm : Form
{
    private MenuItem mnuTextItem;
    private MenuItem mnuText;
    private MainMenu mnuMain;
}

```

```

private TextBox textBox1;

public TextForm( )
{
    mnuTextItem = new MenuItem( );
    mnuTextItem.Text = "Text Manipulation";
    mnuTextItem.Click += new EventHandler(this.HandleTextItem);

    mnuText = new MenuItem( );
    mnuText.Text = "&Text";
    mnuText.MergeOrder = 1; // Merge after File but before Window.
    mnuText.MenuItems.AddRange(new MenuItem[] {mnuTextItem});

    mnuMain = new MainMenu( );
    mnuMain.MenuItems.AddRange(new MenuItem[] {mnuText});
    this.Menu = mnuMain;

    textBox1 = new TextBox( );
    textBox1.Multiline = true;
    textBox1.Dock = System.Windows.Forms.DockStyle.Fill;
    this.Controls.Add (this.textBox1);
}

public void SetFileName(String sFileName)
{
    StreamReader reader = File.OpenText(sFileName);
    textBox1.Text = reader.ReadToEnd( );
    reader.Close( );
    textBox1.SelectionLength = 0;
    this.Text = "TEXT: " + sFileName;
}

protected void HandleTextItem(object sender, EventArgs e)
{
    MessageBox.Show("Handling the text file.");
}
}

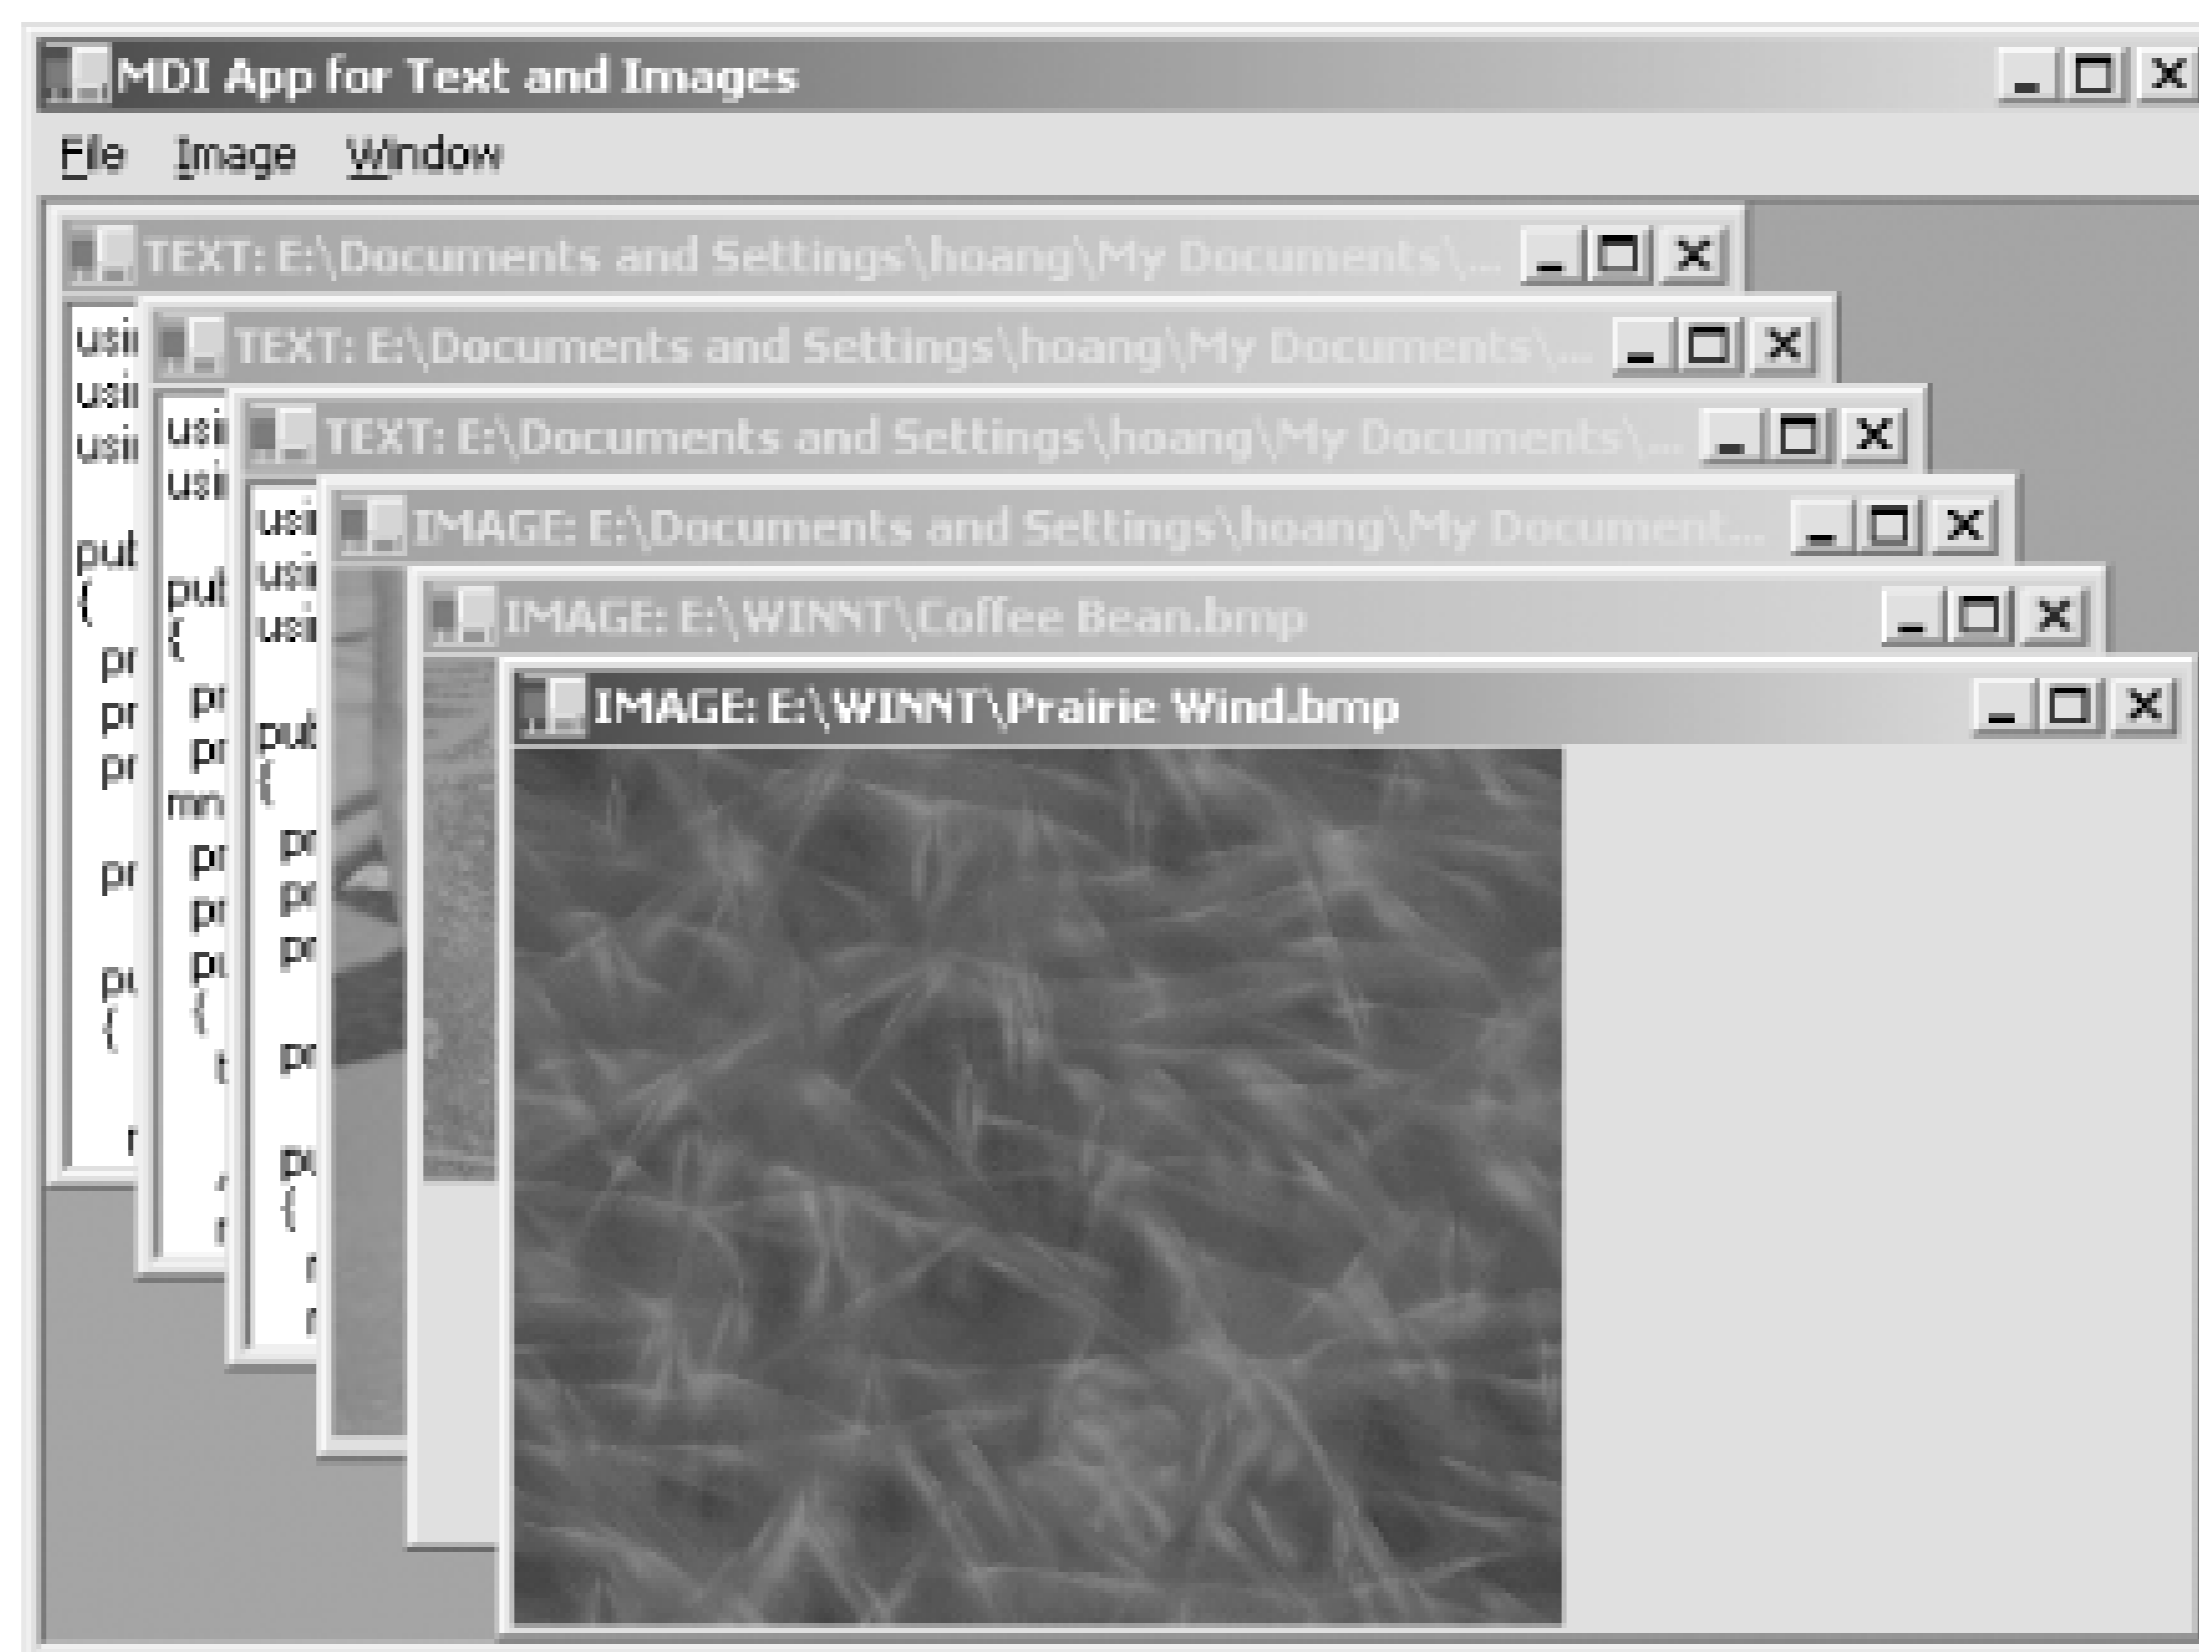
```

Similar to the ImageForm class, the TextForm class also has its menu inserted in the middle of File and Window. When a TextForm becomes the active MDI child form, the menu of the MDI application becomes File, Text, and Window. This menu-merging is done automatically. All we have to do is set up the MergeOrder properties of the menu items.

For the functionality of the TextForm, we have a simple TextBox object. We set its Multiline property to `true` to simulate a simple text editor and have its docking property set to fill the whole form. When the main form passes the text filename to this form, we read the input file and put the content into the text box.

[Figure 8-7](#) illustrates the screen shot for this MDI application at runtime. In this instance, we have three TextForms and three ImageForms open concurrently.

Figure 8-7. The MDI application



The following script is used to build this MDI application. As you can see, the target parameter is set to `winexe` to indicate that the result of the compilation will be an executable instead of `library`, which would result in a DLL. Because we make use of the graphics package for our image rendering, we also have to add the reference to the `System.drawing.dll` assembly. We have three forms in this application: the main form, which is named `MDIApp`, and the two MDI child forms, `ImageForm` and `TextForm` (make sure you type these commands all on one line) [\[10\]](#)

^[10] Again, you can compile the executable without specifying the target type or the references.

```
csc /t:winexe
/r:System.Windows.Forms.dll
/r:system.drawing.dll
MDIApp.cs
ImageForm.cs
TextForm.cs
```

8.3.5 Stage Deployment

Imagine the MDI application from the previous example with hundreds of different types of files supported. For each of the file types, we'd probably have a class similar to the `ImageForm` and `TextForm` classes. Instead of compiling and linking all of these classes into your application, wouldn't it be nice if we could download and install a class on the fly when we use the appropriate file type? In this section, we show how you can convert the previous MDI application to allow just that.

Conceptually, we will have the main executable act as a controller. It will be the first assembly downloaded onto the client machine. Once the user chooses to open a particular file type, this controller determines which supporting DLLs should be downloaded to handle the file type, download and installs the DLL, and then asks the class in the downloaded assembly to display the selected file.

In order for the controller to communicate with all supporting classes, we abstract the commonality out of ImageForm and TextForm to create the abstract class BaseForm:[\[11\]](#)

[11] Because we will demonstrate running MDIApp.exe from a browser later, *BaseForm.dll* has to be explicitly flagged to grant *MDIApp.exe* usage via the AllowPartiallyTrustedCallers attribute. This is to accommodate the Version 1.1 security changes for Microsoft .NET Framework from RCs.

```
using System;
using System.Windows.Forms;
using System.Security;
[assembly:AllowPartiallyTrustedCallers]
namespace BaseForm
{
    abstract public class BaseForm : Form
    {
        abstract public void SetFileName(String sFileName);
    }
}
```

We will rewrite both ImageForm and TextForm to make them inherit from BaseForm. Any other form that you will need to write for other file types will have to also inherit from BaseForm. This way, all the controller module has to do is to know how to talk to BaseForm and everything should be fine.

The followings are the changes to *ImageForm.cs* source file. All we did was to wrap the class inside a namespace call ImageForm, make ImageForm class inherit from BaseForm, and providing the implementation for the abstract method SetFileName(). The rest of the code is the same with the original *ImageForm.cs*.

```
// . . .

namespace ImageForm
{
    public class ImageForm : BaseForm.BaseForm
    {
        // . . .

        override public void SetFileName(String sImageName)
        {
            // . . .
        }

        // . . .
    }
}
```

We apply the same thing to *TextForm.cs* as we do for *ImageForm.cs*. The following summarizes the changes:

```
// . . .

namespace TextForm
```

```

{
    public class TextForm : BaseForm.BaseForm
    {
        // . . .

        override public void SetFileName(String sFileName)
        {
            // . . .
        }

        // . . .
    }
}

```

Now, instead of having the MDIApp compiled and linked with *ImageForm.cs* and *TextForm.cs* and directly using ImageForm and TextForm class in OpenHandler() function, we utilize the reflection namespace to load each assembly from a predetermined URL on the fly. Replace the old code for handling images in OpenHandler():

```

ImageForm imgForm = new ImageForm( );
imgForm.SetFileName(sFN);

```

with:

```

Assembly ass = Assembly.LoadFrom("http://localhost/MDIDLLS/ImageForm.dll");
BaseForm.BaseForm imgForm =
    (BaseForm.BaseForm) ass.CreateInstance("ImageForm.ImageForm");
imgForm.SetFileName(sFN);

```

The three lines of code essentially do the following:

1. Download the assembly from the specified URL if the client does not already have the latest version of the assembly.
2. Create an instance of ImageForm class in ImageForm namespace and cast it to BaseForm.
3. Instruct the BaseForm derived class to deal with the image file.

For the TextForm, replace the following old code for handling text files in OpenHandler():

```

TextForm txtForm = new TextForm( );
txtForm.SetFileFile(sFN);

```

with:

```

Assembly ass = Assembly.LoadFrom("http://localhost/MDIDLLS/TextForm.dll");
BaseForm.BaseForm txtForm =
    (BaseForm.BaseForm) ass.CreateInstance("TextForm.TextForm");
txtForm.SetFileName(sFN);

```

The last thing to do is to make MDIApp use the Reflection namespace by inserting the following line of code:

```
using System.Reflection;
```

You will have to compile *BaseForm.cs*, *ImageForm.cs*, and *TextForm.cs* into *BaseForm.dll*, *ImageForm.dll*, and *TextForm.dll* in that order because ImageForm and TextForm derive from BaseForm.

The following command lines will do the job:

```
csc /t:library /r:System.Windows.Forms.dll baseform.cs
csc /t:library /r:System.Windows.Forms.dll;baseform.dll imageform.cs
csc /t:library /r:System.Windows.Forms.dll;baseform.dll textform.cs
```

You will also have to compile the main module *MDIApp.cs*. Notice that the only thing the *MDIApp.cs* needs to know is BaseForm:

```
csc /t:winexe
    /r:System.Windows.Forms.dll;system.drawing.dll;baseform.dll MDIApp.cs
```

Once you have all the components compiled, copy them all into a virtual directory call MDIDLLs because this is the predefined location where these components are to be downloaded.^[12]

[12] Make sure this virtual directory execute permission does not have executable enabled. In other words, just create the virtual directory with the default settings.

One last step you will have to do is to use the .NET Framework Configuration tool to give file I/O permission to "adjust zone security" and give Local Intranet full trust level. The tool can be found at Administrative Tools/Microsoft .NET Framework Configuration. Choose to Configure Code Access Security and then Adjust Zone Security. Give Local Intranet full trust level for the duration of this experiment. It is not recommended that you use this setting without knowing all consequences; however, this is the simplest thing to do to quickly demonstrate the experiment. For your enterprise application, consult the .NET security documentation for further recommendations.

And now everything is set up. All you have to do is to point your browser to *http://localhost/MDIDLLs/MDIApp.exe* and see how the components are automatically downloaded and run.

Another utility you might want to use to inspect the global assembly cache (where the downloaded components reside) is *gacutil*. Use *gacutil /ldl* to list all downloaded assemblies and *gacutil /cdl* to clear the downloaded assemblies. When you first see the MDIApp main form, the only assembly downloaded is MDIApp. But as soon as you try to open an image file or a text file, ImageForm or TextForm along with BaseForm will be downloaded.

As you can see, with this setup, you won't need to have hundreds of DLLs when you only use a couple of file types. Another cool thing is that when the DLLs on the server are updated, you will automatically have the latest components.

Of course, you can make this MDIApp more production-like by having the main module accessing a Web Service that lists the *file types* this application supports and the *class names* as well as the *assembly filenames* needed to be downloaded. Once this is setup, you can have conditioning code based on the file type; you can pass the *assembly filename URL* to *Assembly.LoadFrom()* method; you can use the class name to create the class instance; all of these make the main module more

generic. Suppose the Web Service that lists the supporting file types for this MDIApp reads information from a database or an XML file. When you need to introduce new file type and supporting assembly to deal with the new file type, all you have to do is add an entry to your database or your XML file and add the assembly file into the virtual directory. We think this is an exercise you should definitely try.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

8.4 Windows Forms and Web Services

Previously, Windows DNA tended not to use fat clients, rich Windows applications on PCs, because of the need for an intrusive installation program. With .NET "copy" deployment, there is no such problem. Now fat Windows clients can interface easily with business logic in the middle tier through XML/HTTP. The rich client application will, in fact, perform better than a web-based front end. However, everything depends on the requirements of the application.

You can still add web references to web services that auto-generate proxy classes for use in your Windows Forms applications just as for Web Forms applications. This is further evidence that .NET as a whole is intended to embrace the Web. In Windows Forms, data binding is automatic once you've set the control's DataSource to the appropriate data source, as seen in the next block of code. Note that there is no explicit call to the DataBind method as in Web Forms data binding:

```
localhost.PubsWS ws = new localhost.PubsWS( );  
DataSet ds = ws.GetAuthors( );  
DataGrid1.DataSource = ds.tables[0].DefaultView;
```

Again, you can also generate the source for the proxy class yourself using the *wSDL.exe* tool, along with the WSDL obtained from the web service. You can then include this source to your Windows Forms project or compile the source into a DLL and add the DLL to the project as a reference.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

8.5 Summary

Windows Forms provides a unified programming model for standard Windows application development. It does not matter what language you are using-you can always be productive because the common substrate has been developed to benefit all.

Windows Forms brings a true object-oriented programming model to Windows GUI development, allowing for an extensible framework that is much cleaner and easier to use compared to previous attempts.

In this chapter, we've shown you the architecture of the Windows Forms application. The Controls and Containers architecture, while very simple, is powerful and flexible for developing Windows-based applications. We have also shown you how to build a Windows Forms application, starting with a simple do-nothing application, then data binding, a multiple-document interface application, a self-downloading client application, and finally a Windows Forms application that pulls data from a web service.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Chapter 9. .NET and Mobile Devices

The increased use by employees of devices other than traditional PCs and laptops, such as personal digital assistants (PDAs) and web-enabled phones, has increased pressure on corporate IT departments to open corporate databases and intranets to a range of non-standard devices. The explosion in use of these devices, and the proliferation of embedded systems and specialized network protocols, has made it important for developers to design their applications in ways that keep it independent of the clients that will access it.

.NET embraces mobile device users and offers developers tools to meet their needs in two ways:

ASP.NET Mobile Controls

Mobile Controls are web server controls that dynamically generate code that renders web content appropriately on devices such as web-enabled phones and PDAs as well as traditional web browsers.

.NET Compact Framework

The .NET Compact Framework is a version of the .NET Platform that runs on Windows CE or Pocket PC mobile devices and can be used to develop clients and applications that run on the devices themselves.

Both ASP.NET Mobile Controls and the .NET Compact Framework are supported by Visual Studio .NET 2003. [\[1\]](#)

[1] In Version 1.0 of the .NET Framework, you have to download separate packages to write ASP.NET mobile applications or applications for smart devices.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

9.1 ASP.NET Mobile Controls

ASP.NET Mobile Controls are special Web Controls that run on the web server and render the appropriate web content targeting a wide array of mobile devices such as mobile phones using Wireless Markup Language (WML) via Wireless Application Protocol (WAP) to PDAs such as Pocket PC Palm Pilot, and so on. The web server receives requests from the web browsers on the mobile device just the same as it would receive requests from a full-fledged browser; however, the Mobile Controls on these specialized web pages know how to render the content based on the capability of the requesting web browsers. For example, if the request came from a WML-capable mobile phone, the Mobile Controls would render WML. If the request came from a Pocket PC web browser, the Mobile Controls would render HTML.

9.1.1 System.Web.UI.MobileControls namespace

In the same way that the `System.Web.UI.WebControls` namespace provides classes for building web pages that can be rendered by traditional web browsers, the `MobileControls` namespace include classes that for constructing web pages that will be rendered by mobile devices. Instead of the Page class discussed in [Chapter 7](#), for example, the `MobileControls` namespace provides a `MobilePage` class. `MobilePage` is the base class to all mobile device Web Forms and derives class from the ASP.NET Page class. All ASP.NET Mobile Controls derive from the `MobileControl` base class, which derives from the corresponding ASP.NET `Control` class.

Here, you also have the choice to use the stock controls provided to you similar to ASP.NET Web Controls and Html Controls. You can also customize controls by deriving and extending stock controls, creating user control or composite server controls or even create your own Mobile Control from deriving directly from `MobileControl` class. Notice that the programming paradigm is still very similar to how you would do it in the normal ASP.NET world.[\[2\]](#)

[2] The programming paradigm for creating user controls or composite server controls is similar to that of ASP.NET. However, if you are to create device-specific custom controls, you will have to be familiar with writing device-specific control adapters to render the control. ASP.NET mobile architecture depends on a device adapter model so that the controls are device-independent. In order for a control to support different devices, a set of adapters are created to adapt the control to the device specific capabilities. We do not cover how to create device-specific custom mobile controls in this book.

9.1.2 ASP.NET Web Server Controls Versus Mobile Controls

[Table 9-1](#) shows the mobile controls and their counterparts in ASP.NET where the mapping is one-to-one.

Table 9-1. ASP.NET to mobile controls

Web Server control	Mobile control
AdRotator	AdRotator
Calendar	Calendar
CompareValidator	CompareValidator
CustomValidator	CustomValidator
DataGrid	ObjectList
Image	Image
HyperLink	Link
Panel	Panel
RangeValidator	RangeValidator
RegularExpressionValidator	RegularExpressionValidator
RequiredFieldValidator	RequiredFieldValidator
TextBox	TextBox
ValidationSummary	ValidationSummary

Although most of Web Server controls map one-to-one to a mobile control, others are mapped down to a simpler control that serves the same purpose. [Table 9-2](#) shows such mappings.

Table 9-2. ASP.NET Web Server controls without one-to-one mappings to Mobile controls

For these ASP.NET controls	Use this Mobile control
Button, ImageButton, LinkButton	Command
DataList, Repeater	List
CheckBox, CheckBoxList, DropDown, DropDownList, ListBox, RadioButton, RadioButtonLists	SelectionList

Other controls, such as PhoneCall, DeviceSpecific, Form, StyleSheet, TextView, are controls that are specific to the mobile environment and do not exist in the standard ASP.NET.

9.1.3 ASP.NET Mobile Web Application

In this section, you will learn how ASP.NET Mobile Controls can be used to create a mobile web application that utilizes the [PubsWS](#) Web Service developed in [Chapter 6](#). This application queries and displays the list of books found in the Pubs database. Using the application, mobile device users can view general information about a book and either buy a hard copy of the title or download it to a

mobile device.

Let's start by opening a new ASP.NET Mobile Web Application in Visual Studio .NET. To open a Mobile Web application, simply select Mobile Web application in the New Project window.

To view information in the Pubs database, you will use an `ObjectList` control. To add the control to your form, you add the following tags in HTML view:

```
<mobile:ObjectList id="bookList" runat="server"
  AutoGenerateFields="False">
  <Field DataField="title"></Field>
  <Field DataField="type" DataFormatString="Type: {0}"></Field>
  <Field DataField="price" DataFormatString="Price: ${0}"></Field>
  <Field DataField="title_id" DataFormatString="RefID: {0}"></Field>
  <Field DataField="notes"></Field>
  <Command Name="buy" Text="Buy"></Command>
  <Command Name="dnload" Text="Download"></Command>
</mobile:ObjectList>
```

The previous code snippet is similar to how you would have a `DataGrid` on an ASP.NET page. In this case the `ObjectList` control is used instead. Similar to the Column tags of the `DataGrid` control, the Field tags are data bound fields that also provide some formatting. The two Command tags represent the two operation that the user will be able to perform on a data item. In this example, the user can buy or download the book.

If you prefer using the VS.NET visual environment, you can drag an `ObjectList` onto the form directly from the Mobile Web Forms menu of the VS.NET Toolbox and then set its properties using the Property window. The preceding code creates an `ObjectList` control, binds it to specific fields of the Pubs database, and provides for two commands-Buy and Download-that that the user can issue on any item selected from the list.

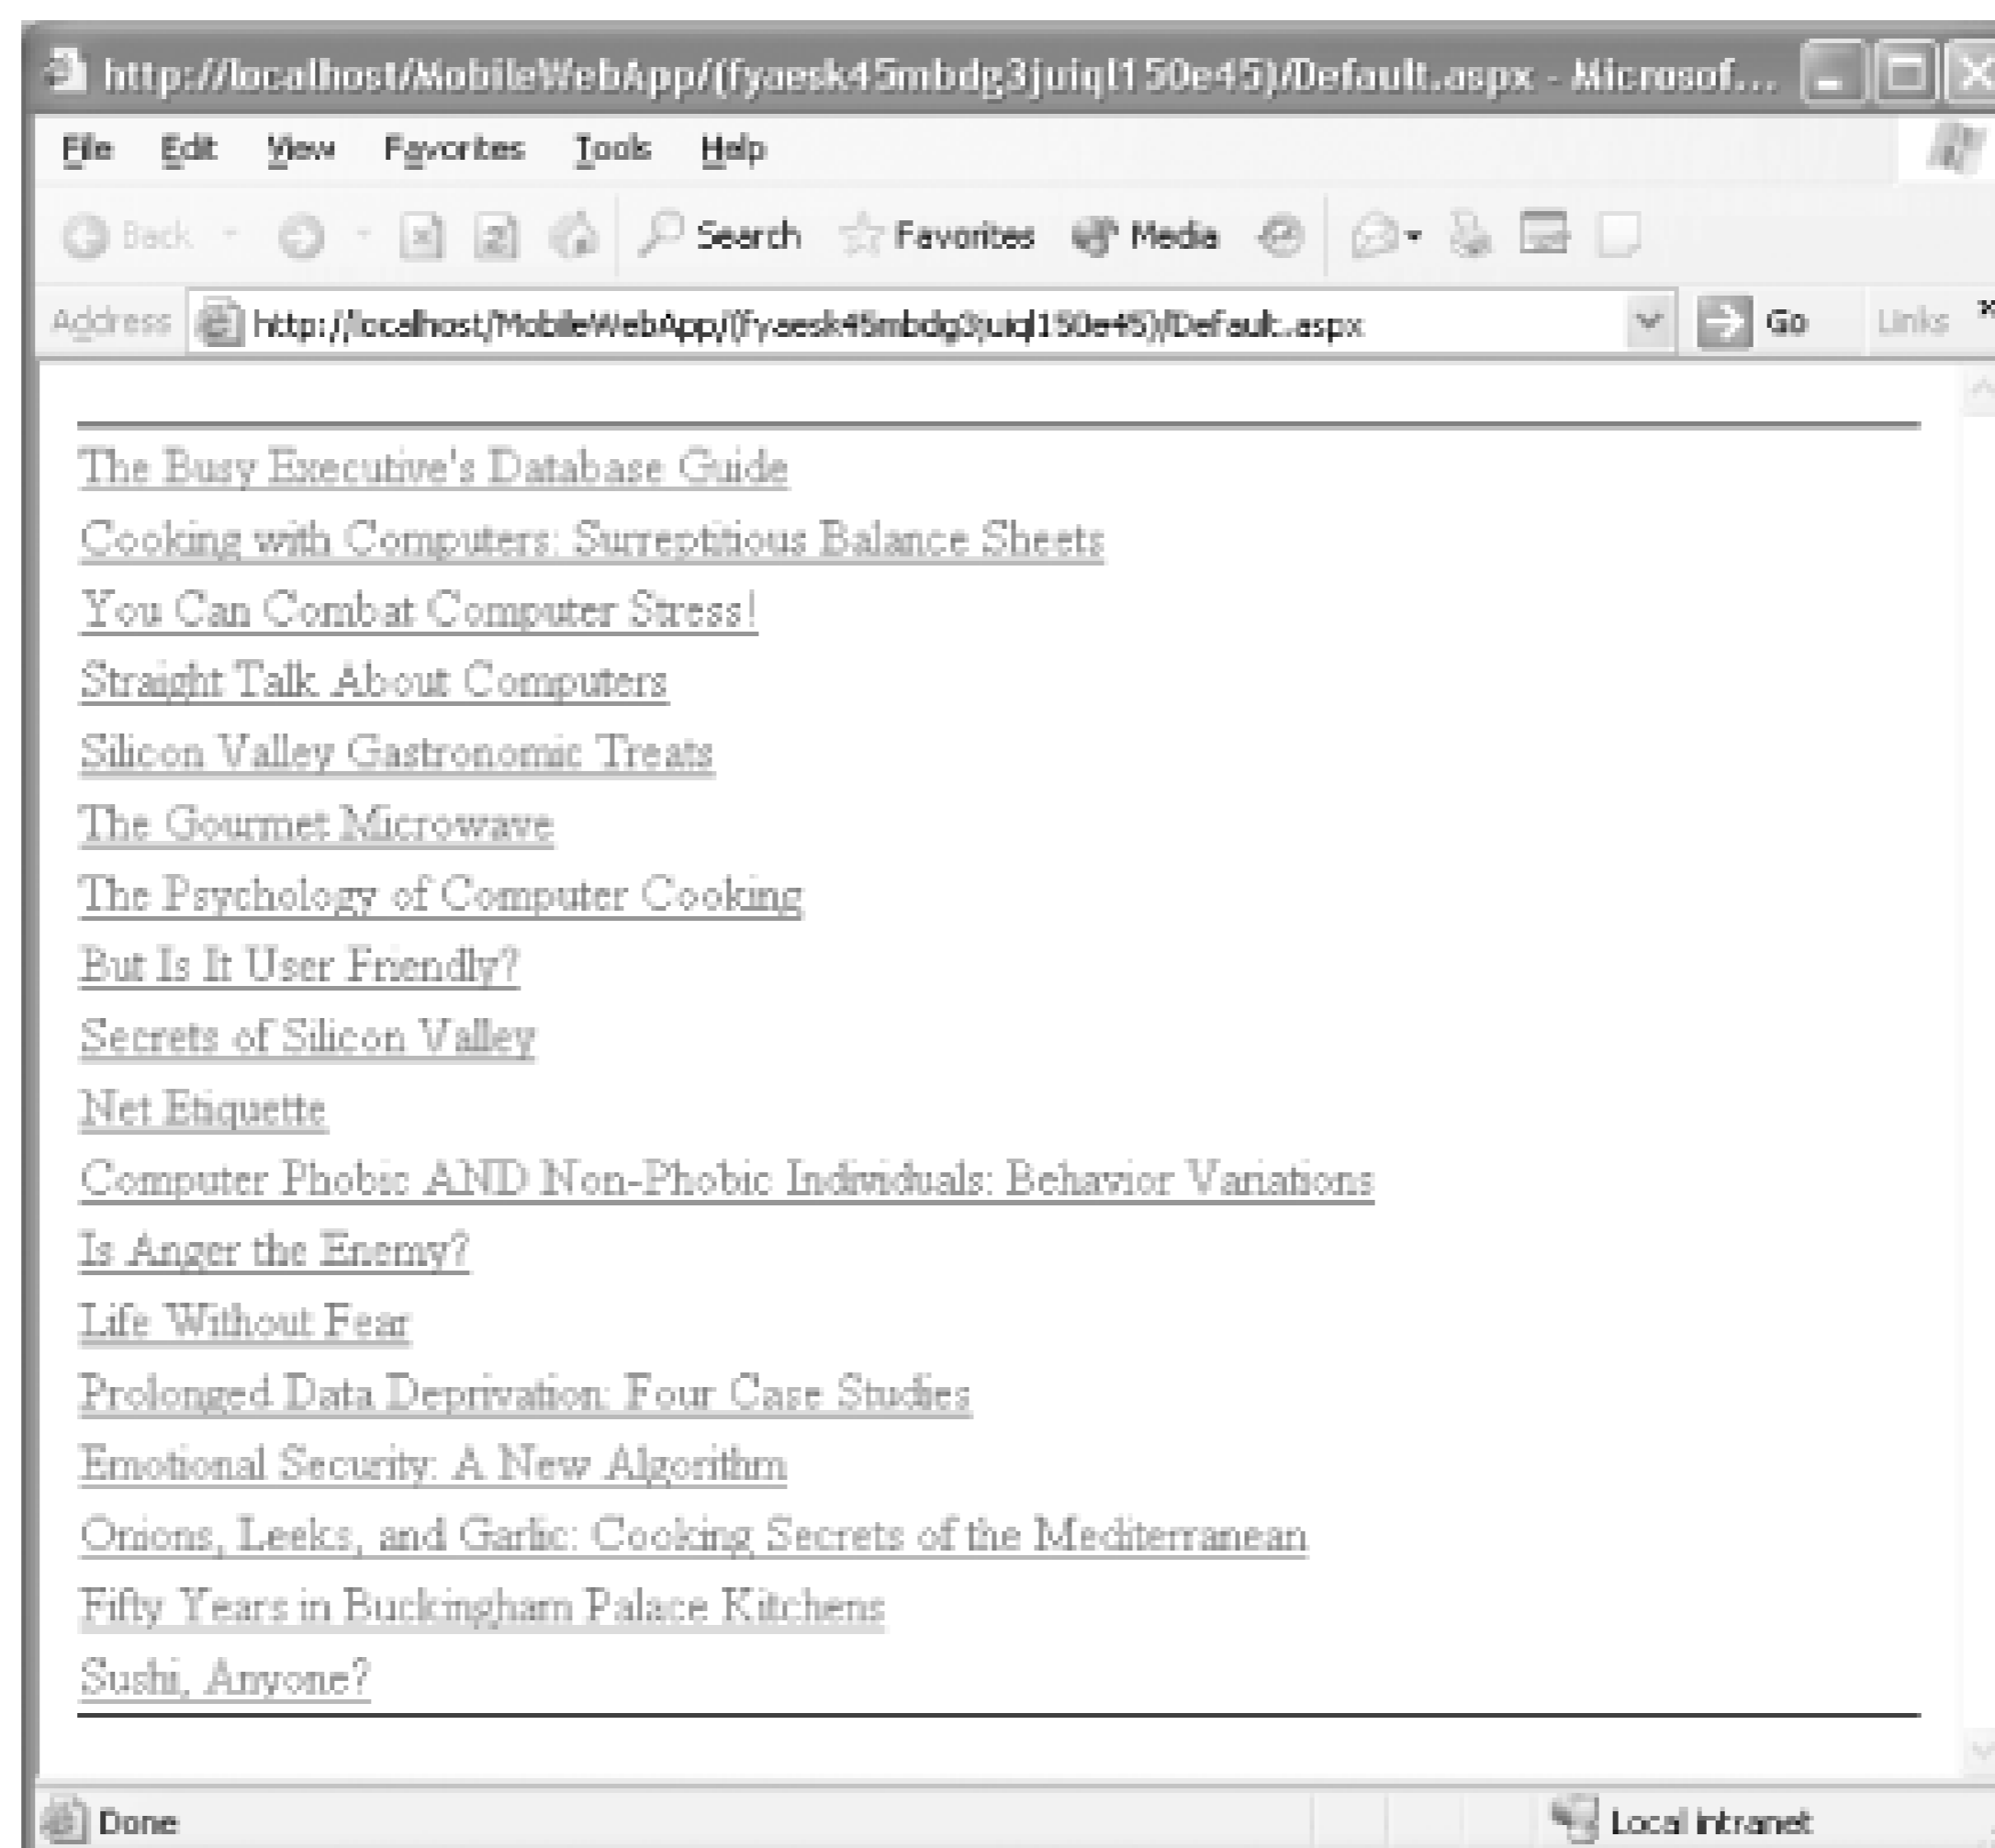
Next, you need to add a web reference to the `PubsWS` Web Service and name it `PubsWSProxy` (see [Chapter 6](#) to review how this is done). Finally, you can switch to the code behind page of your application and write a page load handler that uses the `PubsWS` proxy to load the Pubs data into a `DataSet` that the `ObjectList` can bind to, as shown in the following code fragment:

```
If Not IsPostBack Then
  Dim oProxy As PubsWSProxy.PubsWS = New PubsWSProxy.PubsWS
  oProxy.Credentials = New Net.NetworkCredential("uid", "pwd", "domain")
  Dim oDS As DataSet = oProxy.GetBooks( )
  bookList.DataSource = oDS.Tables(0).DefaultView
  bookList.DataBind( )
End If
```

There is no need to set up the Credentials for the proxy object if your Web Service allows anonymous access.

Now if you run the application, the page shown in [Figure 9-1](#) will display in your browser.

Figure 9-1. Book listing in Microsoft Internet Explorer (IE)

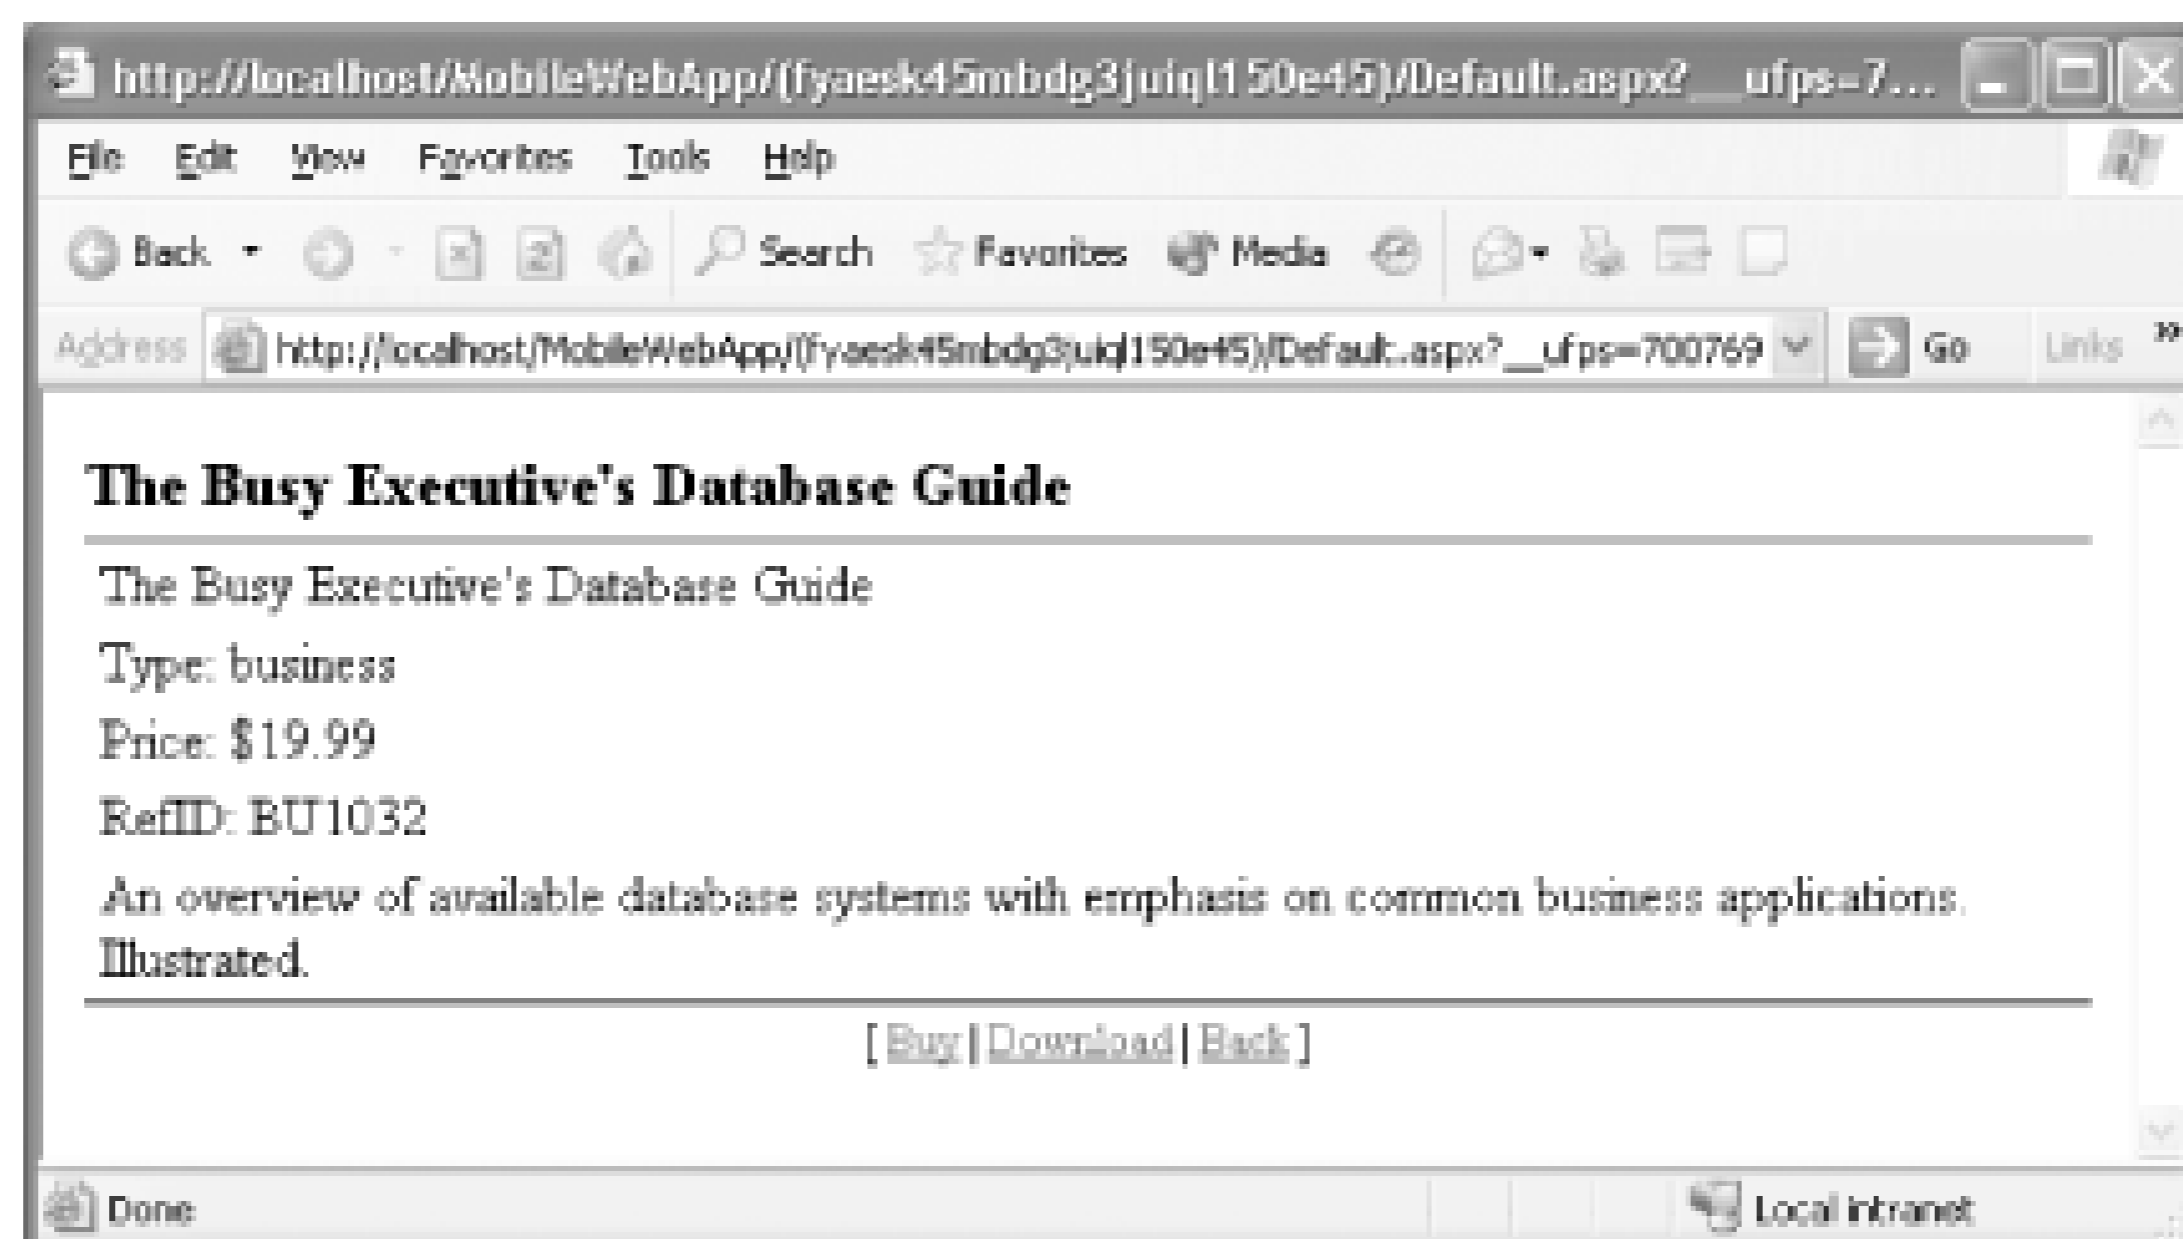


We are viewing this application using a desktop Internet Explorer to test it out. Notice that in [Figure 9-1](#) the `ObjectList` control does not show the bound data as a grid,^[3] but instead as a list of hyperlinks to the first fields of each record.^[4] Clicking on any book in the list will cause all fields of the selected item to be displayed, as shown in [Figure 9-2](#). Note that the commands that apply to the selected item are also shown on the form.

[3] If you want to display an actual table of data on that tiny screen, you will have to set the `TableFields` property of the `ObjectList`. This property takes a set of field names, separated by semicolons. Because the screens of the mobile devices are normally small, by default the `ObjectList` shows the data as a list instead of a table. One other thing to note is that when the page is rendered to WML browsers, the table cannot be displayed, and the `ObjectList` falls back to the list approach.

[4] Depending on the device, this list can be a list of commands instead of hyperlinks.

Figure 9-2. Book detail screen in IE



Now let's add another form to the mobile web application to handle the buying and downloading of the book:

```
<mobile:Form id=FormPurchase runat="server">
  <mobile:TextView id=txtPurchaseInfo runat="server"></mobile:TextView>
  <mobile:Label id=labelCC runat="server">Credit Card #</mobile:Label>
  <mobile:TextBox id=txtCC runat="server"></mobile:TextBox>
  <mobile:Command id=cmdSubmit runat="server">Submit</mobile:Command>
</mobile:Form>
```

The purchase form uses a `TextView` control that displays some simple formatted text and a `TextBox` control to collect a credit card number. The form also makes use of a submit button to post the form data back to the server for processing.

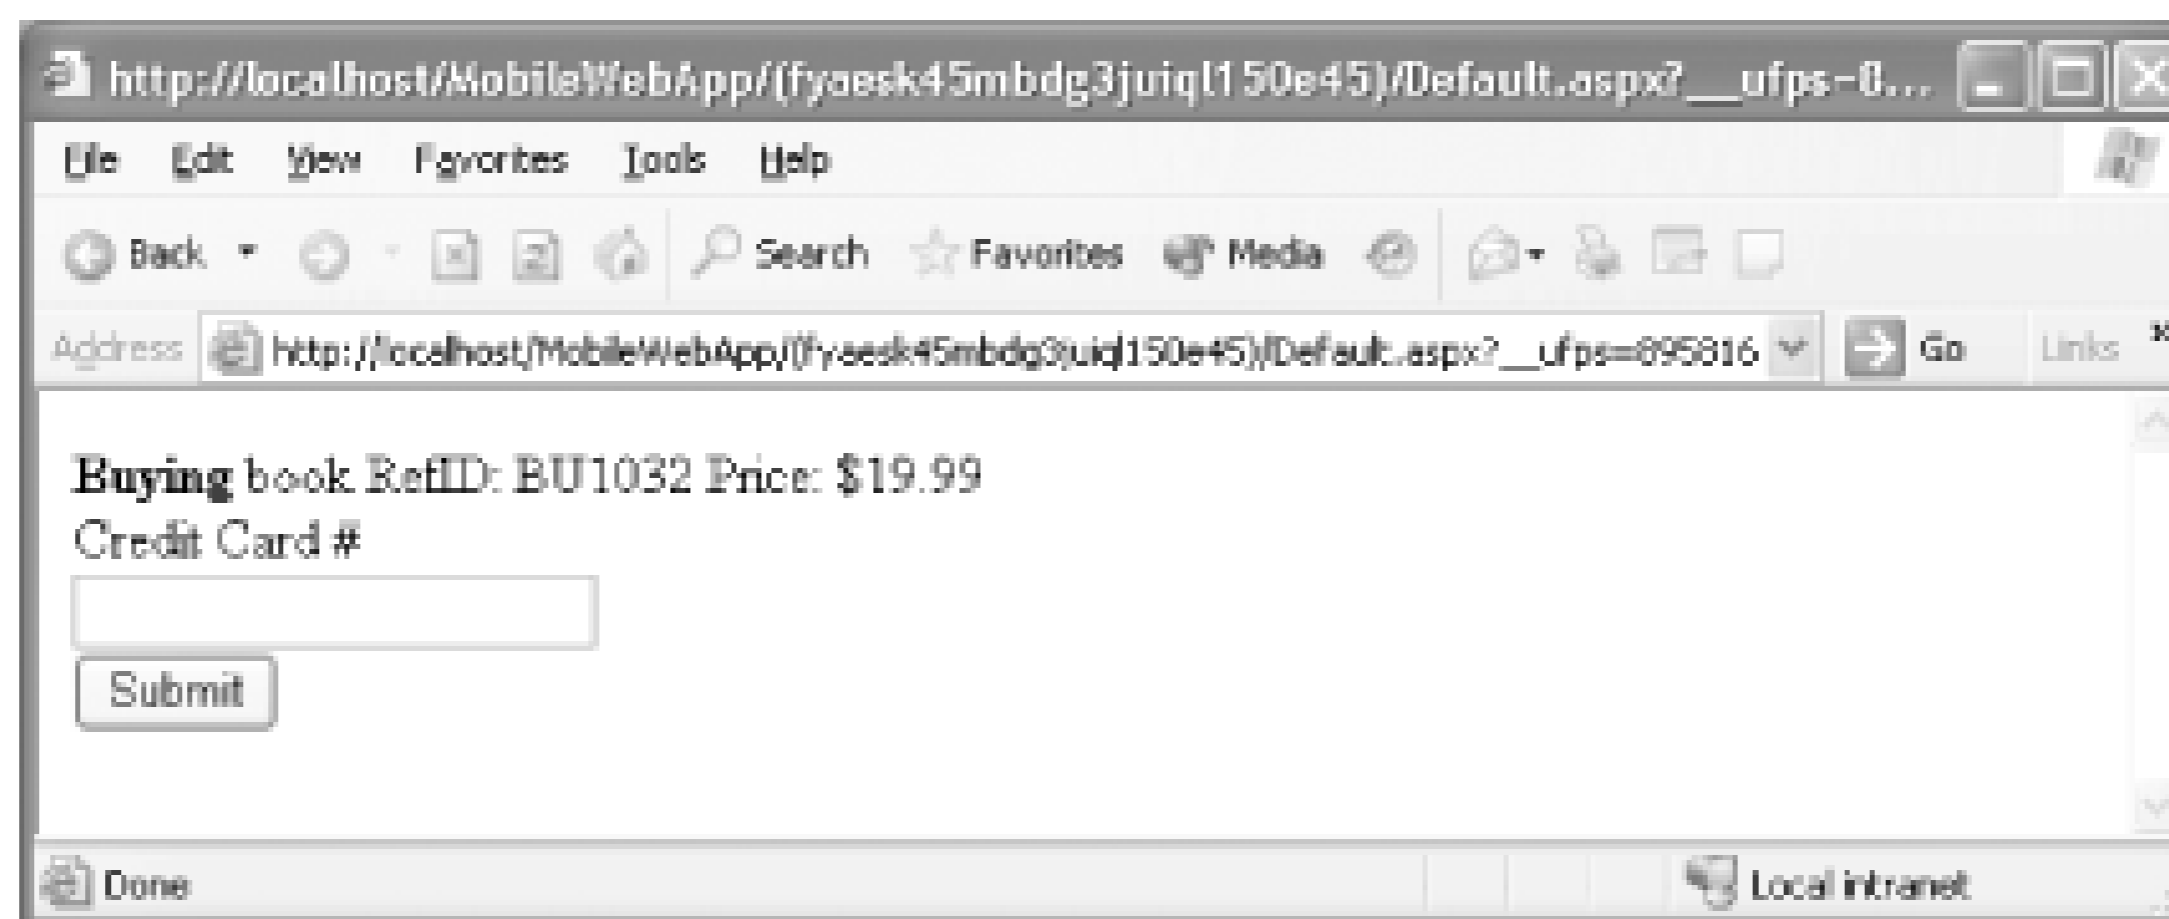
To complete this part of the application, you must write an event handler for the `ItemCommand` event of the `ObjectList` that is triggered when the user chooses to buy or download the book:

```
Private Sub bookList_ItemCommand _
    ( _
    ByVal sender As Object, _
    ByVal e As System.Web.UI.MobileControls.ObjectListCommandEventArgs _
    ) Handles bookList.ItemCommand
    Dim sTmp As String
    Dim oListItem As MobileControls.ObjectListItem = _
        CType(e.CommandArgument, MobileControls.ObjectListItem)
    If e.CommandName = "buy" Then
        sTmp = "<b>Buying</b> "
    ElseIf e.CommandName = "dnload" Then
        sTmp = "<b>Downloading</b> "
    End If
    sTmp += "book " + oListItem.Item("title_id")
    sTmp += " " + oListItem.Item("price")
    CType(FormPurchase.FindControl("txtPurchaseInfo"), _
        MobileControls.TextView).Text = sTmp
    ActiveForm = FormPurchase
End Sub
```


The ItemCommand function inspects the `CommandName` of the `ObjectList`'s event to decide which button the user clicked in order to display the appropriate form.

If the user chooses to buy the first book, the device will display the page shown in [Figure 9-3](#).

Figure 9-3. Purchase form in IE



For the submit click, we don't really process any ordering but just navigate the user back to the default page:

```
Private Sub cmdSubmit_Click _
    ( _
        ByVal sender As Object, ByVal e As System.EventArgs _
    ) Handles cmdSubmit.Click
    'More Processing Code Here
    RedirectToMobilePage("default.aspx")
End Sub
```

A primary motivation for using mobile controls in this sample application is to make it accessible to browsers other than IE. We can use the actual device to test the application or to use the emulator that comes with .NET Framework SDK (see the sidebar [Working with Microsoft Emulators](#)). [Figure 9-4](#) shows how the book list page is rendered by the Pocket PC's Internet Explorer.

Figure 9-4. Book listing on Pocket PC



This display is very similar the output from IE that we've have seen earlier because Pocket PC Internet Browser is capable of displaying HTML. [Figure 9-5](#) shows the page that is displayed when the user selects the first book in the list.

Figure 9-5. Book detail screen on Pocket PC

Finally, [Figure 9-6](#) shows the purchase form displayed when a mobile device user chooses to download the book.

Figure 9-6. Purchase form on Pocket PC



From an emulated web phone (or any WML viewer), however, the screens are drastically different. The reason is that the mobile controls sense the capability of the phone browser and only render WML instead of full blown HTML. [Figure 9-7](#) shows the same application from a phone emulator.

Figure 9-7. Book listing on mobile phone

This is the first screen showing the list of books. You can use the center rocker button to move up and down the list. The currently selected book will be highlighted. When you decided on a book, press the center button to select it. Because the screen for the phone is much smaller than the Pocket PC, the WML translation of the page content decided to show the three choices that you can pick for the selected book (View the details, buy, or download the book), as shown in [Figure 9-8](#).

Figure 9-8. Intermediate form for commands

Let's press the center button one more time to see the book detail screen. Again, this screen is scrollable via the center rocker button. [Figure 9-9](#) and [Figure 9-10](#) show the progression through the data.

Figure 9-9. Book detail screen on mobile phone

Figure 9-10. Book detail screen on mobile phone, part 2

Now that we've seen the detail for this book and it's quite interesting, we press the right button (underneath the word "Back") to go back a screen and press the down button to highlight "Buy," as in [Figure 9-11](#).

Figure 9-11. Choosing Buy option on mobile phone

When the user chooses to buy the book, the purchase form appears. Again, the user will have to use the up/down button to scroll down to select the text box to enter the credit card number, as in [Figure 9-12](#) and [Figure 9-13](#).

Figure 9-12. Purchase form on mobile phone

Figure 9-13. Purchase form on mobile phone, part 2

Once the credit card text box is selected, an input box type of screen shows up to collect the credit card number, as in [Figure 9-14](#) and [Figure 9-15](#).

Figure 9-14. Credit card field input on mobile phone

Figure 9-15. Purchase form with credit card info to be submitted on mobile phone

[\[Team LiB \]](#)

[Team LiB]

9.2 .NET Compact Framework

Working with Microsoft Emulators

Microsoft provides an emulator program that comes with VS.NET to allow developers to emulate Pocket PC 2002 and Windows CE WebPad. The command line to start up the Pocket PC emulator outside of VS.NET is:

```
In directory
C:\Program Files\Microsoft Visual Studio .NET 2003\
CompactFrameworkSDK\ConnectionManager\Bin>
Execute
emulator /CEImage Images\Pocket PC\2002\1033\pPC2002.bin
/ethernet true
/skin Images\Pocket PC\2002\1033\PPC2002.xml
```

In order to communicate with the emulated Pocket PC when you are off-line:

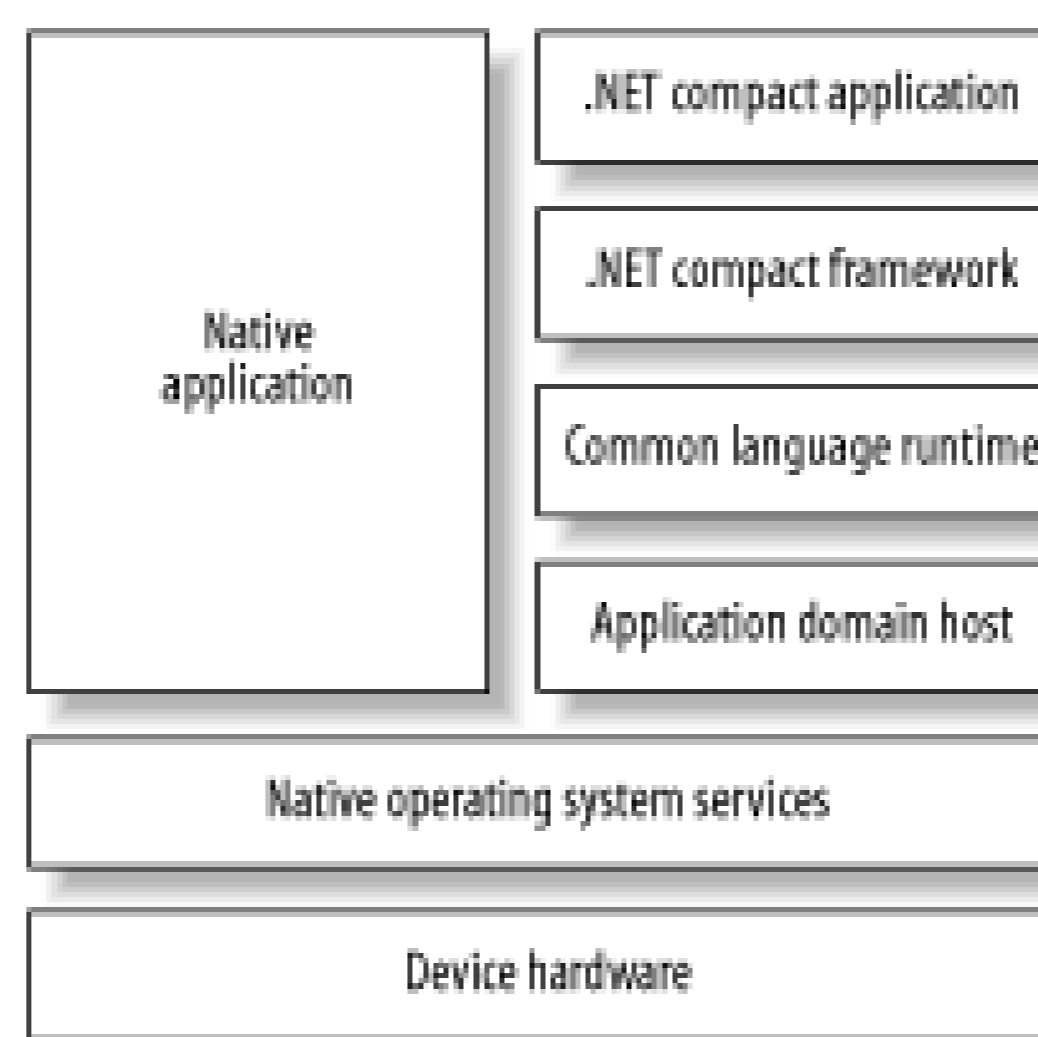
- Install the Microsoft Loopback Adapter in your network setting.
- Configure the IP for this adapter to 192.168.0.1.
- Set the subnet mask to 255.255.255.0.

Now you should be able to create share on your desktop/laptop and access it from the Pocket PC emulator.

Instead of accessing a web application from a browser, the .NET Compact Framework allows you to write rich clients and stand-alone applications that actually run on the mobile devices themselves. These applications allow the device users to be productive without being connected to the network.

Unless you like to deal with the complexity of writing code directly to the device's operating system API, .NET Compact Framework is the answer to writing mobile device applications. Being the subset of the .NET Framework, the Compact Framework provides the same common language runtime and managed code execution. Architecturally, the .NET Compact Framework acts as the abstraction layer on top of the device-specific API so that the mobile applications written on the .NET Framework can be device independent. Figure 9-16 shows the .NET Compact Framework in relation to the mobile application and the device API.

Figure 9-16. NET Compact Framework application architecture



In addition to being abstracted from the native operating system services, these mobile applications also inherit other advantages such as fault isolation, security, and better resource management, because they run within separate application domain hosts.

Currently, the .NET Compact Framework supports Pocket PC and Windows CE .NET-based platforms.

9.2.1 The .NET Framework Versus the .NET Compact Framework

Having a much smaller footprint than the .NET Framework, the .NET Compact Framework provides only the core functionality for writing applications for mobile devices. The following list summarizes the main differences between the two frameworks.

ASP.NET

ASP.NET is not supported by the .NET Compact Framework for the obvious reason that currently, the reason for running web servers on these mobile device are not compelling enough.^[5] This has nothing to do with ASP.NET application running on the servers serving ASP.NET pages to the mobile devices as we've shown you earlier in this chapter.

^[5] This does not mean that it's impossible. Maybe these devices can soon publish their services to a community of services the way JINI works. Maybe we can write a scaled-down web server that is able to serve ASP.NET applications, yet small enough to not hog all the resources of the device.

Data access

Only SQL Server CE .NET data provider is provided. OleDb or Oracle providers are not provided.

Classes

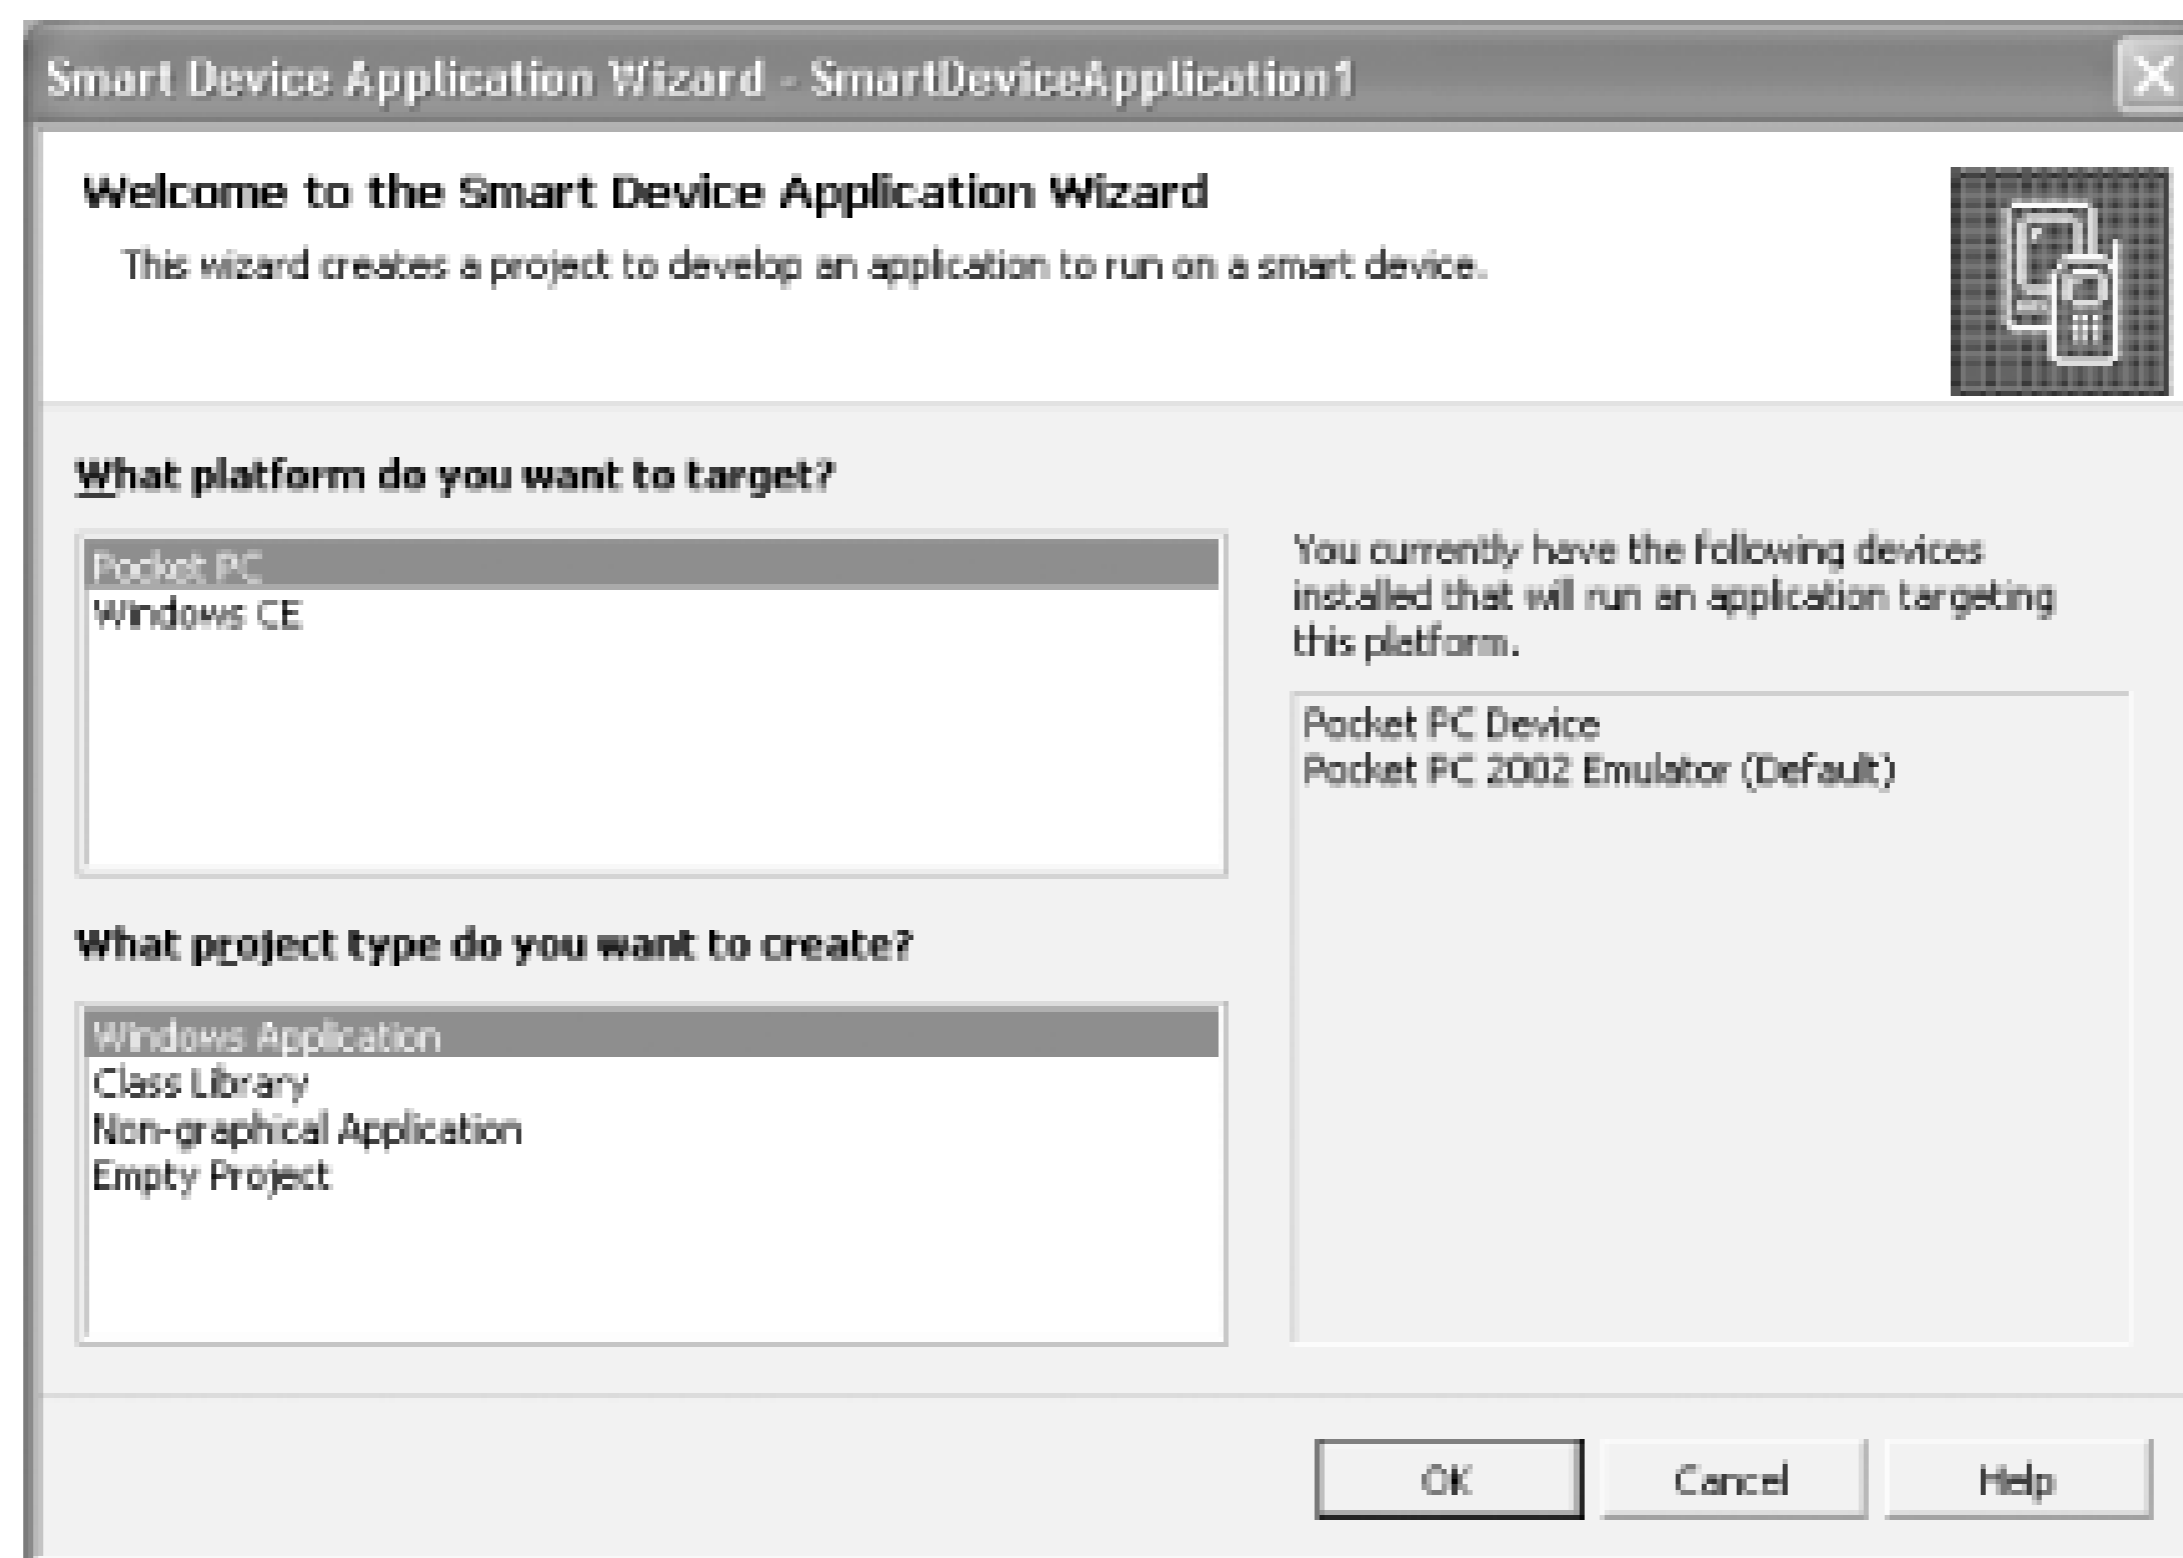
Being less than one tenth of the size of the full .NET Framework, the number of classes would obviously be smaller. We've already previously mentioned that the .NET Compact Framework only provides an essential subset of the full framework for the mobile environment.

Visual Studio .NET 2003 fully support the development and deployment of .NET applications on mobile devices.

To show how easy it is to write a mobile device application with the Compact Framework, let's write and demonstrate (using an emulator) a simple Hello World program that displays the words "Hello World" in a Windows message box on a Pocket PC.

The first thing you do is to create a New Project in Visual Studio .NET 2003. Choose the language of your choice and pick the Smart Device Application template in the New Project window. Click OK. The Smart Device Application Wizard window shown in Figure 9-17 will be displayed.

Figure 9-17. Smart device wizard



The two platforms currently supported are Pocket PC and Windows CE. For each of these platforms, there are a number of project templates you can choose from: Windows Application, Class Library, Non-graphical Application, or Empty Project.

Because you want to write a standard Windows style application that targets a Pocket PC, you will accept the Wizard's default settings by clicking OK.

VS.NET uses the Smart Device Application Template to generate all the files necessary for your application and presents you a blank form with a default menu attached to it. In this first example, you will not want to do anything with this default menu. You can also start to drag and drop controls from the toolbox onto the form and write code that associates with events for these controls the way you've always done it in previous VB environments. Again, in this example, we will not add any controls to the form but only write code to handle the `Load` event of the form and display a message box that says "Hello World!" To do this, just double click on the form itself. VS.NET creates the default handler for the load event and moves you to the code view where you can type the following VB.NET statement:

```
MessageBox.Show("Hello World!")
```

That's all to it. Now you start the debugging process by pressing the F5 key. The message shown in Figure 9-18 will appear to ask you to name the type of device to which you wish to deploy the program. We have defaulted our setting to the emulator so all we have to do at this point is click the Deploy button. VS.NET starts the emulator and deploys the application.^[6]

[6] The first time the emulator starts up, the .NET Compact Framework is installed. The .NET Compact Framework is less than 2MB in size.

Figure 9-18. Smart Device Application Deployment

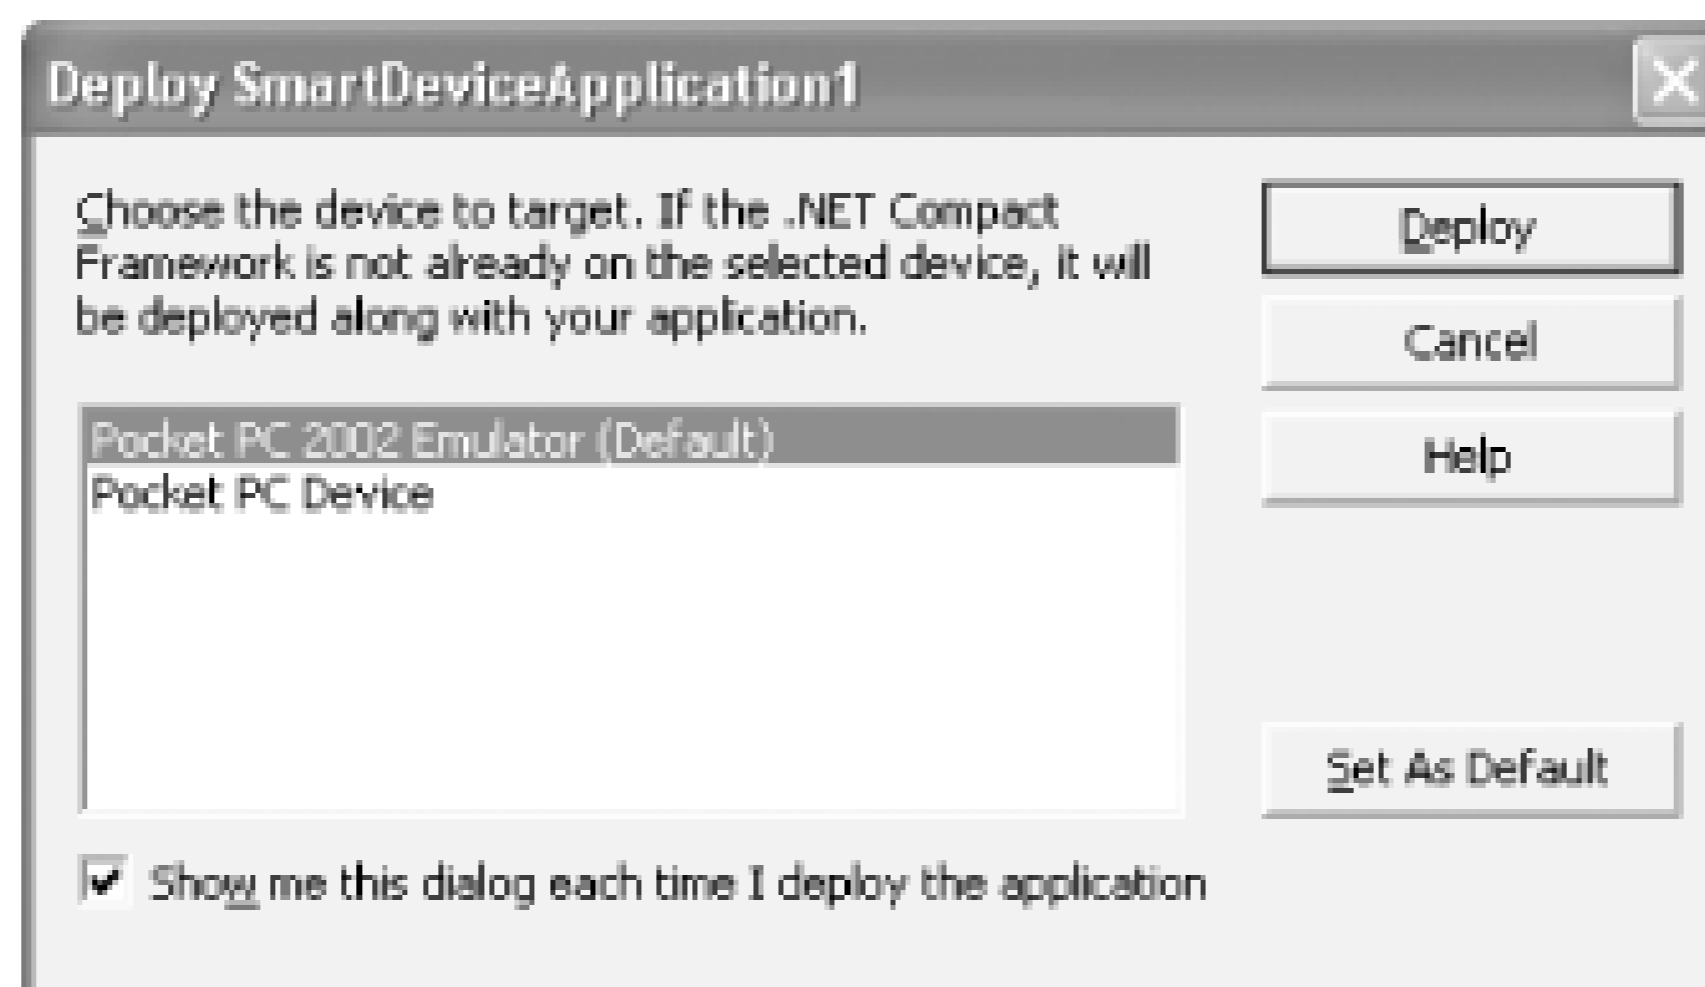


Figure 9-19 shows the output of our first Compact Framework-enabled mobile device application.

Figure 9-19. Hello World application running on Pocket PC



Now that you know the process, you're ready for something a little more involved.

9.2.2 Mobile Devices and SQL Server 2000 CE

A primary motivation for writing stand-alone applications for mobile devices (or any offline device) is because such devices are not always connected to enterprise resources or other data an application needs to do its work. Mobile device applications have to be able to extract the data they need when they are off line. Updates to the disconnected data must be merged enterprise data when it comes time for reconciliation.

Disconnected data can be stored off-line in at least two ways. Data can be stored in a proprietary format file as XML that the application manipulates using custom code or an XML parser, such as the one provided with .NET. With SQL Server 2000 CE, the off-line data can also be stored in a relational database format and manipulated via traditional database management tools.

The following two samples show a movie listing application where you can pick a movie and see where it's shown and the show times or you can pick a nearby theater and see the movies that the selected theater shows (also with show times).

The first sample shows how this is done with just XML. The XML file represents the off-line data store that is downloaded into the mobile device periodically. For simplicity, we do not show how this file is generated (let's assume that there is a Web Service out there somewhere that we can ask for local theater and showing information). The example does not have any update to the data, hence, there is no data reconciliation needed.

The second sample shows how the application can be implemented with SQL Server for CE. Here the `.sdf` file replaces the XML file, but the part of generating the `sdf` file is conveniently omitted in order to simplify the presentation.

SQL Server CE provides three methods for synchronization of data between the device and the enterprise data source: Pull, Push, and SubmitSQL. The Pull and Push methods allow your mobile device application to pull tables from and push tables to the enterprise database. Once the tables are on the mobile device, you don't need to have network connectivity in order for your application to work until you are ready to sync back to the enterprise database. The SubmitSQL method allows the mobile device application to send SQL directly to the enterprise database. This is obviously to keep the database state as synced as possible but it requires constant network connectivity.

You can also setup SQL Server replication so that the application on the mobile device can subscribe and replicate part or the whole database. The replication is similar to how the standard SQL Server Replication works except that your node is on a mobile device and the connectivity between the mobile node and the publisher is through HTTP. As intended, our examples do not cover how to set up your development environment to enable Remote Data Access or SQL Server CE Replication.

9.2.2.1 Storing off-line data as XML

In this example, we omit how the enterprise data (in this case, movie listing for a selected group of local theaters) XML is generated, and synced to the device. We start out with the assumption that the mobile device application will have the xml to consume. This application basically just parses the xml data and provides a GUI where the user queries movies or theaters information.

The XML for the movie listing is in the following format:

```
<?xml version="1.0" encoding="utf-8" ?>
<root>
  <theaters>
    <theater id="1" name="Centreville Multiplex" . . . />
    <theater id="2" name="Lee Highway Multiplex" . . . />
    < . . . >
  </theaters>
  <movies>
    <movie mid="1" name="Narc (R)" summary=" . . . " />
    <movie mid="2" name="25th Hour (R)" summary=" . . . " />
    < . . . >
  </movies>
  <xrefs>
    <xref theaterid="1">
      <movietime refmid="5"
```

```

        time="0:5 Fri - Feb 07: 8:00 12:30 | Sat - Feb 08: 12:30 | . . . " />
        <movietime refmid="2"
            time="0:4 Fri - Feb 07: 7:00 10:30 | Sat - Feb 08: 8:00 11:30 . . . " />
        < . . . >
    </xref>
    <xref theaterid="2">
        <movietime refmid="5" time=" . . . " />
        <movietime refmid="4" time=" . . . " />
        < . . . >
    </xref>
    < . . . >
</xrefs>
</root>

```

There are basically three collections of data: theaters, movies, and xrefs. As the names imply, the theaters and movies collections hold basic information for the theater and movie entities. The xrefs collection serves as the cross-reference between the theaters and the movies to hold the movie show times.

The application starts with the loading of the offline content into the menus and building appropriate data structures to store the theaters, movies, and references between them:

```

XmlDocument doc = new XmlDocument( );
doc.Load(@"\Program Files\XMLMovieListing\listing.xml");
XmlNodeReader oReader = new XmlNodeReader(doc.DocumentElement);

oReader.Read( );
while(oReader.Read( )) {
    if(oReader.Name == "theaters" && oReader.IsStartElement( )) {
        ProcessTheaters(oReader);
    } else if(oReader.Name == "movies" && oReader.IsStartElement( )) {
        ProcessMovies(oReader);
    } else if(oReader.Name == "xrefs" && oReader.IsStartElement( )) {
        ProcessCrossRef(oReader);
    }
}

```

For example, the ProcessTheaters creates the menu with theaters as items, associates the menu click with mnuTheater_Click event handler. This function also creates two hashtables: one to store general theater information keying on the theater id and the other is a hashtable that points to another hashtable that stores movie show times based on movie id for the current theater. Because this application is just an example and the list of local theaters and movies are small, it is ok to store the information in memory. Remember, when you are developing for these mobile devices, memory resource can be scarce and some alternative design should be considered:

```

private void ProcessTheaters(XmlNodeReader oReader) {
    bool bDone = false;
    string sTheaterKey = "";
    while(!bDone && oReader.Read( )) {
        if(oReader.Name == "theater" && oReader.IsStartElement( )) {
            MenuItem mnuItem = new MenuItem( );
            string sText = oReader["id"] + " " + oReader["name"];

```

```

        mnuItem.Text = sText;
        mnuItem.Click += new EventHandler(this.mnuTheater_Click);
        this.mnuTheater.MenuItems.Add(mnuItem);
        sTheaterKey = oReader["id"];
        THEATERS_TIME.Add(sTheaterKey, new Hashtable( ));
        THEATERS.Add(sTheaterKey,
            new Theater(oReader["name"], oReader["addr1"],
                oReader["addr2"], oReader["phone"]));
    }
    if(oReader.Name == "theaters" && oReader.IsStartElement( ) == false) {
        bDone = true;
    }
}
}
}

```

At this point, you probably wonder if you can just rely on the XPath query to find movie/theater associations. As it turns out, because you are using the "compact" framework, some of the functionality that you might be familiar with in the desktop/web world might not be implemented here. In this case, XPath is not implemented so you won't be able to `doSelectNodes` to simplify your life.

Similar to the `ProcessTheaters` function, the `ProcessMovies` creates the menu with movies as items, associates the menu click with `mnuMovie_Click` event handler. This function also creates two hashtables: one to store general movie information keying on the movie id and the other is a hashtable that points to another hashtable that stores movie times based on theater id for the current movie:

```

private void ProcessMovies(XmlNodeReader oReader) {
    bool bDone = false;
    while(!bDone && oReader.Read( )) {
        if(oReader.Name == "movie" && oReader.IsStartElement( )) {
            MenuItem mnuItem = new MenuItem( );
            string sText = oReader["mid"] + " " + oReader["name"];
            mnuItem.Text = sText;
            mnuItem.Click += new EventHandler(this.mnuMovie_Click);
            this.mnuMovie.MenuItems.Add(mnuItem);
            string sMovieKey = oReader["mid"];
            MOVIES_TIME.Add(sMovieKey, new Hashtable( ));
            MOVIES.Add(sMovieKey,
                new Movie(oReader["name"], oReader["summary"]));
        }
        if(oReader.Name == "movies" && oReader.IsStartElement( ) == false) {
            bDone = true;
        }
    }
}
}
}

```

The `ProcessCrossRef` function fills in the two cross-ref structure so that we can list showing times for a particular movie across all theater or showing times for all movies at a particular theater:

```

private void ProcessCrossRef(XmlNodeReader oReader) {
    bool bDone = false;

```

