

THE EXPERT'S VOICE® IN .NET

SECOND EDITION

Beginning C# Object-Oriented Programming

*LEARN TO DESIGN AND DEVELOP
MODERN APPLICATIONS IN C#*

Dan Clark

Apress®

www.it-ebooks.info

For your convenience Apress has placed some of the front matter material after the index. Please use the Bookmarks and Contents at a Glance links to access them.



Apress®

Contents at a Glance

About the Author	xvii
About the Technical Reviewer	xix
Acknowledgments	xxi
Introduction	xxiii
■ Chapter 1: Overview of Object-Oriented Programming	1
■ Chapter 2: Designing OOP Solutions: Identifying the Class Structure	7
■ Chapter 3: Designing OOP Solutions: Modeling the Object Interaction	25
■ Chapter 4: Designing OOP Solutions: A Case Study	43
■ Chapter 5: Introducing the .NET Framework and Visual Studio	59
■ Chapter 6: Creating Classes	83
■ Chapter 7: Creating Class Hierarchies	97
■ Chapter 8: Implementing Object Collaboration.....	119
■ Chapter 9: Working with Collections	143
■ Chapter 10: Implementing the Data Access Layer.....	161
■ Chapter 11: Developing WPF Applications	193
■ Chapter 12: Developing Web Applications.....	221
■ Chapter 13: Developing Windows Store Applications	251
■ Chapter 14: Developing and Consuming Web Services	273
■ Chapter 15: Developing the Office Supply Ordering Application	295
■ Chapter 16: Wrapping Up.....	321

■ Appendix A: Fundamental Programming Concepts	325
■ Appendix B: Exception Handling in C#	341
■ Appendix C: Installing the Required Software	347
Index	353

Introduction

It has been my experience as a .NET trainer and lead programmer that most people do not have trouble picking up the syntax of the C# language. What perplexes and frustrates many people are the higher-level concepts of object-oriented programming methodology and design. To compound the problem, most introductory programming books and training classes skim over these concepts or, worse, don't cover them at all. It is my hope that this book fills this void. My goals in writing this book are twofold. My first goal is to provide you with the information you need to understand the fundamentals of programming in C#. More importantly, my second goal is to present you with the information required to master the higher-level concepts of object-oriented programming methodology and design.

This book provides the knowledge you need to architect an object-oriented programming solution aimed at solving a business problem. As you work your way through the book, you will learn first how to analyze the business requirements of an application. Next, you will model the objects and relationships involved in the solution design. Finally, you will implement the solution using C#. Along the way, you will learn about the fundamentals of software design, the Unified Modeling Language (UML), object-oriented programming, C#, and the .NET Framework.

Because this is an introductory book, it's meant to be a starting point for your study of the topics it presents. As such, this book is not designed to make you an expert in object-oriented programming and UML; nor is it an exhaustive discussion of C# and the .NET Framework; nor is it an in-depth study of Visual Studio. It takes considerable time and effort to become proficient in any one of these areas. It is my hope that by reading this book, your first experiences in object-oriented programming will be enjoyable and comprehensible—and that these experiences will instill a desire for further study.

Target Audience

The target audience for this book is the beginning C# programmer who wants to gain a foundation in object-oriented programming along with C# language basics. Programmers transitioning from a procedure-oriented programming model to an object-oriented model will also benefit from this book. In addition, there are many Visual Basic (VB) programmers who want to transition to C#. Before transitioning to C#, it is imperative that you understand the fundamentals of object-oriented programming.

Because the experience level of a “beginner” can vary immensely, I have included a primer in Appendix A that discusses some basic programming concepts and how they are implemented in C#. I would suggest you review these concepts if you are new to programming.

Activities and Software Requirements

One of the most important aspects of learning is doing. You can't learn to ride a bike without jumping on a bike, and you can't learn to program without cranking out code. Any successful training program needs to include both a theory component and a hands-on component.

I have included both components throughout this book. It is my hope that you will take seriously the activities I have added to each chapter and work through them thoroughly—even repeatedly. Contrary to some students' perception that these activities are “exercises in typing,” the activities are where you get a chance to make the theory

concrete and where true simulation of the concepts occurs. I also encourage you to play as you work through an activity. Don't be afraid to alter some of the code just to see what happens. Some of the best learning experiences occur when students "color outside the lines."

The UML modeling activities in early chapters are designed for someone using UMLet. I chose this program because it's a good diagramming tool to learn with. It lets you create UML diagrams without adding a lot of advanced features associated with the high-end CASE tools. UMLet is a free open-source tool and can be downloaded from www.umlet.com. You can also use another tool such as Visio to complete the activities. However, you don't need a tool to complete these activities; paper and pencil will work just fine.

Once you begin coding, you will need Visual Studio 2012 with C# installed. I encourage you to install the help files and make ample use of them while completing the activities. Chapter 13 deals with creating Windows Store Apps and requires Visual Studio installed on the Windows 8 operating system. Later chapters require Microsoft SQL Server 2008 or higher with the Pubs and Northwind databases installed. Appendix C includes instructions on downloading and installing the sample databases. You can find free Express and trial editions of both Visual Studio and SQL Server along with trial editions of Windows 8 at www.msdn.microsoft.com.



Overview of Object-Oriented Programming

To set the stage for your study of object-oriented programming (OOP) and C#, this chapter will look briefly at the history of object-oriented programming and the characteristics of an object-oriented programming language. You will look at why object-oriented programming has become so important in the development of industrial-strength distributed-software systems. You will also examine how C# has evolved into one of the leading application programming languages.

After reading this chapter, you will be familiar with the following:

- what object-oriented programming is
- why object-oriented programming has become so important in the development of industrial-strength applications
- the characteristics that make a programming language object-oriented
- the history and evolution of C#

What is OOP?

Object-oriented programming is an approach to software development in which the structure of the software is based on objects interacting with each other to accomplish a task. This interaction takes the form of messages passing back and forth between the objects. In response to a message, an object can perform an action.

If you look at how you accomplish tasks in the world around you, you can see that you interact in an object-oriented world. If you want to go to the store, for example, you interact with a car object. A car object consists of other objects that interact with each other to accomplish the task of getting you to the store. You put the key object in the ignition object and turn it. This in turn sends a message (through an electrical signal) to the starter object, which interacts with the engine object to start the car. As a driver, you are isolated from the logic of how the objects of the system work together to start the car. You just initiate the sequence of events by executing the start method of the ignition object with the key. You then wait for a response (message) of success or failure.

Similarly, users of software programs are isolated from the logic needed to accomplish a task. For example, when you print a page in your word processor, you initiate the action by clicking a print button. You are isolated from the internal processing that needs to occur; you just wait for a response telling you if it printed. In the software program, the button object interacts with a printer object, which interacts with the actual printer to accomplish the task of printing the page.

The History of OOP

OOP concepts started surfacing in the mid-1960s with a programming language called Simula and further evolved in the 1970s with advent of Smalltalk. Although software developers did not overwhelmingly embrace these early advances in OOP languages, object-oriented methodologies continued to evolve. In the mid-1980s there was a resurgence of interest in object-oriented methodologies. Specifically, OOP languages such as C++ and Eiffel became popular with mainstream computer programmers. OOP continued to grow in popularity in the 1990s, most notably with the advent of Java and the huge following it attracted. And in 2002, in conjunction with the release of the .NET Framework, Microsoft introduced a new OOP language, C# (pronounced C-sharp) and revamped their widely popular existing language, Visual Basic, so that it is now truly object-oriented. Today OOP languages continue to flourish and are a mainstay of modern programming.

Why Use OOP?

Why has OOP developed into such a widely used paradigm for solving business problems today? During the 1970s and 1980s, procedure-oriented programming languages such as C, Pascal, and Fortran were widely used to develop business-oriented software systems. Procedural languages organize the program in a linear fashion—they run from top to bottom. In other words, the program is a series of steps that run one after another. This type of programming worked fine for small programs that consisted of a few hundred code lines, but as programs became larger they became hard to manage and debug.

In an attempt to manage the ever-increasing size of the programs, structured programming was introduced to break down the code into manageable segments called functions or procedures. This was an improvement, but as programs performed more complex business functionality and interacted with other systems, the following shortcomings of structural programming began to surface:

- Programs became harder to maintain.
- Existing functionality was hard to alter without adversely affecting all of the system's functionality.
- New programs were essentially built from scratch. Consequently, there was little return on the investment of previous efforts.
- Programming was not conducive to team development. Programmers had to know every aspect of how a program worked and could not isolate their efforts on one aspect of a system.
- It was hard to translate business models into programming models.
- Structural programming worked well in isolation but did not integrate well with other systems.

In addition to these shortcomings, some evolutions of computing systems caused further strain on the structural program approach, such as:

- Nonprogrammers demanded and got direct access to programs through the incorporation of graphical user interfaces and their desktop computers.
- Users demanded a more intuitive, less structured approach to interacting with programs.
- Computer systems evolved into a distributed model where the business logic, user interface, and backend database were loosely coupled and accessed over the Internet and intranets.

As a result, many business software developers turned to object-oriented methods and programming languages. The benefits included the following:

- a more intuitive transition from business-analysis models to software-implementation models
- the ability to maintain and implement changes in the programs more efficiently and rapidly
- the ability to create software systems more effectively using a team process, allowing specialists to work on parts of the system
- the ability to reuse code components in other programs and purchase components written by third-party developers to increase the functionality of existing programs with little effort
- better integration with loosely coupled distributed-computing systems
- improved integration with modern operating systems
- the ability to create a more intuitive graphical-user interface for the users

The Characteristics of OOP

In this section you are going to examine some fundamental concepts and terms common to all OOP languages. Don't worry about how these concepts get implemented in any particular programming language; that will come later. My goal is to familiarize you with the concepts and relate them to your everyday experiences so that they make more sense later when you look at OOP design and implementation.

Objects

As I noted earlier, we live in an object-oriented world. You are an object. You interact with other objects. In fact, you are an object with data such as your height and hair color. You also have methods that you perform or that are performed on you, such as eating and walking.

So what are objects? In OOP terms, an object is a structure for incorporating data and the procedures for working with that data. For example, if you were interested in tracking data associated with product inventory, you would create a product object that is responsible for maintaining and using the data pertaining to the products. If you wanted to have printing capabilities in your application, you would work with a printer object that is responsible for the data and methods used to interact with your printers.

Abstraction

When you interact with objects in the world, you are often only concerned with a subset of their properties. Without this ability to abstract or filter out the extraneous properties of objects, you would find it hard to process the plethora of information bombarding you and concentrate on the task at hand.

As a result of abstraction, when two different people interact with the same object, they often deal with a different subset of attributes. When I drive my car, for example, I need to know the speed of the car and the direction it is going. Because the car is using an automatic transmission, I do not need to know the revolutions per minute (RPMs) of the engine, so I filter this information out. On the other hand, this information would be critical to a racecar driver, who would not filter it out.

When constructing objects in OOP applications, it is important to incorporate this concept of abstraction. The objects include only the information that is relevant in the context of the application. If you were building a shipping application, you would construct a product object with attributes such as size and weight. The color of the item would be extraneous information and would be ignored. On the other hand, when constructing an order-entry application, the color could be important and would be included as an attribute of the product object.

Encapsulation

Another important feature of OOP is encapsulation. Encapsulation is the process in which no direct access is granted to the data; instead, it is hidden. If you want to gain access to the data, you have to interact with the object responsible for the data. In the previous inventory example, if you wanted to view or update information on the products, you would have to work through the product object. To read the data, you would send the product object a message. The product object would then read the value and send back a message telling you what the value is. The product object defines which operations can be performed on the product data. If you send a message to modify the data and the product object determines it is a valid request, it will perform the operation for you and send a message back with the result.

You experience encapsulation in your daily life all the time. Think about a human resources department. They encapsulate (hide) the information about employees. They determine how this data can be used and manipulated. Any request for the employee data or request to update the data has to be routed through them. Another example is network security. Any request for security information or a change to a security policy must be made through a network security administrator. The security data is encapsulated from the users of the network.

By encapsulating data, you make the data of your system more secure and reliable. You know how the data is being accessed and what operations are being performed on the data. This makes program maintenance much easier and also greatly simplifies the debugging process. You can also modify the methods used to work on the data, and, if you do not alter how the method is requested and the type of response sent back, you do not have to alter the other objects using the method. Think about when you send a letter in the mail. You make a request to the post office to deliver the letter. How the post office accomplishes this is not exposed to you. If it changes the route it uses to mail the letter, it does not affect how you initiate the sending of the letter. You do not have to know the post office's internal procedures used to deliver the letter.

Polymorphism

Polymorphism is the ability of two different objects to respond to the same request message in their own unique way. For example, I could train my dog to respond to the command *bark* and my bird to respond to the command *chirp*. On the other hand, I could train them to both respond to the command *speak*. Through polymorphism I know that the dog will respond with a bark and the bird will respond with a chirp.

How does this relate to OOP? You can create objects that respond to the same message in their own unique implementations. For example, you could send a print message to a printer object that would print the text on a printer, and you could send the same message to a screen object that would print the text to a window on your computer screen.

Another good example of polymorphism is the use of words in the English language. Words have many different meanings, but through the context of the sentence you can deduce which meaning is intended. You know that someone who says "Give me a break!" is not asking you to break his leg!

In OOP you implement this type of polymorphism through a process called overloading. You can implement different methods of an object that have the same name. The object can then tell which method to implement depending on the context (in other words, the number and type of arguments passed) of the message. For example, you could create two methods of an inventory object to look up the price of a product. Both these methods would be named `getPrice`. Another object could call this method and pass either the name of the product or the product ID. The inventory object could tell which `getPrice` method to run by whether a string value or an integer value was passed with the request.

Inheritance

Most real-life objects can be classified into hierarchies. For example, you can classify all dogs together as having certain common characteristics such as having four legs and fur. Their breeds further classify them into subgroups with common attributes such as size and demeanor. You also classify objects according to their function. For example, there are commercial vehicles and recreational vehicles. There are trucks and passenger cars. You classify cars according to their make and model. To make sense of the world, you need to use object hierarchies and classifications.

You use inheritance in OOP to classify the objects in your programs according to common characteristics and function. This makes working with the objects easier and more intuitive. It also makes programming easier because it enables you to combine general characteristics into a parent object and inherit these characteristics in the child objects. For example, you can define an employee object that defines all the general characteristics of employees in your company. You can then define a manager object that inherits the characteristics of the employee object but also adds characteristics unique to managers in your company. Because of inheritance the manager object will automatically reflect any changes to the characteristics of the employee object.

Aggregation

Aggregation is when an object consists of a composite of other objects that work together. For example, your lawn mower object is a composite of the wheel objects, the engine object, the blade object, and so on. In fact, the engine object is a composite of many other objects. There are many examples of aggregation in the world around us. The ability to use aggregation in OOP is a powerful feature that enables you to accurately model and implement business processes in your programs.

The History of C#

In the 1980s, most applications written to run on the Windows operating system were written in C++. Even though C++ is an OOP language, it's arguably a difficult language to master and the programmer is responsible for dealing with housekeeping tasks such as memory management and security. These housekeeping tasks are difficult to implement and often neglected, which results in bug-filled applications that are difficult to test and maintain.

In the 1990s, the Java programming language became popular. Because it's a managed-programming language, it runs on top of a unified set of class libraries that take care of the low-level programming tasks such as type safety checking, memory management, and destruction of unneeded objects. This allows the programmer to concentrate on the business logic and frees her from having to worry about the error-prone housekeeping code. As a result, programs are more compact, reliable, and easier to debug.

Seeing the success of Java and the increased popularity of the Internet, Microsoft developed its own set of managed-programming languages. Microsoft wanted to make it easier to develop both Windows- and Web-based applications. These managed languages rely on the .NET Framework to provide much of the functionality to perform the housekeeping code required in all applications. During the development of the .NET Framework, the class libraries were written in a new language called C#. The principal designer and lead architect of C# was Anders Hejlsberg. Hejlsberg was previously involved with the design of Turbo Pascal and Delphi. He leveraged this previous experience to design an OOP language that built on the successes of these languages and improved upon their shortcomings. Hejlsberg also incorporated syntax similar to C into the language in order to appeal to the C++ and Java developers. One of the goals of creating the .NET Framework and the C# language was to introduce modern concepts such as object orientation, type safety, garbage collection, and structured-exception handling directly into the platform.

Since releasing C#, Microsoft has continually sought to add additional features and enhancements to the language. For example, version 2.0 added support for *generics* (generics are covered in Chapter 9) and version 3.0 *LINQ* (more about this in Chapter 10) was added to reduce the impedance mismatch between the programming language and the database language used to retrieve and work with data. Today C# 5.0 includes support to make *parallel* and *asynchronous programming* easier for developers to implement (see Chapter 8). With Microsoft's commitment to the continual improvement and evolution of the C#, it will continue to rank as one of the most widely used programming languages in the world.

Microsoft is also committed to providing .NET developers with the tools necessary to have a highly productive and intuitive programming experience. Although you can create C# programs using a text editor, most professional programmers find an integrated development environment (IDE) invaluable in terms of ease of use and increased productivity. Microsoft has developed an exceptional IDE in Visual Studio (VS). Integrated into VS are many features that make programming for the .NET Framework more intuitive (These features are covered in Chapter 5). With the most recent release of Visual Studio 2012, Microsoft has continued to enhance the design-time developing experience. VS 2012 includes new features such as better debugging support for parallel programming and an improved code testing experience. As you work your way through this book, I think you will come to appreciate the power and productivity that Visual Studio and the C# language provide.

Summary

In this chapter, you were introduced to OOP and got a brief history of C#. Now that you have an understanding of what constitutes an OOP language and why OOP languages are so important to enterprise-level application development, your next step is to become familiar with how OOP applications are designed.

In order to meet the needs of the users, successful applications must be carefully planned and developed. The next chapter is the first in a series of three aimed at introducing you to some of the techniques used when designing object-oriented applications. You will look at the process of deciding which objects need to be included in an application and which attributes of these objects are important to the functionality of that application.



Designing OOP Solutions: Identifying the Class Structure

Most software projects you will become involved with as a business software developer will be a team effort. As a programmer on the team, you will be asked to transform the design documents into the actual application code. Additionally, because the design of object-oriented programs is an iterative process, designers depend on the feedback of the software developers to refine and modify the program design. As you gain experience in developing object-oriented software systems, you may even be asked to sit in on the design sessions and contribute to the design process. Therefore, as a software developer, you should be familiar with the purpose and the structure of the various design documents, as well as have some knowledge of how these documents are developed.

This chapter introduces you to some of the common documents used to design the static aspects of the system. (You'll learn how the dynamic aspects of the system are modeled in the next chapter.) To help you understand these documents, this chapter includes some hands-on activities based on a limited case study. You'll find similar activities corresponding to the topics of discussion in most of the chapters in this book.

After reading this chapter, you will be familiar with the following:

- The goals of software design
- The fundamentals of the Unified Modeling Language
- The purpose of a software requirement specification
- How use case diagrams model the services the system will provide
- How class diagrams model the classes of objects that need to be developed

Goals of Software Design

A well-organized approach to system design is essential when developing modern enterprise-level object-oriented programs. The design phase is one of the most important in the software development cycle. You can trace many of the problems associated with failed software projects to poor upfront design and inadequate communication between the system's developers and the system's consumers. Unfortunately, many programmers and program managers do not like getting involved in the design aspects of the system. They view any time not spent cranking out code as unproductive.

To make matters worse, with the advent of “Internet time,” consumers expect increasingly shorter development cycles. So, to meet unrealistic timelines and project scope, developers tend to forgo or cut short the system design phase of development. This is counterproductive to the system’s success. Investing time in the design process will achieve the following:

- Provide an opportunity to review the current business process and fix any inefficiencies or flaws uncovered
- Educate the customers as to how the software development process occurs and incorporate them as partners in this process
- Create realistic project scopes and timelines for completion
- Provide a basis for determining the software testing requirements
- Reduce the cost and time required to implement the software solution

A good analogy to software design is the process of building a home. You would not expect a builder to start working on the house without detailed plans (blueprints) supplied by an architect. You would expect the architect to discuss the home’s design with you before creating any blueprints. It is the architect’s job to listen to your requests and convert them into plans that a builder can use to build the home. A good architect will also educate you as to which features are reasonable for your budget and projected timeline.

Understanding the Unified Modeling Language

To successfully design object-oriented software, you need to follow a proven design methodology. One proven design methodology used in OOP today is the Unified Modeling Language (UML). UML was developed in the early 1980s as a response to the need for a standard, systematic way of modeling the design of object-oriented software. This industry standard was created and is managed by the Object Management Group (OMG). UML has evolved and matured along with the software industry and the current version 2.4.1 was formally released in 2011.

There are many advantages to using UML during the design process. When properly implemented, UML allows you to visualize the software system at various levels of detail. You can verify requirements and scope of the project with the users. It also can be used as the basis for testing the system. UML lends itself well to an incremental development process. UML modeling is also very beneficial for parallel development of large systems. Using the model, each team is aware of how their pieces fit into the system and can convey changes that may affect other teams.

UML consists of a series of textual and graphical models of the proposed solution. These models define the system scope, components of the system, user interaction with the system, and how the system components interact with each other to implement the system functionality.

Some common models used in UML are the following:

- *Software Requirement Specification (SRS)*: a textual description of the overall responsibilities and scope of the system.
- *Use Case*: a textual/graphical description of how the system will behave from the user’s perspective. Users can be human or other systems.
- *Class Diagram*: a visual blueprint of the objects that will be used to construct the system.
- *Sequence Diagram*: a model of the sequence of object interaction as the program executes. Emphasis is placed on the order of the interactions and how they proceed over time.
- *Collaboration Diagram*: a view of how objects are organized to work together as the program executes. Emphasis is placed on the communications that occur between the objects.
- *Activity Diagram*: a visual representation of the flow of execution of a process or operation.

In this chapter, you'll look at the development of the SRS, use cases, and class diagrams. Chapter 3 covers the sequence, collaboration, and activity diagrams.

Developing a SRS

The purpose of the SRS is to do the following:

- Define the functional requirements of the system.
- Identify the boundaries of the system.
- Identify the users of the system.
- Describe the interactions between the system and the external users.
- Establish a common language between the client and the program team for describing the system.
- Provide the basis for modeling use cases.

To produce the SRS, you interview the business owners and the end users of the system. The goals of these interviews are to document clearly the business processes involved and establish the system's scope. The outcome of this process is a formal document (the SRS) detailing the functional requirements of the system. A formal document helps to ensure agreement between the customers and the software developers. The SRS also provides a basis for resolving any disagreements over perceived system scope as development proceeds.

As an example, suppose that the owners of a small commuter airline want customers to be able to view flight information and reserve tickets for flights using a Web registration system. After interviewing the business managers and the ticketing agents, the software designers draft an SRS document that lists the system's functional requirements. The following are some of these requirements:

- Nonregistered web users can browse to the web site to view flight information, but they can't book flights.
- New customers who want to book flights must complete a registration form providing their name, address, company name, phone number, fax number, and e-mail address.
- A customer is classified as either a corporate customer or a retail customer.
- Customers can search for flights based on destination and departure times.
- Customers can book flights indicating the flight number and the number of seats requested.
- The system sends customers a confirmation via e-mail when the flight is booked.
- Corporate customers receive frequent flier miles when their employees book flights.
- Frequent-flier miles are used to discount future purchases.
- Ticket reservations may be canceled up to one week in advance for an 80 percent refund.
- Ticketing agents can view and update flight information.

In this partial SRS document, you can see that several succinct statements define the system scope. They describe the functionality of the system as viewed by the system's users and identify the external entities that will use it. It is important to note that the SRS does not contain references to the technical requirements of the system.

Once the SRS is developed, the functional requirements it contains are transformed into a series of use case diagrams.

Introducing Use Cases

Use cases describe how external entities will use a system. These external entities can be human or other systems (called *actors* in UML terminology). The description emphasizes the users' view of the system and the interaction between the users and the system. Use cases help to further define system scope and boundaries. They are usually in the form of a diagram, along with a textual description of the interaction taking place. Figure 2-1 shows a generic diagram that consists of two actors represented by stick figures, the system represented by a rectangle, and use cases depicted by ovals inside the system boundaries.




Figure 2-1. Generic use case diagram with two actors and three use cases

Use cases are developed from the SRS document. The actor is any outside entity that interacts with the system. An actor could be a human user (for instance, a rental agent), another software system (for instance, a software billing system), or an interface device (for instance, a temperature probe). Each interaction that occurs between an actor and the system is modeled as a use case.

The sample use case shown in Figure 2-2 was developed for the flight-booking application introduced in the previous section. It shows the use case diagram for the requirement “Customers can search for flights based on destination and departure times.”




Figure 2-2. View Flight Info use case

Along with the graphical depiction of the use case, many designers and software developers find it helpful to provide a textual description of the use case. The textual description should be succinct and focused on what is happening, not on how it is occurring. Sometimes any preconditions or postconditions associated with the use case are also identified. The following text further describes the use case diagram shown in Figure 2-2.

- *Description:* A customer views the flight information page. The customer enters flight search information. After submitting the search request, the customer views a list of flights matching the search criteria.
- *Preconditions:* None.
- *Postconditions:* The customer has the opportunity to log in and proceed to the flight-booking page.

As another example, take a look at the Reserve Seat use case shown in Figure 2-3.




Figure 2-3. Reserve Seat use case diagram

The following text further describes the use case diagram shown in Figure 2-3.

- *Description:* The customer enters the flight number and indicates the seats being requested. After the customer submits the request, some confirmation information is displayed.
- *Preconditions:* The customer has looked up the flight information. The customer has logged in and is viewing the flight-booking screen.
- *Postconditions:* The customer is sent a confirmation e-mail outlining the flight details and the cancellation policy.

As you can see from Figure 2-3, certain relationships can exist between use cases. The Reserve Seat use case includes the View Flight Info use case. This relationship is useful because you can use the View Flight Info use case independently of the Reserve Flight use case. This is called inclusion. You cannot use the Reserve Seat use case independently of the View Flight Info use case, however. This is important information that will affect how you model the solution.

Another way that use cases relate to each other is through extension. You might have a general use case that is the base for other use cases. The base use case is extended by other use cases. For example, you might have a Register Customer use case that describes the core process of registering customers. You could then develop Register Corporate Customer and Register Retail Customer use cases that extend the base use case. The difference between extension and inclusion is that, in extension, the base use case being extended is not used on its own. Figure 2-4 demonstrates how you model this in a use case diagram.