```

string sTheaterKey = "";
while(!bDone && oReader.Read( )) {
    if(oReader.Name == "xref" && oReader.IsStartElement( )) {
        sTheaterKey = oReader["theaterid"];
    }
    if(oReader.Name == "movietime" && oReader.IsStartElement( )) {
        Hashtable o = (Hashtable)THEATERS_TIME[sTheaterKey];
        o.Add(oReader["refmid"], oReader["time"]);
        Hashtable o1 = (Hashtable)MOVIES_TIME[oReader["refmid"]];
        o1.Add(sTheaterKey, oReader["time"]);
    }
    if(oReader.Name == "xrefs" && oReader.IsStartElement( ) == false) {
        bDone = true;
    }
}
}
}

```

The menu click handler: `mnuTheater_Click` and `mnuMovie_Click` perform similar tasks, which basically setup the current mode of selection (by theater or by movie), the current theater or movie based on the mode, the list of movie and show times if the mode is by theater, or the list of theater and show times if the mode is by movie. Other supporting functions help in navigating to show movies and show times by theater or theater and show times by movie.^[7]

^[7] Download the complete code online at <http://www.oreilly.com/catalog/dotnetfrm3/> .

Figure 9-20 and Figure 9-21 show the application running on a Pocket PC while having the Movies and Theaters menu expanded.

Figure 9-20. XML movie listing by movie

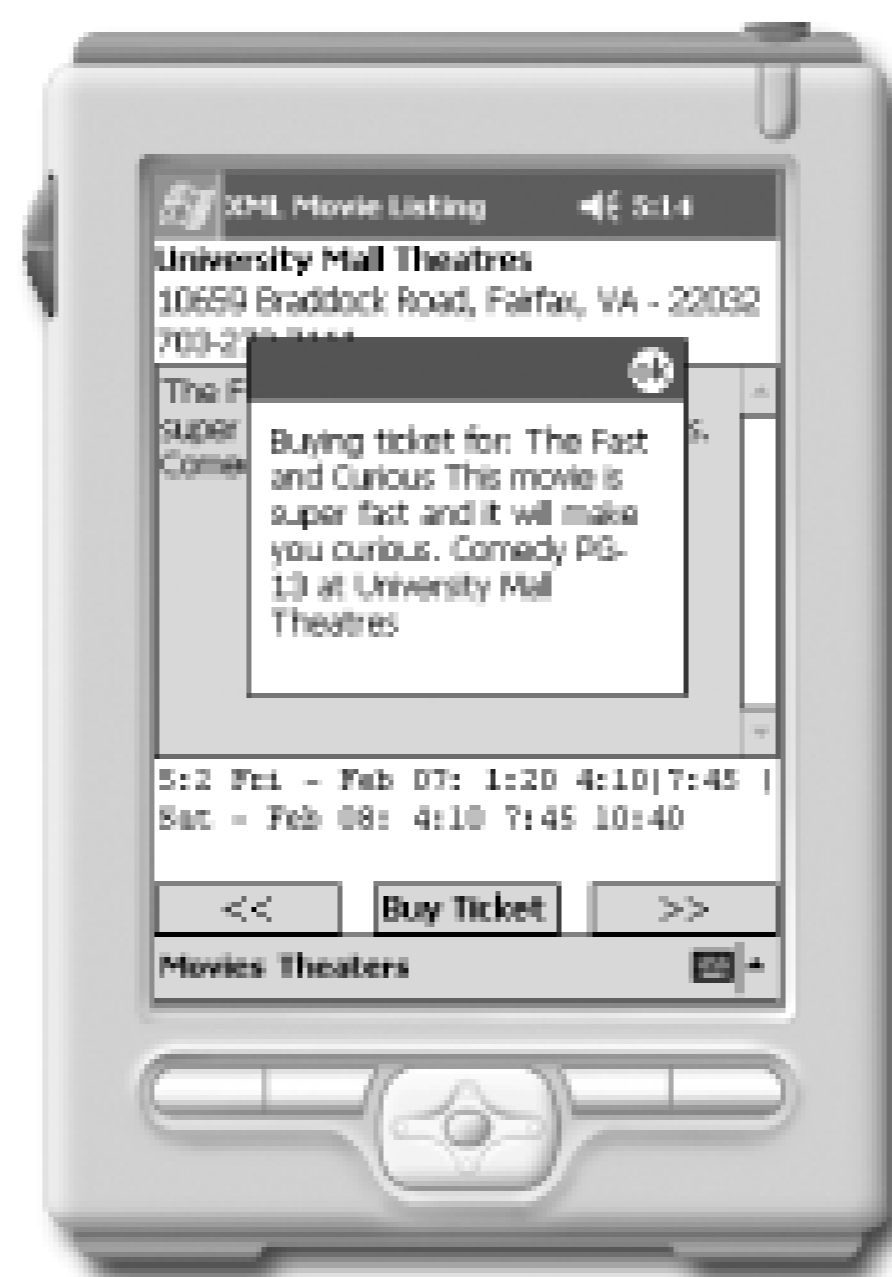
Figure 9-21. XML movie listing by theater



When you select a movie, the application traverses all theaters that show the selected movie and display the show times (see Figure 9-22 and Figure 9-23).

Figure 9-22. By movie

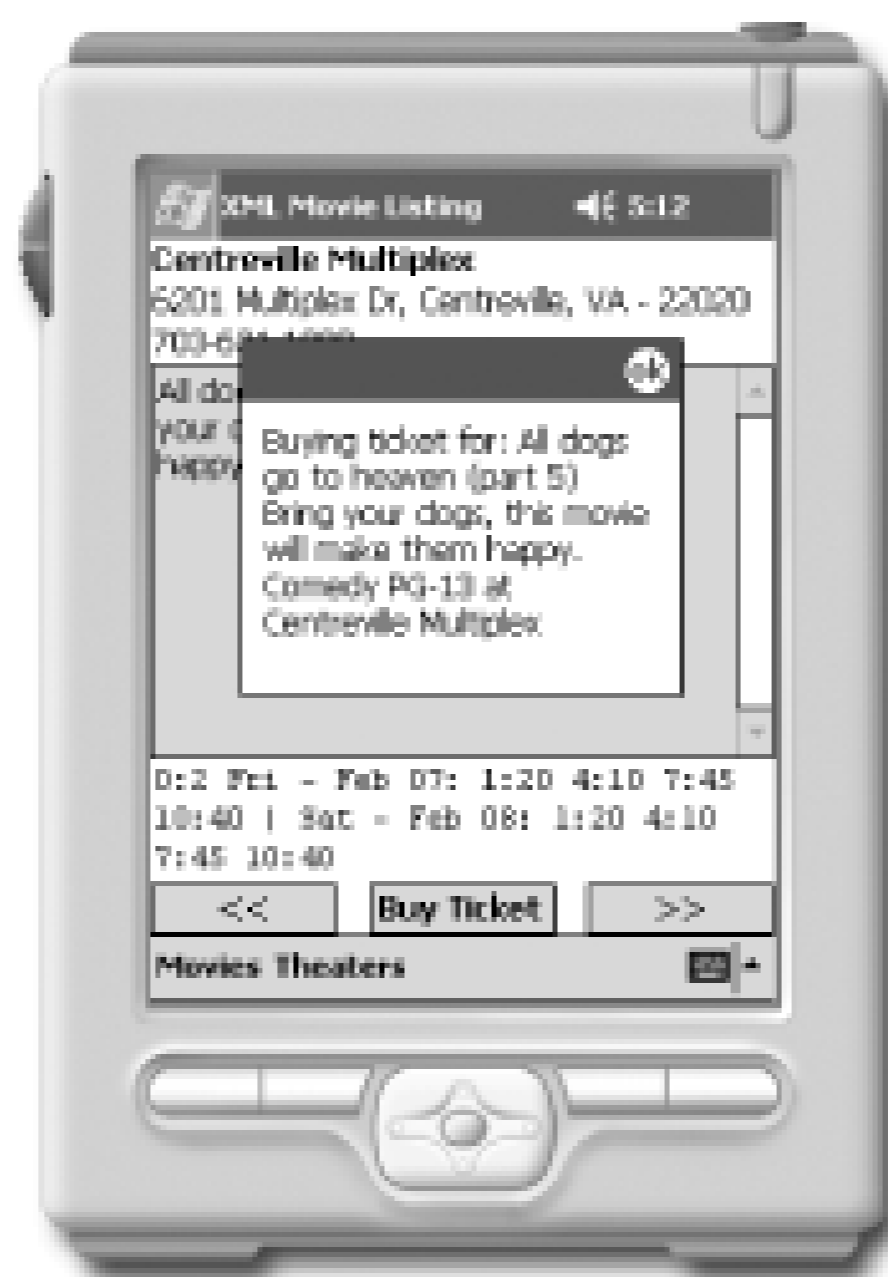
Figure 9-23. Same movie, other theater, also buying ticket



When you pick a theater instead of a movie, the application starts to traverse all movies showing at the selected theater (see Figure 9-24 and Figure 9-25).

Figure 9-24. By theater

Figure 9-25. Same theater, another movie, also buying ticket



9.2.2.2 Storing off-line data in SQL Server for Windows CE

In the next example, you will rely on SQL Server CE for the data instead of an XML file. This change simplifies the code you write but requires the installation of SQL Server CE on the mobile device, thus the total amount of memory used might be higher than the xml version of the application. We can ignore the resource problem and exchange it with the simplification of the code and, fortunately, VS.NET automatically deploys SQL Server CE on the device the same way it does the .NET Compact Framework.

This version of the application starts by opening a database connection and generating the Theater and Movie menus. Unlike the previous version, you don't have to build your own data structure to hold the data and cross references. The tables and their references are managed by the database in this example:

```
m_oConn = new SqlCeConnection(
    "Data Source=\\Program Files\\SQLCEMovieListing\\moviedb.sdf"
);
m_oConn.Open( );
GenerateMenus("select theaterid, name from theater",
    new EventHandler(this.mnuTheater_Click),
    this.mnuTheater);
GenerateMenus("select movieid, name from movie",
    new EventHandler(this.mnuMovie_Click),
    this.mnuMovie);
```

The GenerateMenus function just creates the menu items and associates the event handler to each item:

```
private void GenerateMenus(string sCommand,
    EventHandler handler,
    MenuItem mnuBranch) {
    SqlCeCommand oCmd = new SqlCeCommand( );
    SqlCeDataReader oReader = null;
    oCmd.Connection = m_oConn;
    oCmd.CommandText = sCommand;
```

```

oReader = oCmd.ExecuteReader( );
while(oReader.Read( )) {
    MenuItem mnuItem = new MenuItem( );
    mnuItem.Text = oReader["name"].ToString( );
    mnuItem.Click += handler;
    mnuBranch.MenuItems.Add(mnuItem);
}
oReader.Close( );
}

```

Here are the handlers for Theater and Movie:

```

private void mnuMovie_Click(object o, System.EventArgs e) {
    // Get information from the menu item.
    string sName = ProcessMenuClick((MenuItem)o);
    // Obtain the current movie information.
    SqlCeCommand oCmd = new SqlCeCommand(
        "select summary from movie where name = '" + sName + "'"
        );

    oCmd.Connection = m_oConn;
    SqlCeDataReader oReader = oCmd.ExecuteReader( );
    oReader.Read( );
    m_oCurrentMovie = new Movie(sName, oReader["summary"].ToString( ));
    // Obtaining all theaters and their show times for this movie
    SqlCeDataAdapter oAdapter = new SqlCeDataAdapter
    (
        @"
select t.name, t.addr1, t.addr2, t.phone, st.showtime
from theater t inner join showtime st on t.theaterid = st.theaterid
inner join movie m on m.movieid = st.movieid
where m.name = '" + sName + "'",
        m_oConn
    );
    m_oDS.Clear( );
    oAdapter.Fill(m_oDS);
    // Start in-memory cursor.
    m_iMode = 0;
    m_iCurrentIndex = 0;
    displayCurrent( );
}

private void mnuTheater_Click(object o, System.EventArgs e) {
    // Get information from the menu item.
    string sName = ProcessMenuClick((MenuItem)o);
    // Obtain the current theater information.
    SqlCeCommand oCmd = new SqlCeCommand(
        "select addr1, addr2, phone from theater where name = '" + sName + "'"
        );

    oCmd.Connection = m_oConn;
    SqlCeDataReader oReader = oCmd.ExecuteReader( );
    oReader.Read( );
    m_oCurrentTheater = new Theater(sName,

```

```
oReader["addr1"].ToString( ),
oReader["addr2"].ToString( ),
oReader["phone"].ToString( ));

// Obtain all movie and show times at this theater.
SqlCeDataAdapter oAdapter = new SqlCeDataAdapter
(
@"
select m.name, m.summary, st.showtime
from movie m inner join showtime st on m.movieid = st.movieid
inner join theater t on t.theaterid = st.theaterid
where t.name = '" + sName + "'",
m_oConn
);
m_oDS.Clear( );
oAdapter.Fill(m_oDS);
// Start in-memory cursor.
m_iMode = 1;
m_iCurrentIndex = 0;
displayCurrent( );
}
```

The rest of the code can be found at <http://www.oreilly.com/catalog/dotnetfrm3/>.

[Team LiB]

[\[Team LiB \]](#)

9.3 Summary

As you can see, writing applications for mobile devices can now be fun and productive. Whether you are going to write Web Application targeting mobile devices or applications that actually run on these devices, .NET would have you well covered. The Mobile Controls as well as the Compact Framework are now part of the standard .NET 1.1 to ensure that all aspects of the enterprise are taken care of. Thanks to great tools such as VS.NET, application developers no longer worry about the tedious legwork of setting up the environment, debugging tools, the SDK, emulators, and so on, and devote all their time to developing the business logic for the mobile application.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Appendix A. .NET Languages

This appendix contains two lists of languages (Microsoft-supported and third party) with compilers that generate IL to target the CLR.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

A.1 Microsoft-Supported Languages for .NET

[Table A-1](#) lists commercial languages that Microsoft supports. You can find more information about each of these languages by browsing the provided URLs.

Table A-1. Microsoft-supported languages for .NET

Language	Link
Visual C# .NET	http://msdn.microsoft.com/vcsharp/
Visual J# .NET (Java)	http://msdn.microsoft.com/vjsharp/
Visual C++ .NET (Managed C++)	http://msdn.microsoft.com/visualc/
Visual Basic .NET	http://msdn.microsoft.com/vbasic/
JScript for the CLR	http://msdn.microsoft.com/workshop/languages/clinic/scripting07142000.asp

[\[Team LiB \]](#)

[\[Team LiB \]](#)

A.2 Third-Party Languages for .NET

[Table A-2](#) shows a list of third-party languages with compilers that target the CLR. Some of these are research languages, while others are commercial languages that target .NET. Browse the provided web sites to read more about the languages that interest you. As noted earlier, this list of languages could grow by the time this book hits the market, so be sure to check the following site for the most up-to-date listings: <http://www.gotdotnet.com>.

Table A-2. Third-party languages

Language	Link
APL	http://www.dyadic.com
COBOL	http://www.adtools.com/info/whitepaper/net.html
Component Pascal	http://www.citi.qut.edu.au/research/plas/projects/component_pascal.jsp
Delta Forth	http://www.dataman.ro/dforth/
Eiffel#	http://www.eiffel.com/doc/manuals/technology/dotnet/eiffelsharp/white_paper.html
Fortran	http://www.lahey.com/dotnet.htm , http://www.salfordsoftware.co.uk/compilers/ftn95/dotnet.shtml
Haskell	http://haskell.cs.yale.edu/ghc
Mercury	http://www.cs.mu.oz.au/research/mercury/dotnet.html
Mondrian	http://www.mondrian-script.org
Oberon	http://www.oberon.ethz.ch/lightning
Perl	http://www.activestate.com/ASPNET
Python	http://www.activestate.com/ASPNET
RPG	http://www.asna.com/pr2%5F20%5F01.asp
Scheme	http://rover.cs.nwu.edu/~scheme
Smalltalk	http://www.lessor-software.com/en/content/products/lswst/LSWVSt.htm
S#	http://www.smallscript.net
Standard ML	http://www.research.microsoft.com/Projects/SML.NET
TMT Pascal	http://www.tmt.com/net.htm

For more information, visit the O'Reilly .NET Center at <http://dotnet.oreilly.com> and the .NET DevCenter at <http://www.oreillynet.com/dotnet>.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Appendix B. Common Acronyms

[Table B-1](#) provides a listing of common acronyms or terms that you will come across in .NET-related reading materials and conversations. Some of these acronyms or terms have little relevance to .NET, but appear occasionally in this book, so we've provided them in this table for your convenience.

Table B-1. List of common acronyms

Acronym	Description
.ASMX	File extension for web services source-code files.
.ASPX	File extension for ASP.NET source-code files.
.config	File extension of .NET application configuration files.
ADO	ActiveX Data Objects.
ADO.NET	With goals similar to ADO, ADO.NET is the way to access data sources in .NET.
API	Application Programming Interface.
AppDomain	Short term to mean an application domain.
ASP	Active Server Pages.
ASP.NET	With goals similar to ASP, ASP.NET is the technology that supports the rapid development of Web Forms.
ATL	Active Template Library.
BLOB	Binary Large Object.
CAB Files	Cabinet files.
CCW	COM Callable Wrapper.
CLI	Common Language Infrastructure. This is the specification of the infrastructure and base-class libraries that Microsoft has submitted to ECMA so that a third-party vendor can build a .NET runtime on another platform. The CLR is Microsoft's implementation of the CLI.
CLR	Common Language Runtime.
CLS	Common Language Specification.
CLSID	Class identifier used in COM.
COFF	Common Object File Format.

Acronym	Description
COM	Component Object Model.
COM Interop	Short for COM interoperation.
COM+ 2.0	This term is no longer used as of Beta 1 of the .NET SDK. The new, accepted term is .NET Framework.
COM+ Runtime	This term is no longer used as of Beta 1 of the .NET SDK. The new, accepted term is Common Language Runtime.
CTS	Common Type System.
DB	Database.
DCOM	Distributed Component Object Model.
DHTML	Dynamic HyperText Markup Language.
DISCO	Discovery of web services. A Web Service has one or more .DISCO files that contain information on how to access its WSDL.
DISPID	Dispatch identifier. Used in COM to identify a method or a property for dynamic invocation.
DLL	Dynamically Linked Library.
DNA	Distributed interNet Applications Architecture.
DOM	Document Object Model.
DTD	Data Type Document. This has been superceded by XSD.
EXE	Executable.
GC	Garbage Collector.
GDI	Graphical Device Interface.
GDI+	A .NET library that supports advanced graphics management.
Global.asax	The global configuration file for an ASP.NET application.
GUID	Globally Unique Identifier.
HTML	HyperText Markup Language.
HTTP	HyperText Transfer Protocol.
IDE	Integrated Development Environment.
IDL	Interface Definition Language.
IE	Internet Explorer.
IID	Interface Identifier.
IIS	Internet Information Server.
IJW	It Just Works.

Acronym	Description
IL	Intermediate Language.
ILDASM	Intermediate Language Disassembler.
Inproc	In-Process.
ISAPI	Information Server Application Programming Interface.
Machine.config	Configuration file for administrative policy for an entire machine.
MBR	Marshal-By-Reference.
MBV	Marshal-By-Value.
MFC	Microsoft Foundation Classes.
MSI	Microsoft Windows Installer Package.
MSIL	Microsoft Intermediate Language.
MSVCRT	Microsoft Visual C++ Runtime.
MSXML	Microsoft eXtensible Markup Language.
MTS	Microsoft Transaction Server.
NGWS	Stands for Next Generation Windows Services, the original name for .NET before it was previewed at PDC 2000 in Orlando, Florida. This term is no longer used.
NTFS	NT File System.
N-Tier	Multi-tier.
NTLM	NT Lan Manager.
OBJREF	Object Reference.
Out of proc	Out-of-process.
P/Invoke	Platform Invoke.
PE	Portable Executable.
perm	Permissions.
RAD	Rapid Application Development.
RCW	Runtime Callable Wrapper.
REGASM	Register Assembly tool.
RPC	Remote Procedure Call.
SCL	SOAP Contract Language.
SDK	Software Development Kit.
SEH	Structured Exception Handling.
SMTP	Simple Mail Transfer Protocol.

Acronym	Description
SOAP	Simple Object Access Protocol.
SQL	Structured Query Language.
STL	Standard Template Library.
TCP	Transport Control Protocol.
TLB	Type Library.
TLBEXP	Type Library Exporter Tool.
TLBIMP	Type Library Importer Tool.
UDDI	Universal Description, Discovery, and Integration Service. UDDI is a platform-independent framework for describing and discovering web services.
UDF	Uniform Data Format.
UI	User Interface.
URI	Uniform Resource Identifier.
URL	Uniform Resource Locator.
URT	Universal Runtime. This term is no longer used as of Beta 1 of the .NET SDK. The new, accepted term is the .NET Framework.
VB	Visual Basic.
VBRUN	Visual Basic Runtime.
VES	Virtual Execution System. The VES is a subset of the CLR. The VES doesn't include features such as debugging, profile, and COM interoperation.
VOS	Virtual Object System. This is now called the CTS.
VS.NET	Visual Studio .NET.
WAP	Wireless Application Protocol.
web.config	Configuration file for ASP.NET. You define HTTP modules, handlers, session state management, and other ASP.NET configurable parameters in this file.
WebForms	This term is no longer used as of Beta 1 of the .NET SDK. The new and accepted term is Web Forms (with a space).
WebServices	This term is no longer used as of Beta 1 of the .NET SDK. The new, accepted term is web services (with a space).
Win32	Windows 32-bit.
WinForms	This term is no longer used as of Beta 1 of the .NET SDK.
WML	Wireless Markup Language.
WSDL	Web Service Description Language. Think of this as IDL for web services. Unlike IDL, WSDL is expressed using only XML schemas. SDL is used in Beta1 of the .NET SDK, but WSDL replaced SDL in Beta 2 and later installments.

Acronym	Description
XML	Extensible Markup Language.
XPath	XML Path.
XSD	XML Schema Definition.
XSL	Extensible Stylesheet Language.
XSLT	Extensible Stylesheet Language Transformations.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Appendix C. Common Data Types

Each of the .NET languages can provide its own keywords for the types it supports. For example, a keyword for an integer in VB is `Integer`, whereas in C# or C++ it is `int`; a boolean is `Boolean` in VB, but `bool` in C# or C++. In any case, the integer is mapped to the class `Int32`, and the boolean is mapped to the class `Boolean` in the `System` namespace. Table C-1 lists all simple data types common to the .NET Framework. Non-CLS-compliant types are not guaranteed to interoperate with all CLS-compliant languages.

Table C-1. Common data types

Type	Description
Boolean	True or false.
Byte	8-bit unsigned integer: 0 to 255.
Char	Character. Unicode 16-bit character.
DateTime	Represents a date and time value.
Decimal	Can represent positive and negative values with 28 significant digits: -79,228,162,514,264,337,593,543,950,335 to 79,228,162,514,264,337,593,543,950,335.
Double	Stores 64-bit floating-point values: -1.79769313486231570e308 to 1.79769313486231570e308.
Guid	Represents a globally unique identifier (GUID); this is stored internally as a 128-bit integer. Commonly represented as a series of lowercase hexadecimal digits in groups of 8, 4, 4, 4, and 12 digits and separated by hyphens (e.g., 382c74c3-721d-4f34-80e5-57657b6cbc27).
Int16	Stores 16-bit signed integers: -32,768 to 32,767.
Int32	Stores 32-bit signed integers: -2,147,483,648 to 2,147,483,647.
Int64	Stores 64-bit signed integers: -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807.
SByte	Represents an 8-bit signed integer. The SByte type is not CLS-compliant. -128 to 127.
Single	Represents an IEEE 754f, single precision, 32-bit value: -3.40282346638528859e38 to 3.40282346638528859e38.
String	Represents a string of Unicode characters.

Type	Description
UInt16	Represents a 16-bit unsigned integer. The UInt16 type is not CLS- compliant. 0 to 65,535.
UInt32	Represents a 32-bit unsigned integer. The UInt32 type is not CLS- compliant. 0 to 4,294,967,295
UInt64	Represents a 64-bit unsigned integer. The UInt64 type is not CLS- compliant. The UInt64 data type can represent positive integers with 18 significant digits: 0 to 184,467,440,737,095,551,615.
Void	Void.

[Table C-2](#) shows a number of useful container types that the .NET Framework provides.

Table C-2. Container types

Type	Description
Array	General array construct.
ArrayList	This class implements the IList interface. The array can grow or shrink dynamically in size.
BitArray	This class represents a compact array of bit values. Each element represents a Boolean value (true/false).
HashTable	This class represents a collection of associated keys and values that are organized based on the hash code of the key.
Queue	This class represents a first-in, first-out collection construct.
SortedList	This class is similar to HashTable except that all elements are sorted by their actual keys (not hashed) and elements are accessible through either key or index.
Stack	This class represents a first-in, last-out stack construct.

[[Team LiB](#)]

[\[Team LiB \]](#)

C.1 Usage

This section demonstrates how you can take advantage of container types. We don't illustrate all methods and properties, but we show the important characteristics of these types. All examples in this chapter are in C#; however, you can use these CLS types from any other CLS-compliant languages.

C.1.1 Array

Array, by definition, stores homogeneous information in a structure that allows for indexing, enumerating, and so on. It is very important in the daily programmer's life. In the following code listing, we demonstrate some of the syntax for using Array. In C#:

```
using System;
public class TestArray {

    public static void Main( ) {

        string[] strArray = new string[7] { "Monday", "Tuesday", "Wednesday",
"Thursday", "Friday", "Saturday", "Sunday" };
        for(int i=0; i<strArray.Length; i++) {
            Console.WriteLine(String.Format("item at {0} is {1}",
                i, strArray[i]));
        }

        int[] iArray;
        iArray = new int[7];

        for(int i=0; i<iArray.Length; i++) {
            iArray[i] = i;
        }
        for(int i=0; i<iArray.Length; i++) {
            Console.WriteLine(iArray[i]);
        }

    }
}
```

and in VB:

```
Imports System
Public Class TestArray

    Public Shared Sub Main( )

        Dim i as Integer
```



```

        Dim strArray( ) as String = { "Monday", "Tuesday", "Wednesday",
"Thursday", "Friday", "Saturday", "Sunday" }

        For i = 0 To strArray.Length - 1
            Console.WriteLine(String.Format("item at {0} is {1}", _
                i, strArray(i)))
        Next

        Dim iArray(6) as Integer

        For i = 0 To iArray.Length - 1
            iArray(i) = i
        Next
        For i = 0 to iArray.Length - 1
            Console.WriteLine(iArray(i))
        Next

    End Sub
End Class

```

Notice the differences in the declaration of the array for the highlighted code. In VB, the number represents the upper bound of the array, not the size.

C.1.2 ArrayList

In the following code listing, we demonstrate some of the critical usages of the ArrayList class, such as adding data to the end of the list, inserting data anywhere in the list, iterating through the list, and sorting the list:

```

using System;
using System.Collections;
public class TestArrayList {

    public static void Main( ) {

        ArrayList arrList = new ArrayList( );
        arrList.Add("Monday");
        arrList.Add("Tuesday");
        arrList.Add("Wednesday");
        arrList.Add("Thursday");

        // We'll try to insert Friday afterward.
        // arrList.Add("Friday");

        arrList.Add("Saturday");
        arrList.Add("Sunday");

        int i = 0;
        IEnumerator arrIterator = arrList.GetEnumerator( );
        Console.WriteLine("There are: {0} days in a week.", arrList.Count);
    }
}

```

```

while(arrIterator.MoveNext( )) {
    Console.WriteLine("[{0}] {1}", i++, arrIterator.Current);
}

Console.WriteLine("Insert Friday");
arrList.Insert(4, "Friday");

i = 0;
arrIterator = arrList.GetEnumerator( );
Console.WriteLine("There are: {0} days in a week.", arrList.Count);
while(arrIterator.MoveNext( )) {
    Console.WriteLine("[{0}] {1}", i++, arrIterator.Current);
}

arrList.Sort( );

i = 0;
arrIterator = arrList.GetEnumerator( );
Console.WriteLine("Sorted as text");
while(arrIterator.MoveNext( )) {
    Console.WriteLine("[{0}] {1}", i++, arrIterator.Current);
}

Object oDay = "Friday";
Console.WriteLine("Index for Friday using BinarySearch: {0}",
    arrList.BinarySearch(oDay));
Console.WriteLine("Index for Sunday using BinarySearch: {0}",
    arrList.BinarySearch("Sunday"));
}
}

```

C.1.3 BitArray

The sample code for BitArray is self-explanatory, as shown in the following code listing. We use the bit array to store and retrieve access rights in the following example. You can use the Set and Get methods as well as the [] operator:

```

using System;
using System.Collections;
public class TestBitArray {

    enum Permissions {canRead, canWrite, canCreate, canDestroy};

    public static void Main( ) {

        BitArray bitArr = new BitArray(4);
        bitArr.Set((int)Permissions.canRead, true);
        bitArr[(int)Permissions.canWrite] = false;
        bitArr[(int)Permissions.canCreate] = true;
        bitArr[(int)Permissions.canDestroy] = false;
    }
}

```

```

        Console.WriteLine("bitArr count: {0}\tlength: {1}",
            bitArr.Count,
            bitArr.Length);

Console.WriteLine("Permissions:");
Console.WriteLine("Read: {0}",
    bitArr[(int)Permissions.canRead]);
Console.WriteLine("Write: {0}",
    bitArr[(int)Permissions.canWrite]);
Console.WriteLine("Create: {0}",
    bitArr[(int)Permissions.canCreate]);
Console.WriteLine("Destroy: {0}",
    bitArr[(int)Permissions.canDestroy]);
    }
}

```

C.1.4 HashTable

The HashTable data type is similar to the dictionary object, which is basically an associated array. Each element stored in the table is associated with a key. Because HashTable implements the IDictionaryEnumerator, we can obtain the enumerator to help us iterate through the data collection. As you can see from the sample code, we can also loop through the data using the `keys` or `values` collection:

```

using System;
using System.Collections;
public class TestHashtable {

    public static void Main( ) {

        Hashtable hashTbl = new Hashtable( );
        hashTbl.Add("Param1", "UserName");
        hashTbl.Add("Param2", "Password");

        IDictionaryEnumerator hashEnumerator = hashTbl.GetEnumerator( );
        Console.WriteLine( );
        Console.WriteLine("Loop through with enumerator:");
        while (hashEnumerator.MoveNext( )) {
            Console.WriteLine("Key: {0}\tValue: {1}",
                hashEnumerator.Key,
                hashEnumerator.Value);
        }

        Console.WriteLine( );
        Console.WriteLine("Loop through Keys:");
        foreach(string key in hashTbl.Keys) {
            Console.WriteLine(key);
        }

        Console.WriteLine( );
        Console.WriteLine("Loop through Values:");
    }
}

```

```

        foreach(string val in hashTbl.Values) {
            Console.WriteLine(val);
        }

        Console.WriteLine( );
        Console.WriteLine("Loop through Keys:");
        foreach(string key in hashTbl.Keys) {
            Console.WriteLine("Key: {0}\tValue: {1}", key, hashTbl[key]);
        }
    }
}

```

C.1.5 Queue

To demonstrate the use of a queue Abstract Data Type (ADT), we created a fictitious order-processing code listing. Each enqueued item represents a line item in a typical order. We will then dequeue each line item and perform the total calculation:

```

using System;
using System.Collections;

public class TestQueue {

    public static void Main( ) {
        string sLineItem1, sLineItem2;
        Queue myQueue = new Queue( );

        sLineItem1 = "123\tItem 123\t4\t3.39";
        sLineItem2 = "ABC\tItem ABC\t1\t9.49";
        myQueue.Enqueue(sLineItem1);
        myQueue.Enqueue(sLineItem2);

        Console.WriteLine("\nProcessing Order:\n");
        String sLineItem = "";
        String [] lineItemArr;
        Decimal total = 0;
        while(myQueue.Count > 0) {
            sLineItem = (String)myQueue.Dequeue( );
            Console.WriteLine( "\t{0}", sLineItem);
            lineItemArr = sLineItem.Split(new Char[] { '\t' });
            total += Convert.ToInt16(lineItemArr[2]) *
                Convert.ToDecimal(lineItemArr[3]);
        }
        Console.WriteLine("\nOrder Total: {0}\n", total);
    }
}

```

C.1.6 SortedList

The following code demonstrates the sorted list ADT. A sorted list is similar to a hash table or a

dictionary type. Each item of data is associated with the key with which the list is sorted. Notice that the strings are added to the list in no particular order. However, when we iterate through the list, all strings are sorted by their associated keys:

```
using System;
using System.Collections;
public class TestSortedList {

    public static void Main( ) {

        SortedList mySortedList = new SortedList( );
        mySortedList.Add("AA", "Hello");
        mySortedList.Add("AC", "!");
        mySortedList.Add("AB", "World");

        Console.WriteLine("\nLoop through manually:\n");
        for(int i=0; i< mySortedList.Count; i++) {
            Console.WriteLine("Key: {0}\tValue: {1}",
                mySortedList.GetKey(i),
                mySortedList.GetByIndex(i));
        }

        IDictionaryEnumerator myIterator = mySortedList.GetEnumerator( );
        Console.WriteLine("\nLoop through with enumerator:\n");
        while (myIterator.MoveNext( )) {
            Console.WriteLine("Key: {0}\tValue: {1}",
                myIterator.Key,
                myIterator.Value);
        }
    }
}
```

C.1.7 Stack


The following code demonstrates the first-in, last-out characteristics of the stack abstract data type. The output from the `pop` operation initially shows the fourth item, the third item, and so on:

```
using System;
using System.Collections;
public class TestStack{

    public static void Main( ) {

        Stack myStack = new Stack( );
        myStack.Push("Item 1");
        myStack.Push("Item 2");
        myStack.Push("Item 3");
        myStack.Push("Item 4");

        while(myStack.Count > 0) {
            Console.WriteLine(myStack.Pop( ));
        }
    }
}
```



[\[Team LiB \]](#)

[\[Team LiB \]](#)

Appendix D. Common Utilities

Microsoft .NET Framework provides many tools to help developers make the best use of the Framework. In the following sections, we document the commonly used subset of .NET tools that we've used throughout this book:

- Assembly Generation Utility (*al.exe*)
- Assembly Registration Utility (*gacutil.exe*)
- MSIL Assembler (*ilasm.exe*)
- MSIL Disassembler (*ildasm.exe*)
- C++ Compiler (*cl.exe*)
- C# Compiler (*csc.exe*)
- VB Compiler (*vbc.exe*)
- PE File viewer (*dumpbin.exe*)
- Type Library Exporter (*tlbexp.exe*)
- Type Library Importer (*tlbimp.exe*)
- XML Schema Definition Tool (*xsd.exe*)
- Shared Name Utility (*sn.exe*)
- Web Service Utility (*wsdl.exe*)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.1 Assembly Generation Utility (al.exe)

al.exe is generally used to generate assemblies with manifests. [Table D-1](#) shows some of the common uses of the Assembly Generation Utility.

Table D-1. Assembly Generation Utility common uses

Option	Description
<code>/flags:flags</code>	Specifies a value for the Flags field in the assembly. <code>0x0000</code> : side-by-side compatible <code>0x0010</code> : cannot execute with other versions in the same application domain. <code>0x0020</code> : cannot execute with other versions in the same process <code>0x0030</code> : cannot execute with other versions on the same computer
<code>/help</code> or <code>/?</code>	Use to get help for this command.
<code>/keyfile:keyfilename</code> or <code>/keyf:keyfilename</code>	Use to create shared components. <i>keyfilename</i> contains a key pair generated with the Shared Name Utility (<i>sn.exe</i>). The compiler inserts the public key into the assembly manifest and then signs the assembly with the private key.
<code>/keyname:keycontainer</code> or <code>/keyn:keycontainer</code>	Use to create shared components. <i>keycontainer</i> contains a key pair generated and installed into a key container with the Shared Name Utility (<i>sn.exe</i>). The compiler inserts the public key into the assembly manifest and then signs the assembly with the private key.
<code>/main:entrymethod</code>	Specifies the entry-point method name when converting a module to an executable.
<code>/out:filename</code>	Specifies the output filename.
<code>/target:lib exe win</code> or <code>/t:lib exe win</code>	Specifies the file format of the output file (<code>lib</code> for library, <code>exe</code> for console executable, and <code>win</code> for win32 executable). The default setting is <code>lib</code> .

Option	Description
<i>/version:major.minor.revision.build</i>	Specifies version information for the assembly. The default value is 0.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.2 Assembly Registration Utility (gacutil.exe)

You can use *gacutil.exe* to install and uninstall an assembly, as well as to list the content of the GAC. [Table D-2](#) shows some of the common uses of the Assembly Registration Utility.

Table D-2. Assembly Registration Utility common uses

Option	Description
<code>/l</code>	To list the content of the GAC.
<code>/ldl</code>	To list the content of the downloaded files cache.
<code>/cdl</code>	To clear the content of the downloaded file cache.
<code>/i filename</code>	To install an assembly with file named <i>filename</i> into the GAC.
<code>/u assemblyname</code>	To uninstall an assembly from the GAC by specifying the assembly name. If multiple versions of the same assembly exist, all of them will be removed unless a version is specified with the <i>assemblyname</i> (i.e., <code>gac -u myAssembly,ver=1.0.0.1</code>).
<code>/h or /help or /?</code>	To display command syntax and options.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.3 MSIL Assembler (ilasm.exe)

This tool takes MSIL as input and generates a portable executable (PE) file containing the MSIL and the metadata required to run on the .NET Framework. This is most useful to vendors who would like to create MSIL-compliant compilers. All they have to do is write the compiler to translate the source language to MSIL. *Ilasm.exe* will take the second step to put the MSIL content into the PE format where it can be executed on the .NET Framework. The general syntax for MSIL assembler is:

```
ilasm [options] MSILfilename
```

[Table D-3](#) shows some of the common uses of the assemblers.

Table D-3. Assemblers common uses

Option	Description
<code>/debug</code>	This option ensures that the output PE contains debugging information such as local variables, argument names, and line numbers. This is useful for debug build.
<code>/dll</code>	This option produces a <code>.dll</code> output.
<code>/exe</code>	This option produces an <code>.exe</code> output.
<code>/listing</code>	This option produces a listing of the output on STDOUT.
<code>/output=filename</code>	<code>filename</code> is the output filename.
<code>/?</code>	This option is used to obtain command-line help.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.4 MSIL Disassembler (ildasm.exe)

This tool extracts the MSIL code from a PE file targeted for .NET Framework. The general syntax for this tool is:

```
Ildasm [options] PEfilename
```

[Table D-4](#) shows some of the common uses of the disassembler.

Table D-4. Disassembler common uses

Option	Description
<code>/linenum</code>	This includes references to original source lines.
<code>/output=filename</code>	The output goes to a file instead of in a GUI dialog box.
<code>/source</code>	This shows original source lines as comments.
<code>/text</code>	The output goes in a console window.
<code>/tokens</code>	This shows metadata tokens of classes and members.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.5 C++ Compiler (cl.exe)

[Table D-5](#) shows some of the common uses of the C++ compiler.

Table D-5. C++ compiler common uses

Option	Description
<code>/CLR</code>	This option flags the compiler to compile .NET-runtime managed code.
<code>/entry:methodname</code>	For C++ managed code, this link setting should point to the main entry-point function.
<code>/link</code>	This option combines the compile and link steps.
<code>/subsystem: windows windowsce console </code>	This link option specifies the type of output.
<code>/out:filename</code>	This option allows for the output filename.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.6 C# Compiler (csc.exe)

[Table D-6](#) shows some of the common uses of the C# compiler.

Table D-6. C# compiler common uses

Option	Description
<code>/debug</code>	With this option, the compiler will emit debugging information in the output file.
<code>/define:symbol</code> or <code>/d:symbol</code>	This option is similar to C++. Use this option to define preprocessor symbols.
<code>/doc:docname</code>	<code>docname</code> is the XML output file for the autogenerated XML comment embedded in C# code.
<code>/help</code>	This option shows the command-line help for the C# compiler.
<code>/main:classname</code>	If there is more than one Main entry in different classes, you will have to specify the Main entry in which class you want the entry point of the application.
<code>/out:filename</code>	This option represents the output filename.
<code>/reference:libname</code> or <code>/r:libname</code>	This option allows single or multiple libraries be included with this compilation. For multiple libraries to be included, use a semicolon as the delimiter.
<code>/target:exe library winexe module</code> or <code>/t:exe library winexe module</code>	This option allows you to specify the type of the output: <code>exe</code> for console executables, <code>library</code> for DLLs, and <code>winexe</code> for Windows Form applications. When you set the target to <code>module</code> , the compiler outputs binary for the module but not a .NET assembly. Modules can be added to a .NET assembly later.
<code>/unsafe</code>	If you use unsafe keywords in your C# code, you will have to use this option when compiling your source.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.7 Visual Basic Compiler (vbc.exe)

[Table D-7](#) shows some of the common uses of the Visual Basic compiler.

Table D-7. Visual Basic compiler common uses

Option	Description
<code>/debug</code>	With this option, the compiler will emit debugging information in the output file.
<code>/define:symbol</code> or <code>/d:symbol</code>	Use this option to define preprocessor symbols.
<code>/help</code> or <code>/?</code>	This option shows the command-line help for the Visual Basic compiler.
<code>/keycontainer:keycontainer</code>	<code>keycontainer</code> specifies the key container that contains the key pair for signing the assembly. See <i>sn.exe</i> for information on generating the key container.
<code>/keyfile:keyfile</code>	<code>keyfile</code> specifies the key file that contains the key pair for signing the assembly. See <i>sn.exe</i> for information on generating the key file.
<code>/main:classname</code>	If there is more than one Main entry in different classes, you will have to specify the Main entry in which class you want the entry point of the application.
<code>/out:filename</code>	This option represents the output filename.
<code>optionexplicit[+/-]</code>	Turn on or off <code>optionexplicit</code> to enforce explicit or implicit declaration of variables. The default setting is on.
<code>optionstrict[+/-]</code>	Turn on or off <code>optionstrict</code> to disallow or allow casting with truncation. The default setting is on.
<code>/reference:libname</code> or <code>/r:libname</code>	This option allows single or multiple libraries be included with this compilation. For multiple libraries to be included, use a semicolon as the delimiter.
<code>/target:exe library winexe module</code> or <code>/t:exe library winexe module</code>	This option allows you to specify the type of the output: <code>exe</code> for console executables, <code>library</code> for DLLs, and <code>winexe</code> for Windows Form applications. When you set the target to module, the compiler outputs binary for the module but not a .NET

Option	Description
	assembly. Modules can be added to a .NET assembly later.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.8 PE File Format Viewer (dumpbin.exe)

dumpbin is not a new utility. However, since .NET Framework stores the IL inside the extended PE format, this old utility is still very useful for examining the structure of executable or DLLs, as well as listing import and export entries of the binaries. The general syntax for this utility is:

```
Dumpbin [options] PEfilename
```

[Table D-8](#) shows some of the common *dumpbin* uses.

Table D-8. Dumpbin common uses

Option	Description
<code>/all</code>	Displays all information from the PE file.
<code>/exports</code>	Displays all exports from the PE file.
<code>/header</code>	Displays the header information from the PE file.
<code>/clrheader</code>	Display the .NET Framework header information for an image built with /clr (see CL.exe).
<code>/imports</code>	Displays all imports for the PE file.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.9 Type Library Exporter (tlbexp.exe)

Type library exporter and importer are the two tools necessary for COM interop. The exporter generates a type library for a .NET Framework assembly so that other COM components can interop with .NET components. The general syntax for *tlbexp.exe* is:

```
tlbexp AssemblyName [options]
```

[Table D-9](#) shows some of the common uses of *tlbexp.exe*.

Table D-9. Type Library Exporter common uses

Option	Description
<code>/nologo</code>	This option suppresses the logo of the <code>tlbexp</code> executable.
<code>/out:filename</code>	<i>filename</i> is the name of the type library file.
<code>/silent</code>	This option suppresses all messages from the <code>tlbexp</code> executable.
<code>/verbose</code>	This option displays extra information while converting the component.
<code>/?</code> or <code>/help</code>	This option displays the help information for the tool.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.10 Type Library Importer (tlbimp.exe)

Because it is the reverse tool of the type library exporter, the importer generates a .NET proxy component for a COM component so that .NET components can use legacy COM components. The general syntax for *tlbimp.exe* is:

```
tlbimp PEFile [options]
```

[Table D-10](#) shows some of the common uses of *tlbimp.exe*.

Table D-10. Type Library Importer common uses

Option	Description
<code>/keycontainer: keycontainer</code>	This option signs the resulting assembly with the private key in the <i>keycontainer</i> . The public key in the <i>keyfile</i> will be used in the assembly manifest. See <i>sn.exe</i> for the <i>keycontainer</i> generation.
<code>/keyfile: keyfile</code>	This options signs the resulting assembly with the private key in the <i>keyfile</i> . The public key in the <i>keyfile</i> will be used in the assembly manifest. See <i>sn.exe</i> for <i>keyfile</i> generation.
<code>/nologo</code>	This option suppresses the logo of the <i>tlbimp</i> executable.
<code>/out:filename</code>	<i>filename</i> is the name of the type library file.
<code>/silent</code>	This option suppresses all messages from the <i>tlbimp</i> executable.
<code>/unsafe</code>	This option produces interfaces without .NET Framework security checks.
<code>/verbose</code>	This option displays extra information while converting the component.
<code>/? or /help</code>	This option displays the help information for the tool.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.11 XML Schema Definition Tool (xsd.exe)

XML Schema Definition (XSD) is useful when working with XML schemas that follow the XSD language. With XSD, you can perform the following transformations:

- XDR to XSD
- XML to XSD
- Classes to XSD
- XSD to Classes
- XSD to DataSet

D.11.1 XDR to XSD

To convert an XDR-formatted file to XSD, use the following syntax:

```
xsd [options] file.xdr
```

Note that the file extension *.xdr* dictates the conversion from XDR to XSD.

D.11.2 XML to XSD

To convert an XML-formatted file to XSD, use the following syntax:

```
xsd [options] file.xml
```

Note that the file extension *.xml* dictates the conversion from XML to XSD.

D.11.3 Classes to XSD

You can convert classes to XSD by specifying the runtime assembly file (*.exe* or *.dll* extension) as the filename to the utility. You can also specify a particular type within the assembly you want to convert to XSD using the */type* flag. The *typename* can be a wildcard match. If you omit the */type* flag, all types in the assembly will be converted. The syntax follows:

```
Xsd [/TYPE:typename] assemblyFile
```

or:

```
Xsd [/T:typename] assemblyFile
```

D.11.4 XSD to Classes

To convert XSD back to classes, use the `/classes` or `/c` flag. You can specify a particular element in the XSD schema to be converted to a class. You can also specify the language for the class source file. The general syntax follows:

```
xsd /CLASSES  
    /ELEMENT:element  
    /NAMESPACE:namespace  
    /LANGUAGE:language /URI:uri file.xsd
```

or:

```
xsd /C E:element /N:namespace /L:language /U:uri file.xsd
```

Note that `namespace`, `language`, and `uri` can be specified only once.

D.11.5 XSD to DataSet

To convert XSD to dataset, use the `/dataset` or `/d` flag. Again, you can narrow down to a particular element in the XSD schema to be converted. The general syntax follows:

```
xsd /D [/DATASET] file.xsd
```

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.12 Strong Name Utility (sn.exe)

The Strong Name Utility (*sn.exe*) guarantees unique names for shared components because these components will end up in the GAC. Each shared component is signed with a private key and published with the public key. [Table D-11](#) shows some common uses of *sn.exe*.

Table D-11. Strong Name Utility common uses

Option	Description
<i>/?</i>	This option displays more command-line help.
<i>-d keycontainer</i>	This option is used to remove the <i>keycontainer</i> from the CSP.
<i>-i keyfile keycontainer</i>	This option reads the key pair in <i>keyfile</i> and installs it in the key container <i>keycontainer</i> in the Cryptographic Service Provider (CSP).
<i>-k keyfile</i>	This option generates a new key pair and writes it to <i>keyfile</i> .
<i>-v assembly</i>	This option is used to verify the shared name in an <i>assembly</i> .

[\[Team LiB \]](#)

[\[Team LiB \]](#)

D.13 Web Service Utility (wsdl.exe)

wsdl.exe helps create ASP.NET web services and proxies for their clients. The most common usage of *wsdl.exe* is to generate proxy classes for Web services:

```
wsdl
  /language:language
  /namespace:namespace
  /out:output
  /protocol:protocol
  path
```

The *path* parameter is a local path to a service-description file or URI where the SDL file can be retrieved. The *language* parameter specifies the language for the output-proxy source file. It can be C#, VB, JS or VJS. The generated class will be in the specified namespace. The output source file is controlled by the *output* option. The *protocol* controls which protocol the proxy will use to communicate with the Web Service. The choices of protocols provided by the .NET Framework are SOAP, SOAP12, HttpGet, and HttpPost. You can also have your own protocol here if you've extended the WebClientProtocol or HttpWebClientProtocol class.

For short names options, use the following:

```
wsdl
  /l:language
  /n:namespace
  /o:output
  /protocol:protocol
  path
```

The rest of the syntax can be obtained with `wsdl /?`.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

Colophon

Our look is the result of reader comments, our own experimentation, and feedback from distribution channels. Distinctive covers complement our distinctive approach to technical topics, breathing personality and life into potentially dry subjects.

The animals on the cover of *.NET Framework Essentials*, Third Edition, are shrimp. Different species of shrimp are found in marine and fresh water—shallow and deep—all over the world. Swimming backward by rapidly flexing its abdomen and tail, and with the assistance of specialized legs for swimming, the shrimp feeds on smaller plants and animals, as well as carrion. In fact, several species engage in symbiotic (mutually beneficial or dependent) relationships with other organisms. The coral shrimp (*Stenopus hispidus*) cleans the scales of the coral fish, while the fish in turn swims backward through the shrimp's pincers (presumably to clean them).

Matt Hutchinson was the production editor for *.NET Framework Essentials*, Third Edition. Octal Publishing, Inc. provided production services. Jane Ellin, Phil Dangler, and Claire Cloutier provided quality control.

Ellie Volckhausen designed the cover of this book, based on a series design by Edie Freedman. The cover image is a 19th-century engraving from the Dover Pictorial Archive. Emma Colby produced the cover layout with QuarkXPress 4.1 using Adobe's ITC Garamond font.

David Futato designed the interior layout. This book was converted to FrameMaker 5.5.6 by Andrew Savikas. The text font is Linotype Birka; the heading font is Adobe Myriad Condensed; and the code font is LucasFont's TheSans Mono Condensed. The illustrations that appear in the book were produced by Robert Romano and Jessamyn Read using Macromedia FreeHand 9 and Adobe Photoshop 6. The tip and warning icons were drawn by Christopher Bing. This colophon was written by Jeff Holcomb.

The online edition of this book was created by the Safari production group (John Chodacki, Becki Maisch, and Madeleine Newell) using a set of Frame-to-XML conversion and cleanup tools written and maintained by Erik Ray, Benn Salter, John Chodacki, and Jeff Liggett.

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[#include directive](#)

[#using directive](#)

[%windir% setting](#)

[-k option generating/saving key information](#)

[@Assembly directive](#)

[@Control directive](#)

[@Import directive](#)

[@OutputCache directive](#)

[@Page directive 2nd](#)

[@Register directive](#)

[directive approach and](#)

[pagelets approach and](#)

[@WebService directive](#)

[__gc keywords in C++](#)

[__interface keyword in C++](#)

[_CorExeMain function in CLR header](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[ABCs \(abstract base classes\)](#)

[abstract attribute for .class IL declaration](#)

[Abstract Data Type \(ADT\)](#)

[abstract keyword, signaling the C# compiler with](#)

[abstraction, engineering concept for software](#)

[Acceleration Server 2000](#)

[AcceptChanges\(\) method](#)

[access control lists \(ACLs\)](#)

[acronyms](#)

[Activate\(\) method](#)

[activation information in COM](#)

[Activator class, calling the GetObject method](#)

[Active Server Pages](#) [See ASP]

[Active Template Library](#)

[ActiveMdiChild property](#) 2nd

[ActiveX Data Objects \(ADO\)](#) 2nd

[Add\(\) method](#)

[AddAt\(\) method](#)

[AddMessageFilter Application property](#)

[AddRange\(\) method](#)

[Administrators group](#)

[ADO \(ActiveX Data Objects\)](#) 2nd

[ADO.NET](#)

[architecture](#)

[content components for](#)

[interoperability of](#)

[managed providers for](#)

[performance of](#)

[productivity of](#)

[scalability of](#)

[ADT \(Abstract Data Type\)](#)

[Agents role](#)

[al.exe \(Assembly Generation Utility\)](#)

[using assemblies](#)

[AlternatingItemStyle, controlling visual formatting of DataGrid control](#)

[AlternatingItemTemplate](#)

[anchoring controls inside containers](#)

[ansi attribute](#)

[API \(Application Programming Interface\)](#)

[language](#)

[AppDomain class](#) 2nd

[AppendChild\(\) method, constructing XML document tree](#)

[Application Center 2000 server products](#)

[Application class](#)

[Application Programming Interface](#) [See API]

[application-level configuration, managing sessions](#)

[ApplicationActivation attribute](#)

applications

[ASP.NET development](#)

[security for](#)

architecture

[of .NET framework](#)

[of ADO.NET](#)

[Array class](#) [2nd](#)

[ArrayList class](#) [2nd](#) [3rd](#)

[asmx files](#) [2nd](#)

[ASP \(Active Server Pages\)](#) [2nd](#)

[State NT service call](#)

[ASP.NET](#)

[application development](#)

[event-driven programming](#)

[server controls for](#)

[web services for](#)

[aspx file extension](#) [2nd](#) [3rd](#)

[assemblies](#)

[blocks](#)

[creating](#)

[deployment](#)

[extern keyword](#)

[IL code](#)

[manifests](#)

[public](#)

[shared](#)

[sharing and reuse](#)

[side-by-side execution](#)

[using](#)

[versioning](#)

[Assembly class](#)

[Assembly Generation Utility \(al.exe\)](#)

[using assemblies](#)

[Assembly Registration Utility \(gac.exe\)](#) [2nd](#)

[assemblyIdentity tag \(XML\)](#)

[AssemblyKeyFile attribute \(XML\)](#)

[AssemblyVersion attribute \(XML\)](#)

[ATL \(Active Template Library\)](#) [2nd](#) [3rd](#)

attributes

[COM+, programming with](#)

[inspecting](#)

[transaction support and](#)

[authentication](#) [See also security]

[different modes of](#)

[auto attribute, in common intermediate language code](#)

[AutoResetEvent class](#)

[[Team LiB](#)]

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[B2B \(business-to-business\) 2nd](#)
[B2C \(business-to-consumer\)](#)
[Basic Authentication method](#)
[Basic over SSL Authentication method](#)
[BeginTransaction\(\) method](#)
[Binary Large Objects \(BLOBs\)](#)
[binding data format for web services](#)
[binding XML element](#)
[BitArray class 2nd](#)
[BizTalk servers 2nd](#)
[BLOBs \(Binary Large Objects\)](#)
[boxing \(reference type objects\)](#)
[browsers, generating ASP pages](#)
[btnNext click event](#)
[btnNext_onclick click event](#)
[btnPrev click event](#)
[btnPrev_onclick click event](#)
[BufferResponse \(boolean\) property](#)
[build numbers, adding new component versions](#)
[business-to-business \(B2B\) 2nd](#)
[business-to-consumer \(B2C\)](#)
[button HTML tag](#)
[Button Windows control](#)
[buttons, creating in HTML](#)
[Byte datatype](#)
[bytecode language abstraction layer](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[C# language](#) [2nd](#) [3rd](#)

[code](#)

[compiler \(csc.exe\)](#)

[DataSet, creating with](#)

[language integration of](#)

[web service providers, creating with](#)

[Windows application development and](#)

[C++ language](#) [2nd](#) [3rd](#)

[compiler \(cl.exe\)](#)

[language integration of](#)

[Window application development and](#)

[CAB Files](#)

[CacheDuration property](#)

[call instruction, using Main\(\) method](#)

[CanBePooled\(\) method](#)

[cascade setting for ForeignKeyConstraint rules](#)

[catch blocks](#)

[in C#](#)

[in C++](#)

[CCW \(COM Callable Wrapper\)](#)

[cfg file extension](#) [2nd](#)

[Char class](#) [2nd](#)

[checkboxes in HTML](#)

[ChildRelations of DataTable objects](#)

[CIL \(Common Intermediate Language\)](#) [2nd](#)

[coding](#)

[cl compiler](#)

[cl.exe \(C++ compiler\)](#)

[class attribute for Webservice directive](#)

[class IL declaration](#)

[Class keyword in VB.NET](#)

[classes](#)

[C#, defining in](#)

[CLR support for](#)

[converting XML schema into](#)

[loaders](#) [2nd](#)

[Page](#)

[XML](#)

[CLI \(Common Language Infrastructure\)](#) [2nd](#)

[Click control property](#)

[Client Certificates Authentication method](#)

[Clone\(\) method, constructing XML document tree](#)

[CLR \(Common Language Runtime\)](#) [2nd](#) [3rd](#) [4th](#)

[assemblies and manifests](#)

[environment](#)

[executables](#)

[metadata](#)

- [Portable Executable Files](#)
- [CLS \(Common Language Specification\) 2nd 3rd](#)
 - [class loaders](#)
 - [JIT compilers](#)
 - [verifiers](#)
- [CLSID \(Class identifier\)](#)
- [COBOL language](#)
- [code access](#)
- [code declaration blocks 2nd](#)
- [code identity](#)
- [code manager](#)
- [code verification](#)
- [code-rendering blocks 2nd](#)
- [Codebehind attribute for WebService directive](#)
- [COFF \(Common Object File Format\) 2nd](#)
- [COM \(Component Object Model\) 2nd 3rd 4th](#)
 - [COM+ Services and Managed C++ and windows application development and](#)
- [COM Callable Wrapper \(CCW\)](#)
- [COM Interop \(COM interoperation\)](#)
- [COM+ Explorer \(Component Services Explorer\)](#)
- [COM+ Services](#)
 - [attribute-based programming](#)
 - [role-based security](#)
- [command object](#)
- [comments \(XML\)](#)
- [Commerce Server 2000](#)
- [Common Intermediate Language \[See CIL\]](#)
- [Common Language Infrastructure \(CLI\) 2nd](#)
- [Common Language Runtime \[See CLR\]](#)
- [Common Language Specification \[See CLS\]](#)
- [Common Object File Format \(COFF\) 2nd](#)
- [Common Programming Model](#)
- [Common Type System \(CTS\) 2nd 3rd](#)
- [CommonAppDataRegistry Application property](#)
- [compilers](#)
 - [C#](#)
 - [JIT \(Just-in-Time\) compilers](#)
- [component infrastructure of .NET framework](#)
- [Component Object Model \[See COM\]](#)
- [Component Pascal language](#)
- [Component Services Explorer \(COM+ Explorer\)](#)
- [componentization](#)
- [components](#)
 - [adding new versions of](#)
 - [versus assemblies](#)
 - [distributed](#)
 - [message queuing](#)
 - [private](#)
 - [shared](#)
 - [web forms](#)

[composite controls](#)
[concepts in programming languages](#)
[configuration files \(.cfg\)](#)
[connection objects](#)
[connections \(server\), pooling with ADO.NET](#)
[connectionString method in SessionState class](#)
[Console class, using WriteLine\(\) method](#)
[consumers \(web services\)](#)
 [HTTP GET](#)
 [non-.NET](#)
[containers 2nd 3rd](#)
[content components 2nd](#)
 [DataRelation class and](#)
 [DataSet class and](#)
 [DataTable class and](#)
 [DataView class and](#)
[Context property \(control class\)](#)
[contractRef element](#)
[controls 2nd](#)
 [arranging](#)
 [classes](#)
 [control class](#)
 [custom, inserting into web page](#)
 [derivatives](#)
 [forms, adding to](#)
[Controls property \(control class\)](#)
[Cookieless property in SessionState class](#)
[cookies, managing session states](#)
[CPModel class](#)
[CreateChildControls method 2nd](#)
[CreateHtmlTextWriter method 2nd](#)
[CreateWindow Windows API](#)
[cross-language debugging in VS.NET](#)
[Cryptography child namespace](#)
[csc compiler 2nd](#)
[ctor\(\) method](#)
[CTS \(Common Type System\) 2nd 3rd](#)
[custom attributes, developing/using](#)
[custom controls](#)
 [inserting into web pages](#)
 [syntax](#)
[custom server controls](#)
 [derivatives approach](#)
 [pagelets approach](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[data](#)

[metadata](#)

[data binding](#) [2nd](#)

[expressions](#)

[data reader object](#)

[Data Type document \(DTD\)](#)

[Data types](#)

[data types](#)

[DataAdapter object](#) [2nd](#)

[DataBind method](#) [2nd](#)

[DataColumn class](#)

[DataColumn object](#)

[DataGrid control](#)

[DataList control](#)

[DataRelation class](#) [2nd](#)

[DataRelationCollection type](#)

[DataRow class](#)

[DataSet class](#) [2nd](#) [3rd](#)

[C#, creating with](#)

[XML and](#) [2nd](#)

[DataSetView class](#)

[DataSource property](#)

[DataTable class](#) [2nd](#) [3rd](#)

[DataView class](#) [2nd](#) [3rd](#)

[DateTime datatype](#)

[DB \(database\)](#)

[DbDataAdapter, hierarchy of](#)

[DCOM \(Distributed Component Object Model\)](#) [2nd](#) [3rd](#)

[web services and](#)

[Windows application development and](#)

[Deactivate\(\) method](#) [2nd](#)

[debugging support](#)

[Decimal data type](#)

[definitions element](#)

[delegates](#)

[DELETE command, updating data](#)

[Delete\(\) method](#)

[deployment of .NET framework](#) [2nd](#)

[dequeues](#)

[derivative approach](#)

[DHCP \(Dynamic Host Configuration Protocol\) network](#)

[DHTML \(Dynamic HTML\)](#) [2nd](#) [3rd](#)

[DictionaryBase class](#)

[Digest Authentication method](#)

[directives](#)

[directories, storing private components](#)

[Directory class](#)

[DISCO \(Web Service Discovery\) 2nd](#)
[.disco discovery files](#)
[discovery \(web services\)](#)
[discovery root element](#)
[discoveryRef element](#)
[DispatchMessage Windows API](#)
[DISPID \(Dispatch identifier\)](#)
[Distributed Component Object Model \[See DCOM\]](#)
[distributed components](#)
[Hello application, writing](#)
[distributed computing](#)
[Distributed interNet Applications \(DNA\) 2nd](#)
[DLL \(Dynamically Linked Library\) 2nd](#)
[DNA \(Distributed interNet Applications\) 2nd](#)
[Dns class](#)
[docking controls inside containers](#)
[Document Object Model \[See DOM\]](#)
[DoEvents Application property](#)
[DOM \(Document Object Model\) 2nd](#)
[XmlDocument, generating with](#)
[Double data type](#)
[drag-and-drop, in Web GUIs](#)
[drop-down controls \(HTML\)](#)
[DTD \(Data Type document\)](#)
[dumpbin.exe \(PE File viewer\) 2nd](#)
[dynamic assemblies](#)
[dynamic discovery](#)
[Dynamic Host Configuration Protocol \(DHCP\) network](#)
[Dynamic HTML \(DHTML\) 2nd 3rd](#)
[Dynamically Linked Library \(DLL\) 2nd](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[EditItemStyle, controlling visual formatting of DataGrid control](#)

[EditItemTemplate](#)

[Eiffel# language 2nd](#)

[Enabled control property](#)

[EnableViewState flag \(control class\)](#)

[encapsulation, as a programming concept](#)

[enclosing code-rendering blocks](#)

[enqueues](#)

[enterprise services](#)

[entrypoint directives, using Main\(\) method](#)

[environment \(CLR\)](#)

[Equals\(\) method 2nd](#)

[error handlers, mechanisms for](#)

[ErrorPage property 2nd](#)

[event handlers](#)

[web forms and](#)

[Windows controls, binding with](#)

[event-driven programming](#)

[EventArgs parameter](#)

[Exception class](#)

[exception handling](#)

[as a programming concept](#)

[Exchange Server 2000](#)

[EXE \(executable\)](#)

[building with language integration](#)

[executables \(CLR\)](#)

[ExecuteNonQuery\(\) method, executing commands](#)

[ExecuteReader\(\) method, executing commands](#)

[execution support/management](#)

[Exit Application property](#)

[ExitThread Application property](#)

[extends keyword in C#](#)

[Extensible Markup Language \[See XML\]](#)

[Extensible Stylesheet Language \(XSL\)](#)

[Extensible Stylesheet Language Transformation \(XSLT\)](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[file selectors \(HTML\)](#)

[FileStream class](#)

[Fill\(\) method 2nd](#)

[firewalls, using DCOM with](#)

[FooterStyle, controlling visual formatting of DataGrid control](#)

[FooterTemplate](#)

[ForeignKeyConstraint rules](#)

[Form class 2nd](#)

[LayoutMdi\(\) method, using](#)

[Windows Forms development](#)

forms

[authentication](#)

[class](#)

[controls, adding](#)

[HTML](#)

[Web/Windows](#)

Framework (.NET)

[design goals of](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[GAC \(Global Assembly Cache\) 2nd](#)

[shared components and](#)

[viewing](#)

[gac.exe \(Assembly Registration Utility\) 2nd](#)

[garbage collecting](#)

[allocated objects](#)

[distributed](#)

[terminating web forms](#)

[GC class](#)

[GDI \(Graphical Device Interface\)](#)

[GET methods of HTTP protocol](#)

[GetChanges\(\) method 2nd](#)

[GetHashCode\(\) method](#)

[GetMessage Windows API](#)

[GetMethods\(\) method](#)

[GetObject\(\) method](#)

[GetType\(\) method](#)

[Global Assembly Cache \[See GAC\]](#)

[globally unique identifier \(GUID\)](#)

[Graphical Device Interface \(GDI\)](#)

[Graphical User Interfaces \(GUIs\)](#)

[GUID \(globally unique identifier\)](#)

[Guid datatype](#)

[GUIs \(Graphical User Interfaces\)](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[handleButtonClick event handler](#)

[HasChanges\(\) method 2nd](#)

[HasErrors method](#)

[HashTable class 2nd 3rd 4th](#)

[Haskell language 2nd](#)

[HeaderStyle, controlling visual formatting of DataGrid control](#)

[HeaderTemplate](#)

[hideby sig attribute](#)

[Host Integration Server 2000](#)

[HTML \(Hypertext Markup Language\)](#)

[ASP pages and](#)

[components](#)

[control syntax](#)

[DataList control, laying out lists through the use of](#)

[HtmlControls mapping to tags](#)

[lifecycle of](#)

[Repeater control and](#)

[System.Web.UI.HtmlControls namespace and](#)

[HtmlAnchor HTML Control](#)

[HtmlButton HTML Control](#)

[HtmlControls class 2nd](#)

[HtmlForm HTML Control](#)

[HtmlGenericControl HTML Control](#)

[HtmlGenericControl object](#)

[HtmlImage HTML Control](#)

[HtmlInputButton HTML Control](#)

[HtmlInputCheckBox HTML Control](#)

[HtmlInputFile HTML Control](#)

[HtmlInputHidden HTML Control](#)

[HtmlInputImage HTML Control](#)

[HtmlInputRadioButton HTML Control](#)

[HtmlInputText HTML Control](#)

[HtmlInputText object](#)

[HtmlSelect HTML Control](#)

[HtmlTable HTML Control](#)

[HtmlTableCell HTML Control](#)

[HtmlTableRow HTML Control](#)

[HtmlTextArea HTML Control](#)

[HTTP \(Hypertext Transfer Protocol\)](#)

[HTTP GET](#)

[ASP and](#)

[consumers](#)

[HTTP POST](#)

[ASP and](#)

[non-.NET](#)

[HttpWebRequest class](#)

[hyperlinks \(HTML\)](#)

[Hypertext Markup Language](#) [See HTML]

[Hypertext Transfer Protocol \(HTTP\)](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[ID property \(control class\)](#)

[IDE \(Integrated Development Environment\) 2nd](#)

[identities \(unique\)](#)

[IDL \(Interface Definition Language\) 2nd](#)

[IE \(Internet Explorer\)](#)

[IID \(Interface Identifier\)](#)

[IIS Admin Tool](#)

[pagelets and](#)

[Web Forms, rendering](#)

[IJW \(It Just Works\)](#)

[IL \(Intermediate Language\) 2nd 3rd](#)

[JIT compilers and](#)

[ilasm.exe \(MSIL Assembler\) 2nd](#)

[ildasm.exe \(MSIL Disassembler\) 2nd 3rd 4th](#)

[img HTML image tag](#)

[Implements keyword](#)

[in C#](#)

[in VB.NET](#)

[import directive in C++](#)

[Import keywords in VB.NET](#)

[in attribute](#)

[In-process mode for session-state managing](#)

[indexers, CLR support for](#)

[Information Server Application Program Interface \(ISAPI\)](#)

[inheritance, as a programming concept](#)

[Inherits keyword in VB.NET](#)

[Init event](#)

[InnerHtml property](#)

[Inproc \(In-Process\)](#)

[input type=x HTML tag](#)

[INSERT command, updating data](#)

[insert function](#)

[InsertAfter\(\) method, constructing XML document tree](#)

[InsertBefore\(\) method, constructing XML document tree](#)

[instance attribute for .method declaration](#)

[instpubs.sql](#)

[Int16 data type](#)

[Int32 data type](#)

[Int64 data type](#)

[Integrated Development Environment \(IDE\) 2nd](#)

[Integrated Windows Authentication method](#)

[integration](#)

[Interface Definition Language \(IDL\) 2nd](#)

[Interface Identifier \(IID\)](#)

[Interface keyword in VB.NET](#)

[interfaces](#)

[as a programming concept](#)

[attribute for .class IL declaration](#)

[CLR support for](#)

[Intermediate Language](#) [See IL]

[Internet Explorer \(IE\)](#)

[Internet Security](#) [See also security] [2nd](#)

[interoperability](#)

[of ADO.NET](#)

[support](#)

[IPAddress class](#)

[ISAPI \(Information Server Application Programming Interface\)](#)

[IsCallerInRole\(\) method](#)

[IsInTransaction properties](#)

[IsMdiChild property](#)

[IsMdiContainer property](#)

[IsPostBack property \(Page class\)](#)

[IsValid \(UserControl class\)](#)

[It Just Works \(IJW\)](#)

[ItemStyle, controlling visual formatting of DataGrid control](#)

[ItemTemplate](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[Java Virtual Machine \(JVM\) 2nd](#)

[JavaScript, using in ASP](#)

[JIT \(Just-in-Time\) compilers, using metadata](#)

[JScript](#)

[JVM \(Java Virtual Machine\) 2nd](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[key pairs](#)

[KeyDown control property](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[Label Windows control](#)

[language abstraction layers](#)

[language attribute for WebService directive](#)

[language integration 2nd](#)

[C# and](#)

[Managed C++ and](#)

[VB.NET and](#)

[LayoutMdi\(\) method 2nd](#)

[ldstr instruction, using Main\(\) method](#)

[Lease class, calling the Renew\(\) method](#)

[lease managers](#)

[Left control property](#)

[Load event](#)

[Load method](#)

[LoadControl method 2nd](#)

[LoadPageStateFromPersistenceMedium method \(Page class\)](#)

[LoadViewState method \(control class\)](#)

[LoadXml method](#)

[localhosts](#)

[locals directive in C#](#)

[Login\(\) call](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[machine-level configuration, managing sessions](#)

[Machine.cfg configuration file](#)

[Main\(\) 2nd 3rd](#)

[C# coding and](#)

[language integration and](#)

[VB.NET coding with](#)

[MainApp class](#)

[MainMenu class](#)

[managed providers for ADO.NET 2nd](#)

[Manager role](#)

[manifests 2nd](#)

[assembly metadata and](#)

[deployment](#)

[IL code](#)

[sharing and reuse](#)

[side-by-side execution](#)

[versioning](#)

[MapPath method \(Page class\)](#)

[Marshal-By-Reference \(MBR\) 2nd](#)

[maturity factors of .NET](#)

[maxstack directives, using Main\(\) method](#)

[MBR \(Marshal-By-Reference\) 2nd](#)

[MDI \(Multiple Document Interface\)](#)

[MdiChildren property](#)

[MdiList property](#)

[MdiParent property](#)

[MemoryStream class](#)

[MenuItem class 2nd](#)

[menus in standard forms](#)

[Mercury language](#)

[Merge\(\) method](#)

[MergeOrder property 2nd](#)

[MessageQueue class](#)

[messages](#)

[dequeues](#)

[elements](#)

[enqueues](#)

[queuing](#)

[from web methods](#)

[metadata](#)

[examining](#)

[inspecting/emitting](#)

[interoperability support and](#)

[manifests](#)

[MethodInfo class](#)

[methods](#)

[control class](#)

[defining in C#](#)

[Page class](#)

[MFC \(Microsoft Foundation Classes\) 2nd 3rd](#)

[Microsoft](#)

[.NET, overview of](#)

[Passport authentication service and](#)

[Microsoft Common Object File Format \[See COFF\]](#)

[Microsoft Extensible Markup Language \(MSXML\)](#)

[Microsoft Foundation Classes \(MFC\) 2nd 3rd](#)

[Microsoft Intermediate Language \(MSIL\)](#)

[Microsoft Message Queuing \(MSMQ\)](#)

[Microsoft Transaction Server \(MTS\) 2nd](#)

[Microsoft Visual C++ Runtime \(MSVCRT\)](#)

[mode property in SessionState class](#)

[Module class](#)

[Module keywords in VB.NET](#)

[Mondrian language](#)

[mscorlib assembly 2nd](#)

[msdata\PrimaryKey attribute \(XML\)](#)

[MSI \(Microsoft Window Installer Package\)](#)

[MSIL \(Microsoft Intermediate Language\)](#)

[MSIL Assembler \(ilasm.exe\) 2nd](#)

[MSIL Disassembler \(ildasm.exe\) 2nd 3rd 4th](#)

[MSMQ \(Microsoft Message Queuing\)](#)

[MSVCRT \(Microsoft Visual C++ Runtime\)](#)

[MSXML \(Microsoft Extensible Markup Language\)](#)

[MTS \(Microsoft Transaction Server\) 2nd](#)

[multi-module assemblies](#)

[Multiline property](#)

[multiline text areas in HTML](#)

[Multiple Document Interface \(MDI\)](#)

[multiple-directory deployment](#)

[MustInherit keyword in VB.NET](#)

[MustOverride keyword in VB.NET](#)

[Mutex class](#)

[My Services \(.NET\)](#)

[MyDII.dll](#)

[MyTransactionVote properties](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[N-Tier](#)

[name attribute in assemblyIdentity tag](#)

[namespaces](#)

[IL declaration](#)

[system web services](#)

[NAT \(Network Address Translation\)](#)

[using DCOM with](#)

[NDR \(Network Data Representation\) 2nd](#)

[Network Address Translation \[See NAT\]](#)

[Network Data Representation \(NDR\) 2nd](#)

[newslet attribute for .method declaration](#)

[NextRecordset\(\)](#)

[NGWS](#)

[nodes \(XML\)](#)

[NTFS \(NT Filesystem\)](#)

[NTLM \(NT Lan Manager\)](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[Oberon language](#)

[Object Windows Library \(OWL\)](#)

[object-oriented programming \(OOP\)](#)

[ObjectPooling attribute](#)

[objects](#)

[classes](#)

[parameters](#)

[pooling](#)

[remote](#)

[OBJREF \(Object References\)](#)

[OLE DB managed provider](#)

[OleDbCommand connection property, executing comm](#)

[OleDbCommand in OLE DB managed provider](#)

[OleDbConnection in OLE DB managed provider](#)

[OleDbParameter in OLE DB managed provider](#)

[OOP \(object-oriented programming\)](#)

[OpenHandler event handler](#)

[operating system \(OS\)](#)

[operation XML element](#)

[OrderID field](#)

[OS \(operating system\)](#)

[out attribute](#)

[out-of-process session-state management 2nd](#)

[output HtmlTextWriter object](#)

[OWL \(Object Windows Library\)](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[p-code language abstraction layer](#)

[P/Invoke \(Platform Invoke\)](#)

[Page class](#)

[Web Form events and](#)

[Page_Load event handler](#)

[pagelet approach](#)

[Parent control property](#)

[ParentRelations of DataTable objects](#)

[parsers for XML](#)

[Passport authentication service](#) [2nd](#)

[Path class](#)

[PE \(Portable Executable\)](#) [2nd](#) [3rd](#)

[compiling files with JIT](#)

[PE File viewer \(dumpbin.exe\)](#) [2nd](#)

[performance of ADO.NET](#) [2nd](#)

[Perl language](#)

[Permissions child namespace](#)

[platform of .NET](#)

[plug-and-play](#)

[point type operations](#)

[Policy child namespace](#)

[polymorphism](#) [2nd](#)

[pools](#)

[port endpoint for web services](#)

[Portable Executable](#)

[portType XML element](#)

[POST method](#)

[of HTTP protocol](#)

[pragma compiler statement](#)

[PreJIT compilers](#)

[PrependChild\(\) method, constructing XML document tree](#)

[PreRender event](#)

[PrimaryKey property](#)

[Principal child namespace](#)

[private assemblies](#)

[private attribute](#)

[private components](#)

[private keys](#) [2nd](#)

[private keyword, signaling the C# compiler with](#)

[private queues](#)

[ProcessCmdKey control property](#)

[productivity of ADO.NET](#) [2nd](#)

[programming concepts](#)

[Project HailStorm](#)

[properties](#)

[CLR support for](#)

[defining in C#](#)

[providers \(web service\)](#)

[public assemblies](#)

[public keyword, signaling the C# compiler with](#)

[public methods of Object classes](#)

[publicKeyToken attribute in assemblyIdentity tag](#)

[publickeytoken value](#)

[published assemblies](#)

[Python 2nd](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[Query Analyzer, loading/running instpubs.sql](#)

[query command \(execution type\)](#)

[Queue class](#) [2nd](#) [3rd](#) [4th](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[RAD \(Rapid Application Development\)](#) [2nd](#) [3rd](#)

[ASP.NET and](#)

[radio buttons \(HTML\)](#)

[Rapid Application Development](#) [\[See RAD\]](#)

[RCW \(Runtime Callable Wrapper\)](#)

[Read\(\) method](#)

[ReadXml\(\) method](#)

[ReadXmlData\(\) method](#)

[Receive\(\) method](#)

[recordsets](#)

[reference types](#)

[referenced assemblies, identifying](#)

[Reflection namespace](#)

[Refresh control property](#)

[REGASM \(Register Assembly tool\)](#)

[RegAsm.exe](#)

[RegisterClass Window API](#)

[registered assemblies](#)

[RegisterWellKnownServiceType\(\) method](#)

[RejectChanges\(\) method](#)

[Relations of tables in databases](#) [2nd](#)

[reliability of .NET framework](#)

[remote objects](#)

[Remote Procedure Calling \(RPC\)](#)

[RemoveMessageFilter Application property](#)

[Render method \(control class\)](#)

[Renew\(\) method](#)

[Repeater control](#)

[revision numbers, adding new component versions](#)

[Rogue Wave](#)

[role-based security](#) [2nd](#) [\[See also security\]](#)

[programming](#)

[rows, in HTML tables](#)

[RPC \(Remote Procedure Calling\)](#)

[RPG third-party language](#)

[Run Application property](#)

[Run method](#)

[runat attribute](#) [2nd](#)

[Runtime Callable Wrapper \(RCW\)](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[SavePageStateToPersistenceMedium method \(Page class\)](#)

[SaveViewState method \(control class\)](#)

[SAX \(Simple API for XML\) parsing specification](#)

[SByte datatype](#)

[scalability of ADO 2nd](#)

[Schema Definition](#)

[Scheme](#)

[SCL \(SOAP Contract Language\)](#)

[scrollbars in standard forms](#)

[SDI \(Single Document Interface\)](#)

[SDK \(Software Development Kit\)](#)

[Secure Sockets Layer \(SSL\)](#)

[security 2nd](#)

[application](#)

[properties in IIS Admin Tool](#)

[system](#)

[web services and](#)

[security identifiers \(SIDs\)](#)

[SEH \(Structured Exception Handling\)](#)

[SELECT command \(data retrieval\)](#)

[select HTML tag](#)

[SelectedItemStyle, controlling visual formatting of DataGrid control](#)

[SelectedItemTemplate](#)

[Send\(\) method](#)

[SeparatorTemplate](#)

[server-side object tags](#)

[ServerChange Web Form event](#)

[ServerClick Web Form event](#)

[servers](#)

[ADO.NET, scalability of](#)

[controls 2nd](#)

[distributed Hello application, writing for](#)

[ServicedComponent class](#)

[Services](#)

[network endpoints for Web](#)

[viewing software as](#)

[session state management](#)

[SessionState class, properties of](#)

[SetComplete\(\) method](#)

[SetDefault setting for ForeignKeyConstraint rules](#)

[SetNull setting for ForeignKeyConstraint rules](#)

[shared assemblies \[See also assemblies\] 2nd 3rd 4th](#)

[making components](#)

[shared components \[See also components\] 2nd](#)

[Shared Name Utility \(sn.exe\) 2nd](#)

[Shell Cache Viewer](#)

[adding new component versions](#)

[Show control property](#)
[ShowWindow Windows APIs](#)
[SIDs \(security identifiers\)](#)
[Simple API for XML \(SAX\) parsing specification](#)
[Simple Mail Transfer Protocol \(SMTP\)](#)
[Simple Object Access Protocol \(SOAP\)](#)
[Single data type](#)
[Single Document Interface \(SDI\)](#)
[single-module assemblies](#)
[Smalltalk language 2nd](#)
[SMTP \(Simple Mail Transfer Protocol\)](#)
[sn.exe \(Shared Name Utility\)](#)
[SOAP \(Simple Object Access Protocol\) 2nd](#)
[SOAP Contract Language \(SCL\)](#)
[Software Development Kit \(SDK\)](#)
[SortedList class 2nd 3rd](#)
[sponsors](#)
[SQL managed provider](#)
[SQL SELECT command \(data retrieval\)](#)
[SQL Servers 2nd](#)
[mode for session-state managing](#)
[session-state management for](#)
[SQL UPDATE command, updating data](#)
[SqlClient objects in SQL server managed provider](#)
[SqlCommand connection property, executing commands](#)
[sqlConnectionString property in SessionState class](#)
[SSL \(Secure Sockets Layer\)](#)
[Stack class 2nd 3rd 4th](#)
[Standard ML language](#)
[Standard Template Library \(STL\)](#)
[standard text input \(HTML\)](#)
[StartupPath Application property](#)
[state management](#)
[static assemblies](#)
[Static Discovery](#)
[STL \(Standard Template Library\)](#)
[streams \(XML files\)](#)
[String class 2nd](#)
[Structured Exception Handling \(SEH\)](#)
[styles, controlling visual formatting of DataGrid control](#)
[syntax of Web Forms](#)
[System namespaces](#)
[system security 2nd \[See also security\]](#)
[System.Attribute class, creating custom attributes](#)
[System.Collections namespace](#)
[System.Data namespace 2nd](#)
[System.Data.Common namespace for managed providers](#)
[System.Data.IDbConnection namespace, implementing](#)
[System.Data.OleDb namespace](#)
[System.IO namespace](#)
[System.Messaging namespace](#)
[System.Net namespace](#)

[System.Object 2nd 3rd 4th](#)
[System.Object root class](#)
[System.Reflection namespace](#)
[System.Runtime namespace](#)
[System.Runtime.Remoting namespace](#)
[System.Runtime.Remoting.Channels.Tcp namespace 2nd](#)
[System.Security namespace](#)
[System.Threading namespace](#)
[System.Web.Services namespace 2nd](#)
[System.Web.UI namespace 2nd](#)
 [Page class](#)
 [System.Web.UI.HtmlControls namespace and](#)
 [System.Web.UI.WebControls namespace and](#)
 [System.Windows.Forms namespace and](#)
[System.Web.UI.HtmlControls namespace 2nd](#)
[System.Web.UI.WebControls namespace 2nd](#)
[System.Windows.Forms namespace 2nd](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[table HTML tag](#)

[tables 2nd](#)

[HTML](#)

[tags, mapping to HtmlControls](#)

[TCP \(Transport Control Protocol\)](#)

[td HTML tag](#)

[templates](#)

[controlling display of data items](#)

[DataList control and](#)

[Repeater control and](#)

[text input \(HTML\)](#)

[textarea HTML tag](#)

[TextBox Windows control](#)

[Thread class](#)

[ThreadPool class](#)

[Timeout property in SessionState class](#)

[titlebars in standard forms](#)

[tlb files](#)

[tlbexp.exe](#)

[tlbimp.exe](#)

[TMT Pascal language](#)

[tools, using metadata](#)

[ToString\(\) method](#)

[tr HTML tag](#)

[Trace property \(Page class\)](#)

[TraceContext object](#)

[Transaction objects \(ADO.NET\)](#)

[TransactionOption property](#)

[TranslateMessage Windows API](#)

[Transport Control Protocol \(TCP\)](#)

[try blocks](#)

[in C# 2nd](#)

[in C++](#)

[type container in WSDL structure](#)

[type libraries](#)

[.tlb files 2nd](#)

[Type Library Exporter \(tlbexp.exe\) 2nd 3rd](#)

[Type Library Importer \(tlbimp.exe\) 2nd](#)

[types XML element](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[UDDI \(Universal Description, Discovery, and Integration\) 2nd](#)

[UDF \(Uniform Data Format\)](#)

[UI \(User Interface\) 2nd](#)

[UInt16 data type](#)

[UInt32 data type](#)

[UInt64 data type](#)

[umpbin.exe utility](#)

[unboxing \(reference type objects\)](#)

[Uniform Data Format \(UDF\)](#)

[Uniform Resource Identifier \(URI\)](#)

[Uniform Resource Locator \(URL\)](#)

[unique identities](#)

[UniqueConstraint](#)

[Universal Description, Discovery, and Integration \(UDDI\) 2nd](#)

[Universal Runtime \(URT\)](#)

[Unload event](#)

[update function](#)

[Update\(\) method, modifying databases](#)

[URI \(Uniform Resource Identifier\)](#)

[URL \(Uniform Resource Locator\)](#)

[URT \(Universal Runtime\)](#)

[User \(UserControl class\)](#)

[User Interface \(UI\) 2nd](#)

[UserAppDataRegistry Application property](#)

[Users group, in role-based security](#)

[using keyword in C#](#)

[using namespace statement in C++](#)

[utilities](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[Validate \(UserControl class\)](#)

[Validate method \(Page class\)](#)

[Validators \(UserControl class\)](#)

[Validators property \(Page class\)](#)

[value keyword](#)

[value types](#)

[VB \(Visual Basic\)](#)

[compiler \(vbc.exe\)](#)

[VB.NET \(Visual Basic .NET\)](#)

[code](#)

[language integration of](#)

[vbc compiler 2nd](#)

[vbc.exe \(VB compiler\)](#)

[VBRUN \(Visual Basic Runtime\)](#)

[VBScript, using in ASP](#)

[VEE \(Virtual Execution Engine\)](#)

[verifiers](#)

[versions of components, adding new](#)

[VES \(Virtual Execution System\)](#)

[ViewState property \(control class\)](#)

[Virtual Execution Engine \(VEE\)](#)

[Virtual Execution System \(VES\)](#)

[Virtual Object System \(VOS\)](#)

[Visible property \(control class\)](#)

[Visual Basic \[See VB\]](#)

[Visual Basic .NET \(VB.NET\)](#)

[Visual Basic Runtime \(VBRUN\)](#)

[visual elements in standard forms](#)

[Visual inheritance](#)

[Visual Studio .NET \[See VS.NET\]](#)

[Void datatype](#)

[VOS \(Virtual Object System\)](#)

[vptr](#)

[VS.NET \(Visual Studio.NET\) 2nd](#)

[ADO.NET and](#)

[inheriting classes with](#)

[web services, using](#)

[vtbl](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[WAP \(Wireless Access Protocol\)](#)

[Web](#)

[GUIs](#)

[methods](#)

[Microsoft Services](#)

[paradigm shifts](#)

[services](#)

[Web Forms 2nd](#)

[application development](#)

[ASP.NET and](#)

[components](#)

[configuring](#)

[custom server controls](#)

[events](#)

[life cycle of](#)

[syntax](#)

[System.Web.UI namespace and](#)

[Web Service Discovery \(DISCO\) 2nd](#)

[Web Service Utility \(wsdl.exe\)](#)

[web services 2nd](#)

[consumers](#)

[discovery](#)

[for ASP.NET](#)

[framework](#)

[namespaces](#)

[providers](#)

[security and](#)

[using](#)

[Web Forms and](#)

[wire formats](#)

[WSDL and](#)

[Web Services Description Language 2nd 3rd 4th 5th](#)

[web.config files 2nd](#)

[managing sessions](#)

[setting](#)

[WebControls object](#)

[WebMethod attribute](#)

[WebMethodAttribute class in System.Web.Services namespace](#)

[WebService class in System.Web.Services namespace](#)

[WebServiceAttribute class in System.Web Services namespace](#)

[widgets](#)

[Win32](#)

[Extensions \(Python\)](#)

[Windows 2000](#)

[class loaders and](#)

[DLL Hell, avoiding](#)

[PE file formats and](#)

[role-based security](#)
[security features for](#)
Windows 98
[class loaders and](#)
[PE file formats and](#)
[Windows authentication](#) [See also security] [2nd](#)
[Windows CE](#)
[Windows Forms](#)
[application](#)
[control class](#)
[controls](#)
[developing](#)
[form class](#)
[stage deployment](#)
[System.Windows.Forms namespace and](#)
[web services and](#)
[Windows Me](#)
[class loaders and](#)
[PE file formats and](#)
[Windows NT security features](#)
[Windows XP](#)
[WinForms](#)
[wire formats \(web services\)](#)
[Wireless Access Protocol \(WAP\)](#)
[Wireless Markup Language \(WML\)](#)
[WM_CLOSE](#)
[WM_CREATE](#)
[WM_SIZE](#)
[WML \(Wireless Markup Language\)](#)
[WriteLine\(\) method](#) [2nd](#)
[WriteXml\(\) method](#)
[WriteXmlSchema\(\) method](#)
[WSDL \(Web Services Description Language\)](#) [2nd](#) [3rd](#) [4th](#) [5th](#) [6th](#)
[wsdl.exe \(Web Service Utility\)](#) [2nd](#) [3rd](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[xcopy](#)

[XML \(Extensible Markup Language\)](#)

[C# , generating in](#)

[classes](#)

[classes/DataSet to](#)

[DataSets and 2nd](#)

[parsers](#)

[Schema Definition \(XSD\) 2nd 3rd 4th 5th](#)

[Schema Editor in VS.NET](#)

[XML Schema Definition \[See XSD\]](#)

[XmlAttribute node type](#)

[XmlDocument base class](#)

[XmlDocument node type](#)

[XmlElement node type](#)

[XmlNamedNodeMap collection object](#)

[XmlNode base class](#)

[XmlNodeList collection object](#)

[XmlNodeReader object](#)

[XmlReader object](#)

[XmlText node type](#)

[XmlTextReader object](#)

[XmlWriter object](#)

[XPath \(XML Path\)](#)

[XSD \(XML Schema Definition\) 2nd 3rd 4th 5th](#)

[classes/DataSet to](#)

[XSL \(Extensible Stylesheet Language\)](#)

[XSLT \(Extensible Stylesheet Language Transformation\)](#)

[XsltTransform class](#)

[\[Team LiB \]](#)

[\[Team LiB \]](#)

[\[SYMBOL\]](#) [\[A\]](#) [\[B\]](#) [\[C\]](#) [\[D\]](#) [\[E\]](#) [\[F\]](#) [\[G\]](#) [\[H\]](#) [\[I\]](#) [\[J\]](#) [\[K\]](#) [\[L\]](#) [\[M\]](#) [\[N\]](#) [\[O\]](#) [\[P\]](#) [\[Q\]](#) [\[R\]](#) [\[S\]](#) [\[T\]](#) [\[U\]](#) [\[V\]](#) [\[W\]](#) [\[X\]](#) [\[Z\]](#)

[zApp framework](#)

[\[Team LiB \]](#)