Figure 2-4. Extending use cases

A common mistake when developing use cases is to include actions initiated by the system itself. The emphasis of the use case is on the interaction between external entities and the system. Another common mistake is to include a description of the technical requirements of the system. Remember that use cases do not focus on how the system will perform the functions, but rather on what functions need to be incorporated in the system from the user's standpoint.

After you have developed the use cases of the system, you can begin to identify the internal system objects that will carry out the system's functional requirements. You do this using a class diagram.

ACTIVITY 2-1. CREATING A USE CASE DIAGRAM

After completing this activity, you should be familiar with the following:

- how to produce a use case diagram to define a system's scope
- how to use a UML modeling tool to create and document a use case diagram

Examining the SRS

The software user group you belong to has decided to pool its resources and create a lending library. Lending items include books, movies, and video games. Your task is to develop the application that will keep track of the loan item inventory and the lending of items to the group members. After interviewing the group's members and officers, you have developed a SRS document that includes the following functional requirements:

- Only members of the user group can borrow items.
- Books can be borrowed for four weeks.
- Movies and games can be borrowed for one week.
- Items can be renewed if no one is waiting to borrow them.
- Members can only borrow up to four items at the same time.
- A reminder is e-mailed to members when an item becomes overdue.
- A fine is charged for overdue items.
- Members with outstanding overdue items or fines can't borrow new items.
- A secretary is in charge of maintaining item inventory and purchasing items to add to the inventory.
- A librarian has been appointed to track lending and send overdue notices.
- The librarian is also responsible for collecting fines and updating fine information.

The next steps are to analyze the SRS to identify the actors and use cases.

1. By examining the SRS document, identify which of the following will be among the principal actors interacting with the system:
 - A. Member
 - B. Librarian
 - C. Book
 - D. Treasurer
 - E. Inventory
 - F. E-mail
 - G. Secretary
2. Once you have identified the principal actors, you need to identify the use cases for the actors. Identify the actor associated with the following use cases:
 - A. Request Item
 - B. Catalog Item
 - C. Lend Item
 - D. Process Fine

See the end of the chapter for Activity 2-1 answers.

Creating a Use Case Diagram

Although it is possible to create UML diagrams by hand or on a whiteboard, most programmers will eventually turn to a diagramming tool or a Computer-Aided Software Engineering (CASE) tool. CASE tools help you construct professional-quality diagrams and enable team members to easily share and augment the diagrams. There are many CASE tools on the market, including Microsoft Visio. Before choosing a CASE tool, you should thoroughly evaluate if it meets your needs and is sufficiently flexible. A lot of the advanced features associated with high-end CASE tools are difficult to work with, so you spend more time figuring out how the CASE tool works than documenting your design.

A good diagramming tool to learn on is UMLet. UMLet enables you to create UML diagrams without adding a lot of advanced features associated with the high-end CASE tools. Best of all, UMLet is a free open-source tool and can be downloaded from www.umlet.com.

■ **Note** These activities use the UMLet 11.5.1 stand-alone edition. This also requires Java 1.6 available at www.java.com.

After downloading and installing UMLet, you can complete the following steps (if you do not want to use a tool, you can create the Use Case diagram by hand):

1. Start UMLet by double clicking the UMLet.jar file. You are presented with three windows. The main window is the design surface, the upper right window contains the UML object templates, and the lower right window is where you change or add properties to the objects.
2. Locate the actor template in the upper right window (see Figure 2-5). Double click the actor template. An actor will appear in the upper left corner of the design surface.




Figure 2-5. Locating the actor template

3. If not already selected, select the actor shape on the design surface. In the lower right window, change the name of the actor shape to Member.
4. Repeat the procedures to add a Secretary and a Librarian actor.
5. From the Template window, double click the Use case 1 shape to add it to the design surface. Change the name of the use case to Request Item.
6. Repeat step 5 for two more use cases. Include a Catalog Item use case that will occur when the Secretary adds new items to the library inventory database. Add a Lend Item use case that will occur when the Librarian processes a request for an item.
7. From the Template window, double click the Empty Package shape and change the name to Library Loan System. Right click on the shape in the design surface and change the background color to white. Move the use case shapes inside the Library Loan System shape (see Figure 2-6).




Figure 2-6. Placing the use cases inside the system boundary

8. From the Template window, double click on the Communications Link shape. It is the line with no arrow heads (see Figure 2-7). On the design surface, attach one end to the Member shape and the other end to the Request Item shape.




Figure 2-7. Locating the Communications Link shape

9. Repeat step 8 two times to create a Communication Link shape between the Librarian and the Lend Item shapes as well as a Communication Link shape between the Secretary and the Catalog Item shapes.
10. From the Templates window, double click the Extends Relationship arrow. Attach the tail end of the Extends arrow to the Lend Item use case and attach the head of the arrow to the Request Item use case.
11. Your completed diagram should be similar to the one shown in Figure 2-8. Save the file as UMLAct2_1 and exit UMLet.




Figure 2-8. Completed use case diagram

Understanding Class Diagrams

The concepts of classes and objects are fundamental to OOP. An object is a structure for incorporating data and the procedures for working with the data. These objects implement the functionality of an object-oriented program. Think of a class as a blueprint for an object and an object as an instance of a class. A class defines the structure and the methods that objects based on the class type will contain.

Designers identify a potential list of classes that they will need to develop from the SRS and the use case diagrams. One way you identify the classes is by looking at the noun phrases in the SRS document and the use case descriptions. If you look at the documentation developed thus far for the airline booking application, you can begin to identify the classes that will make up the system. For example, you can develop a Customer class to work with the customer data and a Flight class to work with the flight data.

A class is responsible for managing data. When defining the class structure, you must determine what data the class is responsible for maintaining. The class attributes define this information. For example, the Flight class will have attributes for identifying the flight number, departure time and date, flight duration, destination, capacity, and seats available. The class structure must also define any operations that will be performed on the data. An example of an operation the Flight class is responsible for is updating the seats available when a seat is reserved.

A class diagram can help you visualize the attributes and operations of a class. Figure 2-9 is an example of the class diagram for the Flight class used in the flight booking system example. A rectangle divided into three sections represents the class. The top section of the rectangle shows the name of the class, the middle section lists the attributes of the class, and the bottom section lists the operations performed by the class.

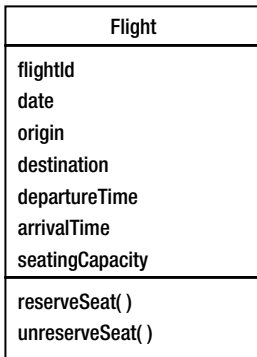


Figure 2-9. Flight class diagram

Modeling Object Relationships

In OOP, when the program executes, the various objects work together to accomplish the programming tasks. For example, in the flight booking application, in order to reserve a seat on the flight, a Reservation object must interact with the Flight object. A relationship exists between the two objects, and this relationship must be modeled in the class structure of the program. The relationships among the classes that make up the program are modeled in the class diagram. Analyzing the verb phrases in the SRS often reveals these relationships (this is discussed in more detail in Chapter 3). The following sections examine some of the common relationships that can occur between classes and how the class diagram represents them.

Association

When one class refers to or uses another class, the classes form an association. You draw a line between the two classes to represent the association and add a label to indicate the name of the association. For example, a seat is associated with a flight in the flight-booking-application, as shown in Figure 2-10.



Figure 2-10. Class associations

Sometimes a single instance of one class associates with multiple instances of another class. This is indicated on the line connecting the two classes. For example, when a customer makes a reservation, there is an association between the Customer class and the Reservation class. A single instance of the Customer class may be associated with multiple instances of the Reservation class. The *n* placed near the Reservation class indicates this multiplicity, as shown in Figure 2-11.




Figure 2-11. Indicating multiplicity in a class diagram

A situation may also exist where an instance of a class may be associated with multiple instances of the same class. For example, an instance of the `Pilot` class represents the captain while another instance of the `Pilot` class represents the co-pilot. The pilot manages the co-pilot. This scenario is referred to as a self-association and is modeled by drawing the association line from the class back to itself, as shown in Figure 2-12.




Figure 2-12. A self-associating class

Inheritance

When multiple classes share some of the same operations and attributes, a base class can encapsulate the commonality. The child class then inherits from the base class. This is represented in the class diagram by a solid line with an open arrowhead pointing to the base class. For example, a `CorporateCustomer` class and a `RetailCustomer` class could inherit common attributes and operations from a base `Customer` class, as shown in Figure 2-13.




Figure 2-13. Documenting inheritance

Aggregation

When a class is formed by a composition of other classes, they are classified as an aggregation. This is represented with a solid line connecting the classes in a hierarchical structure. Placing a diamond on the line next to a class in the diagram indicates the top level of the hierarchy. For example, an inventory application designed to track plane parts for the plane maintenance department could contain a `Plane` class that is a composite of various part classes, as shown in Figure 2-14.




Figure 2-14. Depicting aggregations

Association Classes

As the classes and the associations for a program are developed, there may be a situation where an attribute can't be assigned to any one class, but is a result of an association between classes. For example, the parts inventory application mentioned previously may have a `Part` class and a `Supplier` class. Because a part can have more than one supplier and the supplier supplies more than one part, where should the price attribute be located? It does not fit nicely as an attribute for either class, and it should not be duplicated in both classes. The solution is to develop an association class that manages the data that is a product of the association. In this case, you would develop a `PartPrice` class. The relationship is modeled with a dashed line drawn between the association and the association class, as shown in Figure 2-15.




Figure 2-15. An association class

Figure 2-16 shows the evolving class diagram for the flight booking application. It includes the classes, attributes, and relationships that have been identified for the system. The operations associated with the classes will be developed in Chapter 3.

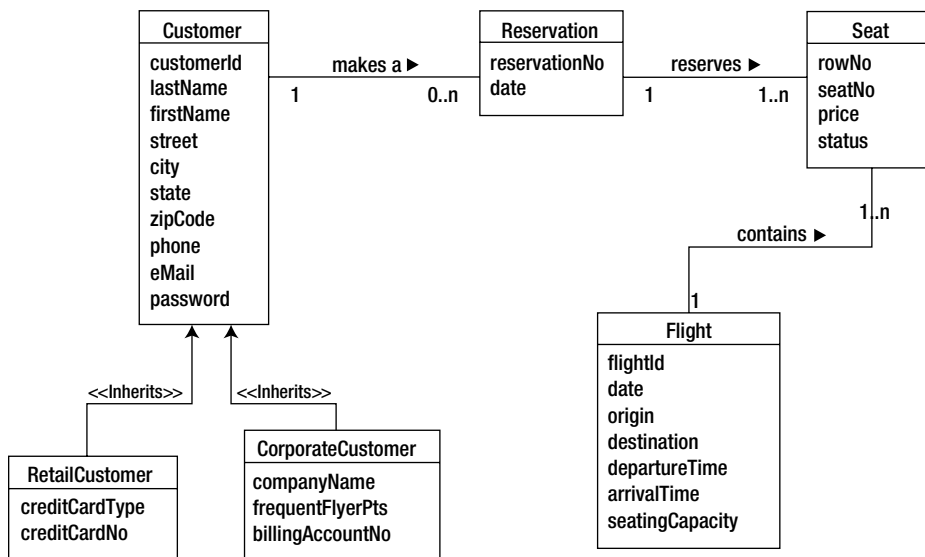


Figure 2-16. Flight booking class diagram

ACTIVITY 2-2. CREATING A CLASS DIAGRAM

After completing this activity, you should be familiar with the following:

- how to determine the classes that need to be constructed by examining the use case and the system scope documentation
- how to use a UML modeling tool to create a class diagram

Identifying Classes and Attributes

Examine the following scenario developed for a use case from the user group library application:

After viewing the list of available loan items, members request an item to check out on loan. The librarian enters the member number and retrieves information about outstanding loans and any unpaid fines. If the member has fewer than four outstanding loans and does not have any outstanding fines, the loan is processed. The librarian retrieves information about the loan item to determine if it is currently on loan. If the item is available, it is checked out to the member.

1. By identifying the nouns and noun phrases in the use case scenario, you can get an idea of what classes you must include in the system to perform the tasks. Which of the following items would make good candidate classes for the system?
 - A. Member
 - B. Item
 - C. Librarian
 - D. Number
 - E. Fine
 - F. Checkout
 - G. Loan
2. At this point, you can start identifying attributes associated with the classes being developed. A Loan class will be developed to encapsulate data associated with an item out on loan. Which of the following would be possible attributes for the Loan class?
 - A. MemberNumber
 - B. MemberPhone
 - C. ItemNumber
 - D. ReturnDate
 - E. ItemCost
 - F. ItemType

See the end of the chapter for Activity 2-2 answers.

Creating a Class Diagram

To create a class diagram using UML Modeler, follow these steps (you can also create it by hand):

1. Start UMLet. You are presented with three windows. The main window is the design surface, the upper right window contains the UML object templates, and the lower right window is where you change or add properties to the objects.
2. Locate the SimpleClass template in the upper right window (see Figure 2-17). Double click the SimpleClass template. A SimpleClass will appear in the upper left corner of the design surface.

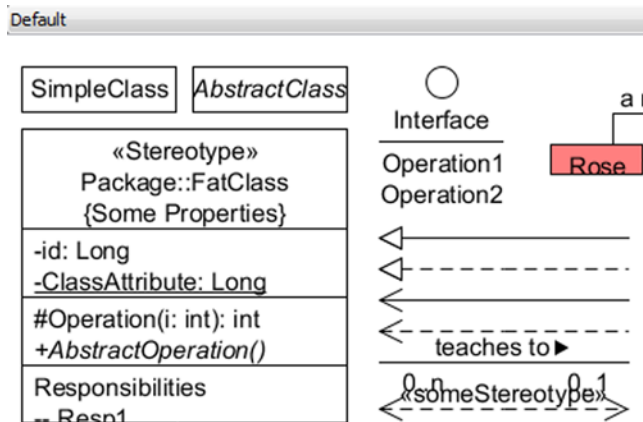


Figure 2-17. Adding a class shape

3. In the lower right properties window, change the class name to Member.
4. Repeat the procedure for a Loan, Item, Book, and Movie class.
5. Locate the association template in the upper right window (see Figure 2-18). Double click the association template. An association will appear in the upper left corner of the design surface.




Figure 2-18. Adding an association shape

- Attach the left end of the association shape to the `Member` class and the right end to the `Loan` class shape. Select the association shape and update the properties in the properties window so that they match Figure 2-19.




Figure 2-19. Updating association properties

- Repeat steps 5 and 6 to create a “Contains a” association shape between the `Loan` class and the `Item` class. This should be a one-to-one association.
- Locate the generalization arrow template in the upper right window (see Figure 2-20). Double click the generalization shape. A generalization shape will appear in the upper left corner of the design surface.




Figure 2-20. Adding a generalization arrow

- Attach the tail end of the generalization arrow to the `Book` class and the head end to the `Item` class shape. Select the generalization arrow and update the properties in the properties window so that they match Figure 2-21.




Figure 2-21. Updating generalization properties

10. Repeat steps 8 and 9 to show that the **Movie** class inherits from the **Item** class.
11. Click on the **Member** class in the design window. In the properties window, add the `memberNumber`, `firstName`, `lastName`, and `eMail` attributes as shown in Figure 2-22.




Figure 2-22. Adding class attributes

12. Your completed diagram should be similar to Figure 2-23. Save the file as `UMLAct2_2`.




Figure 2-23. Completed class diagram

Summary

In this chapter, you were introduced to the goals of the object-oriented design process and UML. You learned about some of the design documents and diagrams produced using UML. These include the SRS, which defines the scope of the system; use case diagrams, which define the system boundaries and identify the external entities that will use the system; and class diagrams, which model the structure of the classes that will be developed to implement the system.

You saw how modeling the class structure of your applications includes identifying the necessary classes, identifying the attributes of these classes, and establishing the structural relationships required among the classes. In Chapter 3, you will continue your study of object-oriented design. In particular, you will look at modeling how the objects in your applications will collaborate to carry out the functionality of the application.

ACTIVITY ANSWERS

Activity 2–1 Answers

1. A, B, G. The actors are Member, Librarian, and Secretary.
2. A. Member, B. Secretary, C. Librarian, D. Librarian. The Request Item use case goes with Member, the Catalog Item use case goes with Secretary, the Lend Item use case goes with Librarian, and the Process Fine use case goes with Librarian.

Activity 2–2 Answers

1. A, B, C, E, G. The candidate classes are Member, Item, Librarian, Fine, and Loan.
 2. A, C, D. The attributes associated with the Loan class are MemberNumber, ItemNumber, and ReturnDate.
-



Designing OOP Solutions: Modeling the Object Interaction

The previous chapter focused on modeling the static (organizational) aspects of an OOP solution. It introduced and discussed the methodologies of the UML. It also looked at the purpose and structure of use case diagrams and class diagrams. This chapter continues the discussion of UML modeling techniques and focuses on modeling the dynamic (behavioral) aspects of an OOP solution. The focus in this chapter is on how the objects in the system must interact with each other and what activities must occur to implement the solution. This is an important aspect of the modeling processes. These models will serve as the basis for coding the various methods of the classes (identified in Chapter 2) that will make up the software application.

After reading this chapter, you should be familiar with the following:

- the purpose of scenarios and how they extend the use case models
- how sequence diagrams model the time-dependent interaction of the objects in the system
- how activity diagrams map the flow of activities during application processing
- the importance of graphical user interface design and how it fits into the object-oriented design process

Understanding Scenarios

Scenarios help determine the dynamic interactions that will take place between the objects (class instances) of the system. A scenario is a textual description of the internal processing needed to implement the functionality documented by a use case. Remember that a use case describes the functionality of the system from the viewpoint of the system's external users. A scenario details the execution of the use case. In other words, its purpose is to describe the steps that must be carried out internally by the objects making up the system.

Figure 3-1 shows a Process Movie Rental use case for a video rental application. The following text describes the use case:

- **Preconditions:** The customer makes a request to rent a movie from the rental clerk. The customer has a membership in the video club and supplies the rental clerk with her membership card and personal identification number (PIN). The customer's membership is verified. The customer information is displayed, and the customer's account is verified to be in good standing.
- **Description:** The movie is confirmed to be in stock. Rental information is recorded, and the customer is informed of the due date.
- **Post conditions:** None.




Figure 3-1. *Process Movie Rental use case*

The following scenario describes the internal processing of the Process Movie Rental use case:

- The movie is verified to be in stock.
- The number of available copies in stock is decremented.
- The due date is determined.
- The rental information is recorded. This information includes the movie title, copy number, current date, and due date.
- The customer is informed of the rental information.

This scenario describes the best possible execution of the use case. Because exceptions can occur, a single use case can spawn multiple scenarios. For example, another scenario created for the Process Movie Rental use case could describe what happens when a movie is not in stock.

After you map out the various scenarios for a use case, you can create interaction diagrams to determine which classes of objects will be involved in carrying out the functionality of the scenarios. The interaction diagram also reveals what operations will be required of these classes of objects. Interaction diagrams come in two flavors: sequence diagrams and collaboration diagrams.

Introducing Sequence Diagrams

A sequence diagram models how the classes of objects interact with each other over time as the system runs. The sequence diagram is a visual, two-dimensional model of the interaction taking place and is based on a scenario. Figure 3-2 shows a generic sequence diagram.




Figure 3-2. *Generic sequence diagram*

As Figure 3-2 demonstrates, the flow of messages from object to object is represented horizontally. The time flow of the interactions taking place is depicted vertically, starting from the top and progressing downward. Objects are placed next to each other from left to right according to the calling order. A dashed line extends from each of them downward. This dashed line represents the lifeline of the object. Rectangles on the lifeline represent activations of the object. The height of the rectangle represents the duration of the object's activation.

In OOP, objects interact by passing messages to each other. An arrow starting at the initiating object and ending at the receiving object depicts the interaction. A dashed arrow drawn back to the initiating object represents a return message. The messages depicted in the sequence diagram will form the basis of the methods of the classes of the system. Figure 3-3 shows a sample sequence diagram for the Process Movie Rental scenario presented in the previous section. At this point, the diagram only models the case where the movie is in stock.

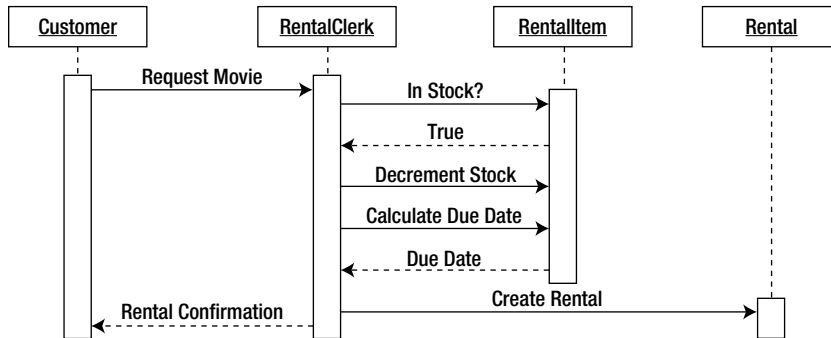


Figure 3-3. Process Movie Rental sequence diagram

As you analyze the sequence diagram, you gain an understanding of the classes of objects that will be involved in carrying out the program processing; you will also understand what methods you will need to create and attach to those classes. You should also model the classes and methods depicted in the sequence diagram in the class diagram. These design documents must be continually cross-referenced and revised as necessary.

The sequence diagram in Figure 3-3 reveals that there will be four objects involved in carrying out the Process Movie Rental scenario.

- The Customer object is an instance of the Customer class and is responsible for encapsulating and maintaining the information pertaining to a customer.
- The RentalClerk object is an instance of the RentalClerk class and is responsible for managing the processing involved in renting a movie.
- The RentalItem object is an instance of the RentalItem class and is responsible for encapsulating and maintaining the information pertaining to a video available for rent.
- The Rental object is an instance of the Rental class and is responsible for encapsulating and maintaining the information pertaining to a video currently being rented.

Message Types

By analyzing the sequence diagram, you can determine what messages must be passed between the objects involved in the processing. In OOP, messages are passed synchronously or asynchronously.

When messages are passed synchronously, the sending object suspends processing and waits for a response before continuing. A line drawn with a closed arrowhead in the sequence diagram represents synchronous messaging.

When an object sends an asynchronous message, the object continues processing and is not expecting an immediate response from the receiving object. A line drawn with an open arrowhead in the sequence diagram represents asynchronous messaging. A dashed arrow usually depicts a response message. These lines are shown in Figure 3-4.




Figure 3-4. Different types of messages

By studying the sequence diagram for the Process Movie Rental scenario shown in Figure 3-3, you can see the types of messages that must be passed. For example, the RentalClerk object initiates a synchronous message with the RentalItem object, requesting information as to whether a copy of the movie is in stock. The RentalItem object then sends a response back to the RentalClerk object, indicating a copy is in stock. This needs to be synchronous because the RentalClerk is waiting for the response to process the request. An example of an asynchronous message would be if someone signs up for an e-mail alert when a movie is returned. In this case, a message would be sent to fire off the e-mail but there is no need for a response.

Recursive Messages

In OOP, it is not uncommon for an object to have an operation that invokes another object instance of itself. This is referred to as recursion. A message arrow that loops back toward the calling object represents recursion in the sequence diagram. The end of the arrow points to a smaller activation rectangle, representing a second object activation drawn on top of the original activation rectangle (see Figure 3-5). For example, an Account object calculates compound interest for overdue payments. To calculate the interest over several compound periods, it needs to invoke itself several times.

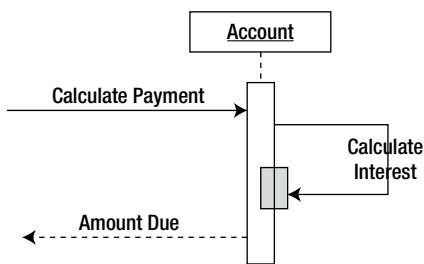


Figure 3-5. Diagramming a recursive message

Message Iteration

Sometimes, a message call is repeated until a condition is met. For example, when totaling rental charges, an Add method is called repeatedly until all rentals charged to the customer have been added to the total. In programming terminology, this is iteration. A rectangle drawn around the iterating messages represents an iteration in a sequence diagram. The binding condition of the iteration is depicted in the upper-left corner of the rectangle. Figure 3-6 shows an example of an iteration depicted in a sequence diagram.




Figure 3-6. Depicting an iterative message

Message Constraints

Message calls between objects may have a conditional constraint attached to them. For example, customers must be in good standing in order to be allowed to rent a movie. You place the condition of the constraint within brackets ([]) in the sequence. The message will be sent only if the condition evaluates to true (see Figure 3-7).




Figure 3-7. Identifying conditional constraints

Message Branching

When conditional constraints are tied to message calling, you often run into a branching situation where, depending on the condition, different messages may be invoked. Figure 3-8 represents a conditional constraint when requesting a movie rental. If the status of the rental item is “in stock,” a message is sent to the Rental object to create a rental. If the status of the rental item is “out of stock,” a message is sent to the Reservation object to create a reservation. A rectangle (labeled as alt in Figure 3-8) drawn around the messages shows the alternate paths that can occur depending on the condition.

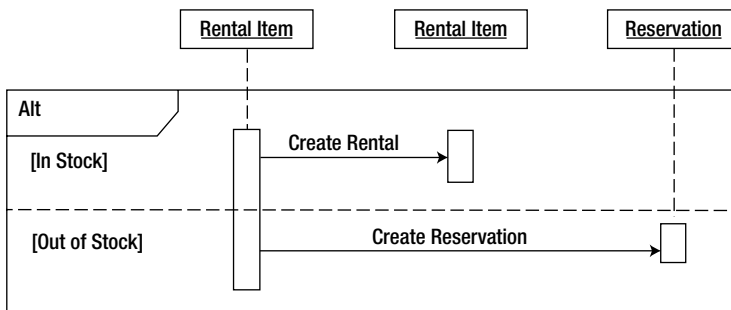


Figure 3-8. Branching messages in a sequence diagram

ACTIVITY 3-1. CREATING A SEQUENCE DIAGRAM

After completing this activity, you should be familiar with the following:

- producing a sequence diagram to model object interaction
- using a UML modeling tool to create a sequence diagram
- adding methods to the class diagram

Examining the Scenario

The following scenario was created for a use case in the user group library application introduced in Activity 2-1. It describes the processing involved when a member borrows an item from the library.

When a member makes a request to borrow an item, the librarian checks the member's records to make sure no outstanding fines exist. Once the member passes these checks, the item is checked to see if it is available. Once the item availability has been confirmed, a loan is created recording the item number, member number, checkout date, and return date.

1. By examining the noun phrases in the scenario, you can identify which objects will be involved in carrying out the processing. The objects identified should also have a corresponding class depicted in the class diagram that has been previously created. From the scenario depicted, identify five objects that will carry out the processing.
2. After the objects have been identified and cross-referenced with the class diagram, the next step is to identify the messaging that must occur between these objects to carry out the task. You can look at the verb phrases in the scenario to help identify these messages. For example, the "request to borrow item" phrase indicates a message interaction between the Member object and the Librarian object. What are the other interactions depicted in the scenario?

See the end of the chapter for Activity 3-1 answers.

Creating a Sequence Diagram

Follow these steps to create a sequence diagram using UMLet:

1. Start UMLet. Locate the drop-down list at the top of the template window. Change the template type to Sequence (see Figure 3-9).

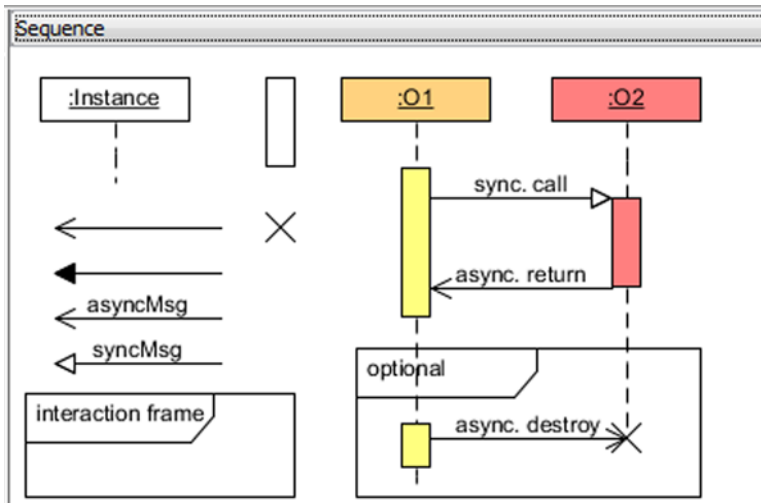


Figure 3-9. Changing template shape types

2. Double-click the Instance shape in the template window. An Instance shape will appear in the upper left corner of the design surface. In the properties window, change the name of the shape to Member.
3. From the shapes window, locate the lifeline and activation shapes (see Figure 3-10) and add them to the Member instance, as shown in Figure 3-11.




Figure 3-10. Locating the lifeline and activation shapes




Figure 3-11. Adding shapes to the sequence diagram

- Repeat steps 2 and 3 to add a Librarian, LoanHistory, Item, and Loan object to the diagram. Lay them out from left to right as shown in Figure 3-12.

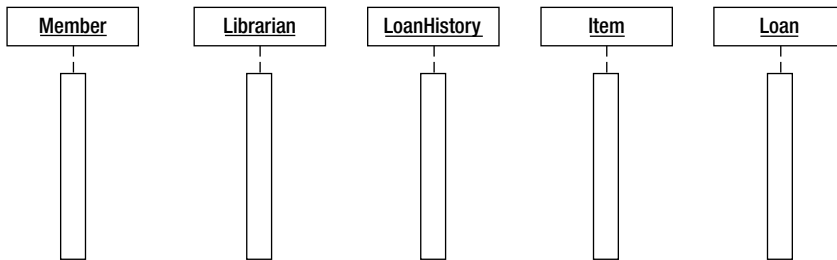


Figure 3-12. Object layout in the sequence diagram

- From the shapes template window, double-click the Sequence Message arrow shape. Attach the tail end of the arrow to the Member object's lifeline and the head of the arrow to the Librarian object's lifeline. In the properties window, change the name of the message to "Request Item."
- To create a return arrow, double-click on the solid arrow with the open arrow head in the shapes template window. In the properties window, change the first line to `lt=.` This should change the arrow from solid to dash. Attach the tail end to the Librarian object and the head end to the Member object. Change the name to "Return Loan Info." Your diagram should look similar to Figure 3-13.




Figure 3-13. Message layout in the sequence diagram

- Repeat steps 5 and 6 to create a message from the Librarian object to the LoanHistory object. Name the calling message (the solid line) "Check History." Name the return message (the dashed line) "Return History."
- Create a message from the Librarian object to the Item object. Name the calling message "Check Availability." Name the return message "Return Availability Info."
- Create a message from the Librarian object to the Item object. Name the calling message "Update Status." Name the return message "Return Update Confirmation."
- Create a message from the Librarian object to the Loan object. Name the calling message "Create Loan." Name the return message "return loan confirmation."

11. Rearrange the shapes so that your diagram looks similar to Figure 3-14. Save the diagram as UML_Act3_1.

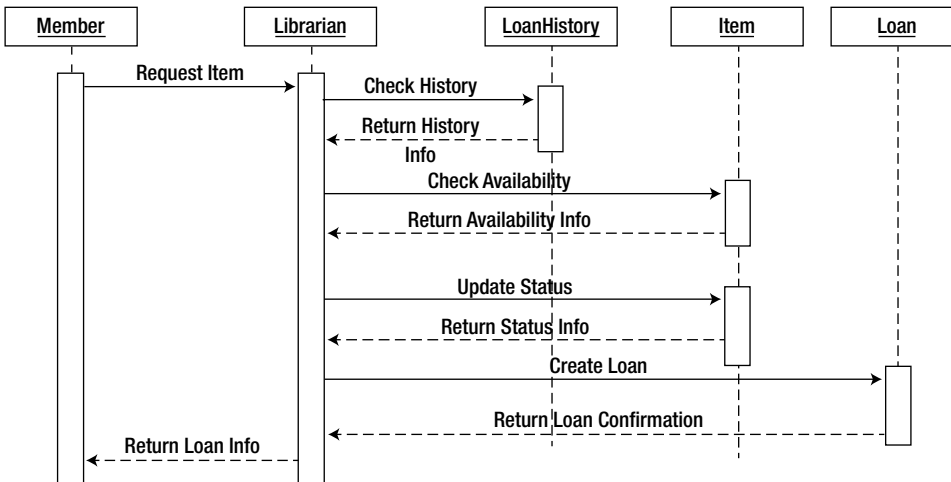


Figure 3-14. Completed sequence diagram

Adding Methods to the Class Diagram

After you have developed the sequence diagram, you begin to gain an understanding of the methods that must be included in the various classes of the application. You achieve the message interaction depicted in the sequence diagram by a method call from the initiating object (client) to the receiving object (server). The method being called is defined in the class that the server object is instantiated as. For example, the “check availability” message in the sequence diagram indicates that the `Item` class needs a method that processes this message call.

Follow these steps to add the methods:

1. In UMLet, choose **File** ► **New** to create a new diagram. Locate the drop-down list at the top of the template window. Change the template type to **Class**.
2. Double-click on the **Simple Class** shape template. Select the shape in the design window.
3. In the properties window, change the name of the class to `Item`. Underneath the name in the properties window, enter two dashes. This will create a new section in the class shape. This section is where you enter the attributes of the class.
4. In the properties window, add the `ItemNumber` attribute to the class followed by two more dashes. This creates a third section in the class shape that is used to add the methods of the class.
5. Add a `CheckAvailability` and an `UpdateStatus` method to the class as shown in Figure 3-15.

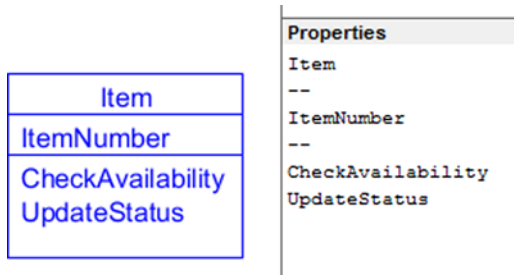


Figure 3-15. Adding methods to a class

6. Save the diagram as UML_Act3_1b.

Understanding Activity Diagrams

An activity diagram illustrates the flow of activities that need to occur during an operation or process. You can construct the activity diagram to view the workflow at various levels of focus.

- A high, system-level focus represents each use case as an activity and diagrams the workflow among the different use cases.
- A mid-level focus diagrams the workflow occurring within a particular use case.
- A low-level focus diagrams the workflow that occurs within a particular operation of one of the classes of the system.

The activity diagram consists of the starting point of the process represented by a solid circle and transition arrows representing the flow or transition from one activity to the next. Rounded rectangles represent the activities, and a bull's eye circle represents the ending point of the process. For example, Figure 3-16 shows a generic activity diagram that represents a process that starts with activity A, proceeds to activity B, and concludes.




Figure 3-16. Generic activity diagram

Decision Points and Guard Conditions

Often, one activity will conditionally follow another. For example, in order to rent a video, a PIN verifies membership. An activity diagram represents conditionality by a decision point (represented by a diamond) with the guard condition (the condition that must be met to proceed) in brackets next to the flow line (see Figure 3-17).




Figure 3-17. Indicating decision points and guard conditions

Parallel Processing

In some cases, two or more activities can run in parallel instead of sequentially. A solid, bold line drawn perpendicularly to the transition arrow represents the splitting of the paths. After the split, a second solid, bold line represents the merge. Figure 3-18 shows an activity diagram for the processing of a movie return. The order in which the Increment Inventory and the Remove Rental activities occur does not matter. The parallel paths in the diagram represent this parallel processing.




Figure 3-18. Parallel processing depicted in an activity diagram

Activity Ownership

The activity diagram's purpose is to model the control flow from activity to activity as the program processes. The diagrams shown thus far do not indicate which objects have responsibility for these activities. To signify object ownership of the activities, you segment the activity diagram into a series of vertical partitions (also called swim lanes). The object role at the top of the partition is responsible for the activities in that partition. Figure 3-19 shows an activity diagram for processing a movie rental, with swim lanes included.

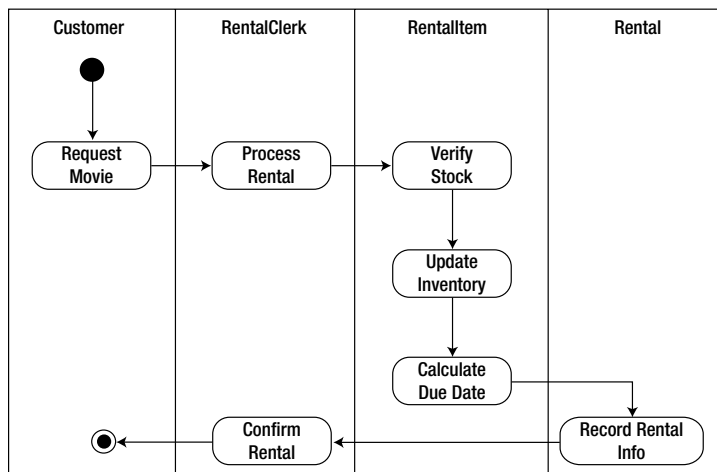


Figure 3-19. Swim lanes in an activity diagram

ACTIVITY 3-2. CREATING AN ACTIVITY DIAGRAM

After completing this activity, you should be familiar with the following:

- using an activity diagram to model control flow as the program completes an activity
- using a UML modeling tool to create an activity class diagram

Identifying Objects and Activities

Examine the following scenario developed for a use case from the user group library application:

After viewing the list of available loan items, a member requests an item to check out on loan. The librarian enters the member number and retrieves information about outstanding loans and any unpaid fines. If the member has fewer than four outstanding loans and does not have any outstanding fines, the loan is processed. The librarian retrieves information on the loan item to determine if it is currently on loan. If the item is available, it is checked out to the member.

By identifying the nouns and noun phrases in the use case scenario, you can get an idea of what objects will perform the tasks in carrying out the activities. Remember that these objects are instances of the classes identified in the class diagram. The following objects will be involved in carrying out the activities: Member, Librarian, LoanHistory, Item, and Loan.

The verb phrases help identify the activities carried out by the objects. These activities should correspond to the methods of the classes in the system. Match the following activities to the appropriate objects:

- A. Request Movie
- B. Process Rental
- C. Check Availability
- D. Check Member's Loan Status
- E. Update Item Status
- F. Calculate Due Date
- G. Record Rental Info
- H. Confirm Rental

See the end of the chapter for Activity 3-2 answers.

Creating an Activity Diagram

Follow these steps to create an activity diagram using UMLet:

1. Start UMLet. Locate the drop-down list at the top of the template window. Change the template type to Activity.
2. Double-click the System box shape in the template window. A System box shape will appear in the upper left corner of the design surface. In the properties window, change the name of the shape to Member to represent the Member partition.

- Repeat step 2 to add a partition for the Librarian, LoanHistory, Item, and Loan objects. Align the partitions from left to right as in Figure 3-20.

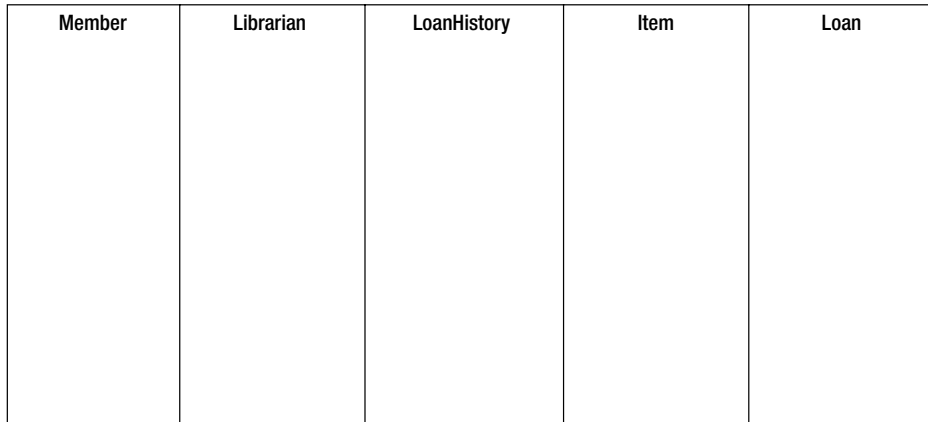


Figure 3-20. Creating the activity diagram partitions

- From the Shapes window, double-click the Initial State shape and add it to the Member partition. Below the Initial State in the Member partition, add a State shape. Rename the State to “request item.” Add a transition shape (arrow) from the Initial State to the Request Item action state.
- Under the Librarian partition, add a Process Loan state and a Transition shape from the Request Item state to the Process Loan state.
- Under the LoanHistory partition, add a Check Member Status action state and a Transition shape from the Process Loan action to the Check Member Status action state.
- From the Shapes window, double-click the Conditional Branch shape (diamond) and add it to the LoanHistory partition below the Check Member Status action state. Add a Transition shape from the Check Member Status state to the Conditional Branch. From the Conditional Branch, add a Transition shape to a Deny Loan state under the Librarian partition. Add a label to the Transition shape with a condition of “fail.” Also add a Transition shape to a Check Item Status action state under the Item partition with a label condition of “pass.” Your diagram should be similar to Figure 3-21.




Figure 3-21. Adding a branching condition

8. Repeat step 7 to create a Conditional Branch from the Check Item Status state. If the item is in stock, add a Transition shape to an Update Item Status state under the Item partition. If the item is out of stock, add a Transition shape to the Deny Loan state under the Librarian partition.
9. From the Update Item Status state, add a Transition shape to a Record Loan Info state under the Loan partition.
10. From the Record Loan Info state, add a Transition shape to a Confirm Loan state under the Librarian partition.
11. From the Shapes window, click the Final State shape and add it to the bottom of the Member partition. Add a Transition shape from the Deny Loan state to the Final action state. Add another Transition shape from the Confirm Loan state to the Final action state.

Your completed diagram should resemble the one shown in Figure 3-22. Save the diagram as UMLAct3_2 and exit UMLet.




Figure 3-22. Completed activity diagram

Exploring GUI Design

Thus far, the discussions of object-oriented analysis and design have focused on modeling the functional design and the internal processing of the application. Successful modern software applications rely on a rich set of graphical user interfaces (GUIs) to expose this functionality to the users of the application.

In modern software systems, one of the most important aspects of an application is how well it interacts with the users. Gone are the days when users would interact with the application by typing cryptic commands at the DOS prompt. Modern operating systems employ GUIs that are, for the most part, intuitive to use. Users have also grown accustomed to the polished interfaces of the commercial office-productivity applications. Users have come to expect the same ease of use and intuitiveness built into custom applications developed in-house by their IT department.

The design of the user interface should not be done haphazardly; rather, it should be planned in conjunction with the business logic design. The success of most applications is judged by the response of the users toward the application. If users are not comfortable when interacting with the application and the application does not improve the productivity of the user, it is doomed to failure. To the user, the interface *is* the application. It does not matter how pristine and clever the business logic code may be; if the user interface is poorly designed and implemented, the application will not be acceptable to the users. Developers would do well to remember that it is the user who drives software development.

Although UML was not specifically designed for GUI design, many software architects and programmers employ UML diagrams to help model the user interface of an application.

GUI Activity Diagrams

The first step in developing a user interface design is to perform a task analysis to discover how users will interact with the system. The task analysis is based on the use cases and scenarios that have been modeled previously. You can then develop activity diagrams to model how the interaction between the user and the system will take place. Using the previous movie rental application example, Figure 3-23 shows an activity diagram modeling the activities the rental clerk goes through to record rental information.

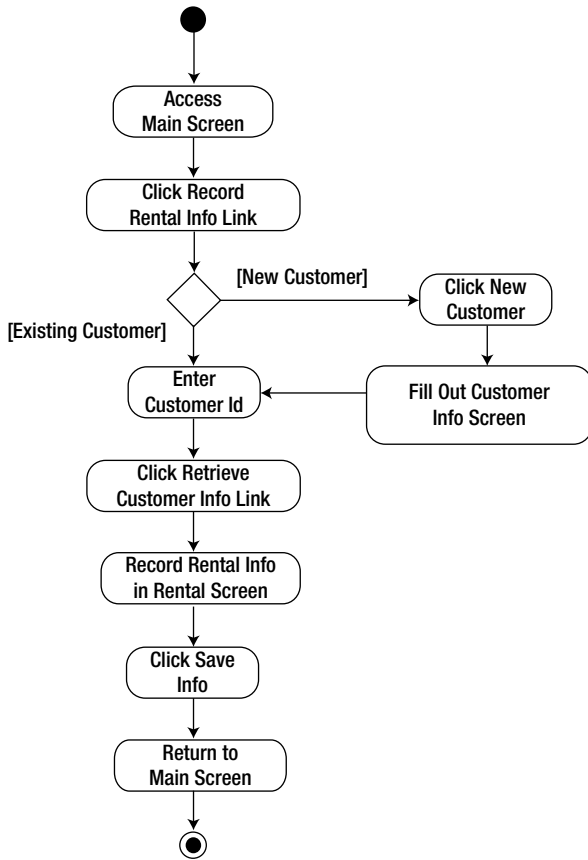


Figure 3-23. GUI modeling with an activity diagram

Interface Prototyping

After you have identified and prioritized the necessary tasks, you can develop a prototype sketch of the various screens that will make up the user interface. Figure 3-24 shows a prototype sketch of the Customer Info screen. Although you can use paper and pencil to develop your diagrams, there are some nice GUI prototyping tools available that offer common GUI design templates and the ability to link screens, plus other useful features. Some popular commercial prototyping tools include Microsoft's SketchFlow, Balsamiq Mockups, and AppMockupTools. Some popular free tools to check out are Pencil Project, Mockingbird, and Prototype Composer. The prototype shown in Figure 3-24 was created using the Pencil Project tool.




Figure 3-24. GUI prototype sketch

Interface Flow Diagrams

Once you have prototyped the various screens, you can use interface flow diagrams to model the relationships and flow patterns among the screens that make up the user interface. Figure 3-25 shows a partial interface flow diagram for the video rental application.

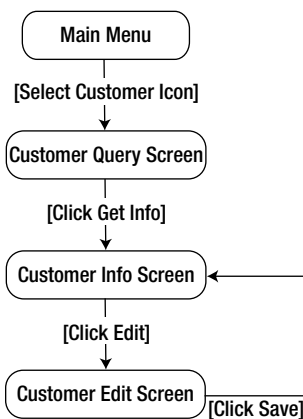


Figure 3-25. Interface flow diagramming

Application Prototyping

Once you have roughed out the screen layout and the design of the user interface, you can develop a simple prototype. The prototype should contain skeleton code that simulates the functionality of the system. At this point, you should not put a great effort into integrating the user interface front end with the business functionality of the application. The idea is to let the users interact with a working prototype to generate feedback on the user interface. It is important

that the users are comfortable with the interface design and that it is intuitive to use. If the interface is cumbersome for the users and does not increase their productivity, the application is doomed to failure.

The processes of refining and testing the user interface will be iterative and most likely will continue through several cycles. Once the user interface design and the internal functional design of the application have been completed and prototyped, the next step in the application development cycle is to start coding the application.

Summary

This chapter introduced scenarios, sequence diagrams, collaboration diagrams, and activity diagrams. You saw how to use these diagrams for modeling object interaction. Additionally, you learned how some of the UML diagrams might be used to help model the user interface of the application.

The goal of this and the previous chapters has been to introduce you to some of the common modeling diagrams and concepts involved in software design and UML. In Chapter 4, you will take the concepts developed thus far and use them to implement a solution design for a sample case study.

ACTIVITY ANSWERS

Activity 3-1 Answers

Member, Librarian, Item, Loan, LoanHistory. These five objects are involved in the processing depicted in the scenario.

The other messaging interactions depicted in the scenario are as follows:

1. The Librarian object checks the lending history of the Member object with the LoanHistory object.
2. The Librarian object checks the availability of the item through the Item object.
3. The Librarian object updates the availability of the item through the Item object.
4. The Librarian creates a Loan object containing loan information.
5. The Librarian returns loan information to the Member object.

Activity 3-2 Answers

A. Member, B. Librarian, C. Item, D. LoanHistory, E. Item, F. Loan, G. Loan, H. Librarian.

The Member object is responsible for the Request Movie activity. The Librarian object is responsible for the Process Rental and Confirm Rental activities. The LoanHistory object is responsible for the Check Member's Loan Status activity. The Item object is responsible for the Check Availability and Update Item Status activities. The Loan object is responsible for the Calculate Due Date and Record Rental Info activities.



Designing OOP Solutions: A Case Study

Designing solutions for an application is not an easy endeavor. Becoming an accomplished designer takes time and conscious effort, which explains why many developers avoid it like the plague. You can study the theories and know the buzzwords, but the only way to truly develop your modeling skills is to roll up your sleeves, get your hands dirty, and start modeling. In this chapter, you follow the process of modeling an office-supply ordering system. Although this is not a complex application, it includes several good use cases and will consist of multiple classes with ample class interactions. By analyzing the case study, you will gain a better understanding of how a model is developed and how the pieces fit together.

After reading this chapter, you should be familiar with the following:

- how to model an OOP solution using UML
- some common OOP design pitfalls to avoid

Developing an OOP Solution

In the case-study scenario, your company currently has no standard way for departments to order office supplies. Each department separately implements its own ordering process. As a result, it is next to impossible to track company-wide spending on supplies, which impacts the ability to forecast budgeting and identify abuses. Another problem with the current system is that it does not allow for a single contact person who could negotiate better deals with the various vendors.

As a result, you have been asked to help develop a company-wide office-supply ordering (OSO) application. To model this system you will complete the following steps:

- create a Software Requirements Specification (SRS)
- develop the use cases
- diagram the use cases
- model the classes
- model the user interface design

Creating the System Requirement Specification

After interviewing the various clients of the proposed system, you develop the SRS. Remember from Chapter 2 that the SRS scopes the system requirements, defines the system boundaries, and identifies the users of the system.

You have identified the following system users:

- *Purchaser*: initiates a request for supplies
- *Department manager*: tracks and approves supply requests from department purchasers
- *Supply vendor processing application*: receives order files generated by the system
- *Purchase manager*: updates the supply catalog, tracks supply requests, and checks in delivered items

You have identified the following system requirements:

- Users must log in to the system by supplying a username and password.
- Purchasers will view a list of supplies that are available to be ordered.
- Purchasers will be able to filter the list of supplies by category.
- Purchasers can request multiple supplies in a single purchase request.
- A department manager can request general supplies for the department.
- Department managers must approve or deny supply requests for their department at the end of each week.
- If department managers deny a request, they must supply a short explanation outlining the reason for the denial.
- Department managers must track spending within their departments and ensure there are sufficient funds for approved supply requests.
- A purchase manager maintains the supply catalog and ensures it is accurate and current.
- A purchase manager checks in the supplies when they are received and organizes the supplies for distribution.
- Supply requests that have been requested but not approved are marked with a status of *pending*.
- Supply requests that have been approved are marked with a status of *approved* and an order is generated.
- Once an order is generated, a file containing the order details is placed in an order queue. Once the order has been placed in the queue, it is marked with a status of *placed*.
- A separate supply-vendor-processing application will retrieve the order files from the queue, parse the documents, and distribute the line items to the appropriate vendor queues. Periodically, the supply-vendor-processing application will retrieve the orders from a vendor queue and send them to the vendor. (This is being developed by a separate team.)
- When all the items of an order are checked in, the order is marked with a status of *fulfilled* and the purchaser is informed that the order is ready for pick up.

Developing the Use Cases

After generating the SRS and getting the appropriate system users to sign off on it, the next task is to develop the use cases, which will define how the system will function from the users' perspective. The first step in developing the use cases is to define the actors. Remember from Chapter 2 that the actors represent the external entities (human or other systems) that will interact with the system. From the SRS, you can identify the following actors that will interact with the system:

- purchaser
- department manager
- purchase manager
- supply vendor processing application

Now that you have identified the actors, the next step is to identify the various use cases with which the actors will be involved. By examining the requirement statements made in the SRS, you can identify the various use cases. For example, the statement "Users must log in to the system by supplying a username and password" indicates the need for a Login use case. Table 4-1 identifies the use cases for the OSO application.

Table 4-1. Use Cases for the OSO Application

Name	Actor(s)	Description
Login	Purchaser, Department manager, Purchase manager	Users see a login screen. They then enter their username and password. They either click Login or Cancel. After login, they see a screen containing product information.
View Supply Catalog	Purchaser, Department manager, Purchase manager	Users see a catalog table that contains a list of supplies. The table contains information such as the supply name, category, description, and cost. Users can filter supplies by category.
Purchase Request	Purchaser, Department manager	Purchasers select items in the table and click a button to add them to their cart. A separate table shows the items in their cart, the number of each item requested and the cost, as well as the total cost of the request.
Department Purchase Request	Department manager	Department managers select items in the table and click a button to add them to their cart. A separate table shows the items in their cart, the number of each item requested and the cost, as well as the total cost of the request.
Request Review	Department manager	Department managers see a screen that lists all pending supply requests for members of their department. They review the requests and mark them as approved or denied. If they deny the request, they enter a brief explanation.
Track Spending	Department manager	Department managers see a screen that lists the monthly spending of department members as well as the running total of the department.

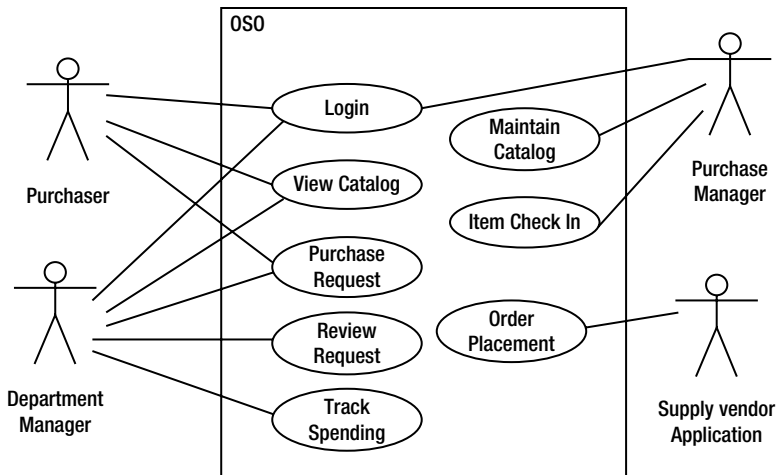
(continued)

Table 4-1. (continued)

Name	Actor(s)	Description
Maintain Catalog	Purchase manager	The purchase manager has the ability to update product information, add products, or mark products as discontinued. The administrator can also update category information, add categories, and mark categories as discontinued.
Item Check In	Purchase manager	The purchase manager sees a screen for entering the order number. The purchase manager then sees the line items listed for the order. The items that have been received are marked. When all the items for an order are received, it is marked as fulfilled.
Order Placement	Supply-vendor-processing application	The supply-vendor-processing application checks the queue for outgoing order files. Files are retrieved, parsed, and sent to the appropriate vendor queue.

Diagramming the Use Cases

Now that you have identified the various use cases and actors, you are ready to construct a diagram of the use cases. Figure 4-1 shows a preliminary use case model developed with UMLet, which was introduced in Chapter 2.

**Figure 4-1.** Preliminary OSO use case diagram

After you have diagrammed the use cases, you now look for any relationships that may exist between the use cases. Two relationships that may exist are the *includes* relationship and the *extends* relationship. Remember from the discussions in Chapter 2 that when a use case includes another use case, the use case being included needs to run as a precondition. For example, the Login use case of the OSO application needs to be included in the View Supply Catalog use case. The reason you make Login a separate use case is that the Login use case can be reused by one or more other use cases. In the OSO application, the Login use case will also be included with the Track Spending use case. Figure 4-2 depicts this includes relationship.




Figure 4-2. Including the Login use case

■ **Note** In some modeling tools, the includes relationship may be indicated in the use case diagram by the uses keyword.

The extends relationship exists between two use cases when, depending on a condition, a use case will extend the behavior of the initial use case. In the OSO application, when a manager is making a purchase request, she can indicate that she will be requesting a purchase for the department. In this case, the Department Purchase Request use case becomes an extension of the Purchase Request use case. Figure 4-3 diagrams this extension.




Figure 4-3. Extending the Purchase Request use case

After analyzing the system requirements and use cases, you can make the system development more manageable by breaking up the application and developing it in phases. For example, you can develop the Purchase Request portion of the application first. Next, you can develop the Request Review portion, and then the Item Check In portion. The rest of this chapter focuses on the Purchase Request portion of the application. Employees and department managers will use this part of the application to make purchase requests. Figure 4-4 shows the use case diagram for this phase.




Figure 4-4. Purchase Request use case diagram

Developing the Class Model

Developing the class model involves several tasks. You begin by identifying the classes, and then you add attributes, associations, and behaviors.

Identifying the Classes

After you have identified the various use cases, you can start identifying the classes the system needs to include to carry out the functionality described in the use cases. To identify the classes, you drill down into each use case and define a series of steps needed to carry it out. It is also helpful to identify the noun phrases in the use case descriptions. The noun phrases are often good indicators of the classes that will be needed.

For example, the following steps describe the View Supply Catalog use case:

- User has logged in and been assigned a user status level. (This is the precondition.)
- Users are presented with a catalog table that contains a list of supplies. The table contains information such as the supply name, category, description, and cost.
- Users can filter supplies by category.
- Users are given the choice of logging out or making a purchase request. (This is the post condition.)

From this description, you can identify a class that will be responsible for retrieving product information from the database and filtering the products being displayed. The name of this class will be the ProductCatalog class.

Examining the noun phrases in the use case descriptions dealing with making purchase requests reveals the candidate classes for the OSO application, as listed in Table 4-2.

Table 4-2. Candidate Classes Used to Make Purchase Requests


Use Case	Candidate Classes
Login	User, username, password, success, failure
View Supply Catalog	User, catalog table, list of supplies, information, supply name, category, description, cost
Purchase Request	Purchaser, items, cart, number, item requested, cost, total cost
Department Purchase Request	Department manager, items, cart, number, item requested, cost, total cost, department purchase request

Now that you have identified the candidate classes, you need to eliminate the classes that indicate redundancy. For example, a reference to items and line items would represent the same abstraction. You can also eliminate classes that represent attributes rather than objects. Username, password, and cost are examples of noun phrases that represent attributes. Some classes are vague or generalizations of other classes. *User* is actually a generalization of *purchaser* and *manager*. Classes may also actually refer to the same object abstraction but indicate a different state of the object. For example, a supply request and an order represent the same abstraction before and after approval. You should also filter out classes that represent implementation constructs such as list and table. For example, a cart is really a collection of order items that makes up an order.

Using these elimination criteria, you can whittle down the class list to the following candidate classes:

- Employee
- DepartmentManager
- Order
- OrderItem
- ProductCatalog
- Product

You can now start formulating the class diagram for the Purchase Request portion of the OSO application. Figure 4-5 shows the preliminary class diagram for the OSO application.

**Figure 4-5.** Preliminary OSO class diagram

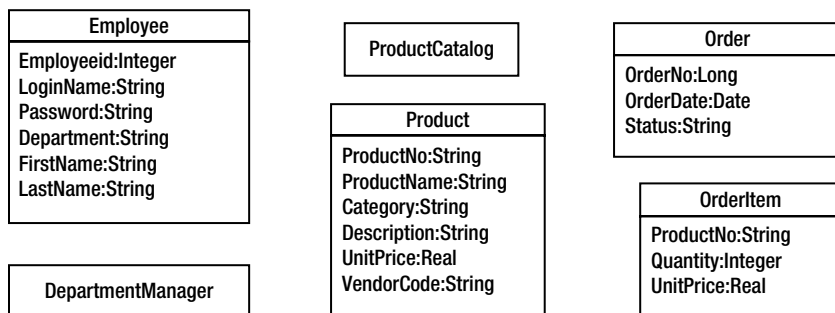
Adding Attributes to the Classes

The next stage in the development of the class model is to identify the level of abstraction the classes must implement. You determine what state information is relevant to the OSO application. This required state information will be implemented through the attributes of the class. Analyzing the system requirements for the *Employee* class reveals the need for a login name, password, and department. You also need an identifier such as an employee ID to uniquely identify various employees. An interview with managers revealed the need to include the first and last names of the employees so that they can track spending by name. Table 4-3 summarizes the attributes that will be included in the OSO classes.

Table 4-3. *OSO Class Attributes*

Class	Attribute	Type
Employee	EmployeeID	Integer
	LoginName	String
	Password	String
	Department	String
	FirstName	String
	LastName	String
DepartmentManager	EmployeeID	Integer
	LoginName	String
	Password	String
	Department	String
	FirstName	String
	LastName	String
Order	OrderNumber	Long
	OrderDate	Date
	Status	String
OrderItem	ProductNumber	String
	Quantity	Integer
	UnitPrice	Decimal
Product	ProductNumber	String
	ProductName	String
	Description	String
	UnitPrice	Decimal
	Category	String
	VendorCode	String
ProductCatalog	None	

Figure 4-6 shows the OSO class diagram with the class attributes identified so far.

**Figure 4-6.** *The Purchase Request component class diagram with attributes added*

Identifying Class Associations

The next stage in the development process is to model the class associations that will exist in the OSO application. If you study the use cases and SRS, you can gain an understanding of what types of associations you need to incorporate into the class structural design.

■ **Note** You may find that you need to further refine the SRS to expose the class associations.

For example, an employee will be associated with an order. By examining the multiplicity of the association, you discover that an employee can have multiple orders, but an order can be associated with only one employee. Figure 4-7 models this association.




Figure 4-7. Depicting the association between the *Employee* class and the *Order* class

As you start to identify the class attributes, you will notice that the *Employee* class and the *DepartmentManager* class have many of the same attributes. This makes sense, because a manager is also an employee. For the purpose of this application, a manager represents an employee with specialized behavior. This specialization is represented by an inheritance relationship, as shown in Figure 4-8.




Figure 4-8. The *DepartmentManager* class inheriting from the *Employee* class

The following statements sum up the associations in the OSO class structure:

- An *Order* is a collection of *OrderItem* objects.
- An *Employee* can have multiple *Order* objects.
- An *Order* is associated with one *Employee*.
- The *ProductCatalog* is associated with multiple *Product* objects.
- A *Product* is associated with the *ProductCatalog*.
- An *OrderItem* is associated with one *Product*.
- A *Product* may be associated with multiple *OrderItem* objects.
- A *DepartmentManager* is an *Employee* with specialized behavior.

Figure 4-9 shows these various associations (excluding the class attributes for clarity).

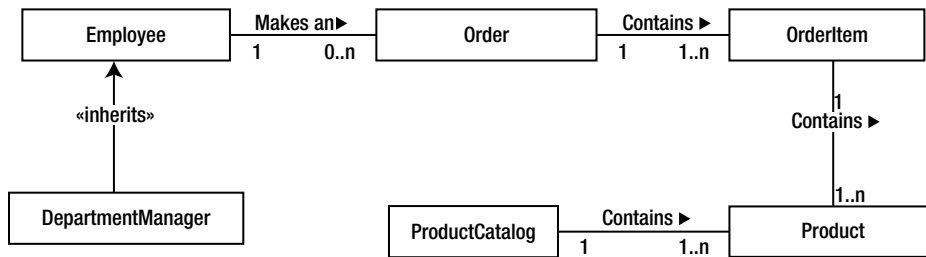


Figure 4-9. The Purchase Request component class diagram with associations added

Modeling the Class Behaviors

Now that you have sketched out the preliminary structure of the classes, you are ready to model how these classes will interact and collaborate. The first step in this process is to drill down into the use case descriptions and create a more detailed scenario of how the use case will be carried out. The following scenario describes one possible sequence for carrying out the Login use case.

1. The user is presented with a login dialog box.
2. The user enters a login name and a password.
3. The user submits the information.
4. The name and password are checked and verified.
5. The user is presented with a supply request screen.

Although this scenario depicts the most common processing involved with the Login use case, you may need other scenarios to describe anticipated alternate outcomes. The following scenario describes an alternate processing of the Login use case:

1. The user is presented with a login dialog box.
2. The user enters a login name and a password.
3. The user submits the information.
4. The name and password are checked but cannot be verified.
5. The user is informed of the incorrect login information.
6. The user is presented with a login dialog box again.
7. The user either tries again or cancels the login request.

At this point, it may help to create a visual representation of the scenarios outlined for the use case. Remember from Chapter 3 that activity diagrams are often used to visualize use case processing. Figure 4-10 shows an activity diagram constructed for the Login use case scenarios.

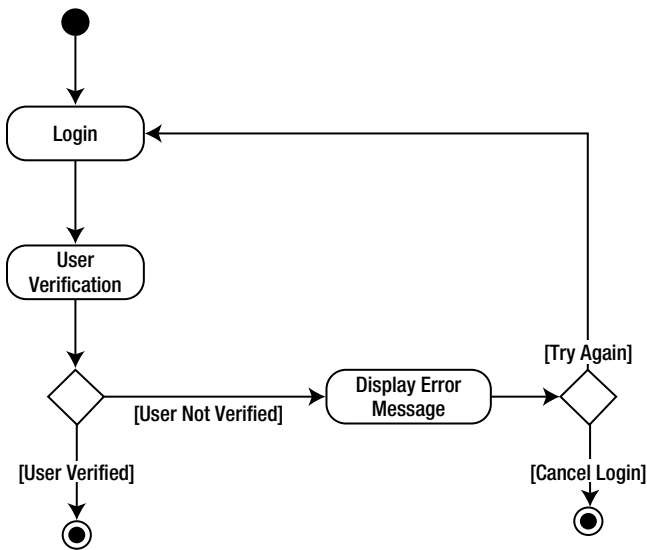


Figure 4-10. An activity diagram depicting the Login use case scenarios

After analyzing the process involved in the use case scenarios, you can now turn your attention to assigning the necessary behaviors to the classes of the system. To help identify the class behaviors and interactions that need to occur, you construct a sequence diagram, as discussed in Chapter 3.

Figure 4-11 shows a sequence diagram for the Login use case scenarios. The Purchaser (UI) class calls the Login method that has been assigned to the Employee class. The message returns information that will indicate whether the login has been verified.




Figure 4-11. A sequence diagram depicting the Login use case scenarios

Next, let's analyze the View Supply Catalog use case. The following scenario describes the use case:

1. User has logged in and has been verified.
2. User views a catalog table that contains product information, including the supply name, category, description, and price.
3. User chooses to filter the table by category, selects a category, and refreshes the table.

From this scenario, you can see that you need a method of the ProductCatalog class that will return a listing of product categories. The Purchaser class will invoke this method. Another method the ProductCatalog class needs is one that will return a product list filtered by category. The sequence diagram in Figure 4-12 shows the interaction that occurs between the Purchaser (UI) class and the ProductCatalog class.

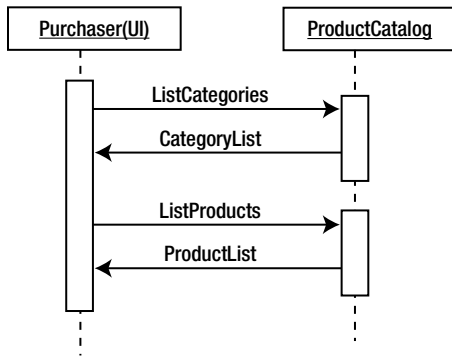


Figure 4-12. A sequence diagram depicting the View Supply Catalog scenario

The following scenario was developed for the Purchase Request use case:

1. A purchaser has logged in and has been verified as an employee.
2. The purchaser selects items from the product catalog and adds them to the order request (shopping cart), indicating the number of each item requested.
3. After completing the item selections for the order, the purchaser submits the order.
4. Order request information is updated, and an order ID is generated and returned to the purchaser.

From the scenario, you can identify an `AddItem` method of the `Order` class that needs to be created. This method will accept a product ID and a quantity, and then return the subtotal of the product ordered. The `Order` class will need to call a method of the `OrderItem` class, which will create an instance of an order item. You also need a `SubmitOrder` method of the `Order` class that will submit the request and the return order ID of the generated order. Figure 4-13 shows the associated sequence diagram for this scenario.

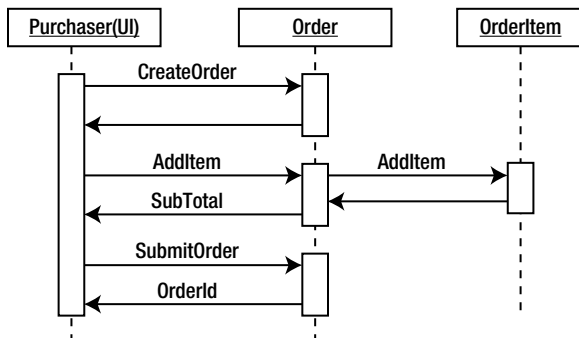


Figure 4-13. A sequence diagram depicting the Purchase Request scenario

Some other scenarios that need to be included are deleting an item from the shopping cart, changing the quantity of an item in the cart, and canceling the order process. You will also need to include similar scenarios and create similar methods for the Department Purchase Request use case. After analyzing the scenarios and interactions that need to take place, you can develop a class diagram for the Purchase Request portion of the application, as shown in Figure 4-14.




Figure 4-14. Purchase Request class diagram

Developing the User Interface Model Design

At this point in the application design process, you don't want to commit to a particular GUI implementation (in other words, a technology-specific one). It is helpful, however, to model some of the common elements and functionality required of a GUI for the application. This will help you create a prototype user interface that you can use to verify the business logic design that has been developed. The users will be able to interact with the prototype and provide feedback and verification of the logical design.

The first prototype screen that you need to implement is the one for logging in. You can construct an activity diagram to help define the activities the user needs to perform when logging in to the system, as shown in Figure 4-15.




Figure 4-15. An activity diagram depicting user login activities

Analyzing the activity diagram reveals that you can implement the login screen as a fairly generic interface. This screen should allow the user to enter a username and password. It should include a way to indicate that the user is logging in as either an employee or a manager. A manager must log in as an employee to make a purchase request and as a manager to approve requests. The final requirement is to include a way for the user to abort the login process. Figure 4-16 shows a prototype of the login screen.

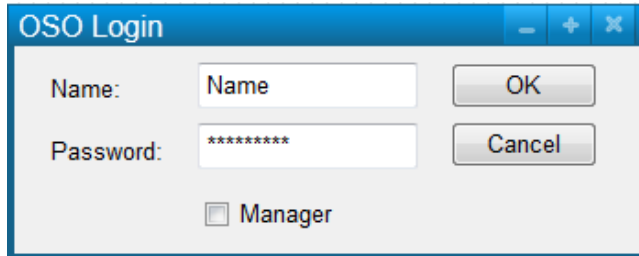


Figure 4-16. Login screen prototype

The next screen you need to consider is the product catalog screen. Figure 4-17 depicts the activity diagram for viewing and filtering the products.

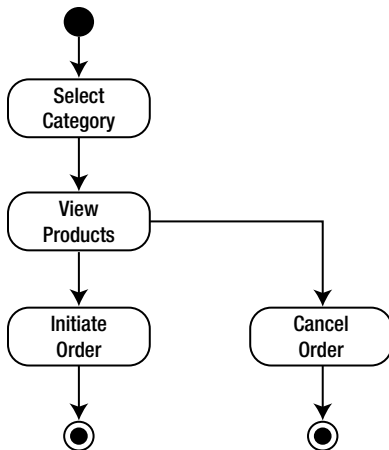


Figure 4-17. An activity diagram depicting activities for viewing products

The activity diagram reveals that the screen needs to show a table or list of products and product information. Users must be able to filter the products by category, which can be initiated by selecting a category from a category list. Users also need to be able to initiate an order request or exit the application. Figure 4-18 shows a prototype screen that can be used to view the products. Clicking the Add to Order button adds the product to the order and launches the order details screen, where the user can adjust the quantity of the order.

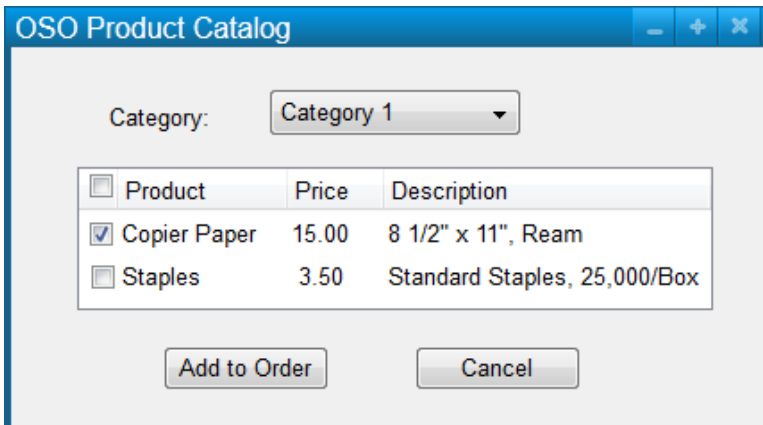


Figure 4-18. View products screen prototype

The final screen that needs to be prototyped for this part of the application is the shopping cart interface. This will facilitate the adding and removing of items from an order request. It also needs to allow the user to submit the order or abort an order request. Figure 4-19 shows a prototype of the order request screen. Clicking on the Add button will launch the previous product catalog screen, which allows the user to add additional line items. The Remove button will remove line items checked. Users can change the quantity of the line item by clicking the up and down arrows.

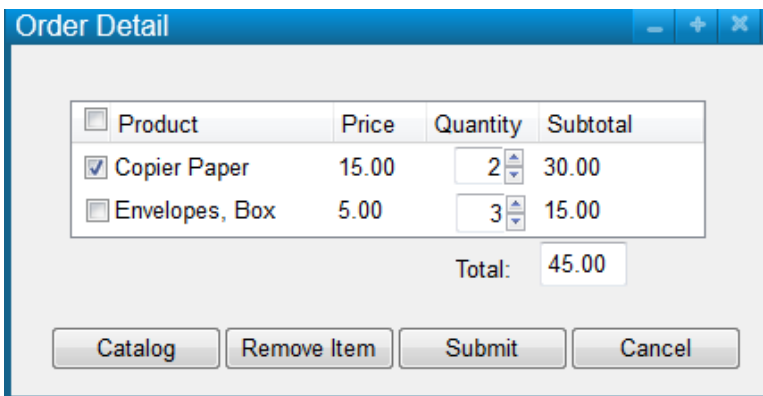


Figure 4-19. Order detail screen prototype

That completes the preliminary design for this phase of the OSO application. You applied what you learned in Chapters 2 and 3 to model the design. Next, let's review some common mistakes to avoid during this process.

Avoiding Some Common OOP Design Pitfalls

When you start to model your own OOP designs, you want to be sure to follow best practices. The following are some of the common traps that you should avoid

- *Not involving the users in the process:* It is worth emphasizing that users are the consumers of your product. They are the ones who define the business processes and functional requirements of the system.
- *Trying to model the whole solution at one time:* When developing complex systems, break up the system design and development into manageable components. Plan to produce the software in phases. This will provide for faster modeling, developing, testing, and release cycles.
- *Striving to create a perfect model:* No model will be perfect from the start. Successful modelers understand that the modeling process is iterative, and models are continuously updated and revised throughout the application development cycle.
- *Thinking there is only one true modeling methodology:* Just as there are many different equally viable OOP languages, there are many equally valid modeling methodologies for developing software. Choose the one that works best for you and the project at hand.
- *Reinventing the wheel:* Look for patterns and reusability. If you analyze many of the business processes that applications attempt to solve, a consistent set of modeling patterns emerge. Create a repository where you can leverage these existing patterns from project to project and from programmer to programmer.
- *Letting the data model drive the business logic model:* It is generally a bad idea to develop the data model (database structure) first and then build the business logic design on top of it. The solution designer should first ask what business problem needs to be solved, and then build a data model to solve the problem.
- *Confusing the problem domain model with the implementation model:* You should develop two distinct but complementary models when designing applications. A domain model design describes the scope of the project and the processing involved in implementing the business solutions. This includes what objects will be involved, their properties and behaviors, and how they interact and relate to each other. The domain model should be implementation-agnostic. You should be able to use the same domain model as a basis for several different architecturally specific implementations. In other words, you should be able to take the same domain model and implement it using a Visual Basic rich-client, two-tier architecture or a C# (or Java, for that matter) n-tier distributed web application.

Summary

Now that you have analyzed the domain model of an OOP application, you are ready to transform the design into an actual implementation. The next part of this book will introduce you to the C# language. You will look at the .NET Framework and see how C# applications are built on top of the framework. You will be introduced to working in the Visual Studio IDE and become familiar with the syntax of the C# language. The next section will also demonstrate the process of implementing OOP constructs such as class structures, object instantiation, inheritance, and polymorphism in C#. With your newfound skills, you will revisit the case study introduced in this chapter and apply these skills to transform the application design into a functioning application.



Introducing the .NET Framework and Visual Studio

Business application programming has evolved from a two-tier, tightly coupled model into a multitiered, loosely coupled model, often involving data transfer over the Internet or a corporate intranet. In an effort to allow programmers to be more productive and deal with the complexities of this type of model, Microsoft developed the .NET Framework. To effectively program in C#, you need to understand the underlying framework upon which it is built.

After reading this chapter, you should be familiar with the following:

- the .NET Framework
- the features of the Common Language Runtime (CLR)
- how the just-in-time (JIT) compiler works
- the .NET Framework base class library
- namespaces and assemblies
- the features of the Visual Studio integrated development environment

Introducing the .NET Framework

The .NET Framework is a collection of fundamental classes designed to provide the common services needed to run applications. Let's look at the goals of the .NET Framework and then review its components.

Goals of the .NET Framework

Microsoft designed the .NET Framework with certain goals in mind. The following sections examine these goals and how the .NET Framework achieves them.

Support of Industry Standards

Microsoft wanted the .NET Framework to be based on industry standards and practices. As a result, the framework relies heavily on industry standards such as the Extensible Markup Language (XML), HTML 5, and OData. Microsoft

has also submitted a Common Language Infrastructure (CLI) Working Document to the European Computer Manufacturers Association (ECMA), which oversees many of the common standards in the computer industry.

The CLI is a set of specifications needed to create compilers that conform to the .NET Framework. Third-party vendors can use these specifications to create .NET-compliant language compilers; for example, Interactive Software Engineering (ISE) has created a .NET compiler for Eiffel. Third-party vendors can also create a CLR that will allow .NET-compliant languages to run on different platforms. One example, Mono is an open source, cross platform implementation of the CLR that gives C# applications the ability to run on the Linux platform.

Extensibility

To create a highly productive environment in which to program, Microsoft realized the .NET Framework had to be extensible. As a result, Microsoft exposed the framework class hierarchy to developers. Through inheritance and interfaces, you can easily access and extend the functionality of these classes. For example, you could create a button control class that not only inherits its base functionality from the button class exposed by the .NET Framework, but also extends the base functionality in the unique way required by your application.

Microsoft has also made it much easier to work with the underlying operating system. By repackaging and implementing the Windows operating system application programming interface (API) functions in a class-based hierarchy, Microsoft has made it more intuitive and easier for OOP programmers to work with the functionality exposed by the underlying operating system.

Unified Programming Models

Another important goal Microsoft incorporated into the .NET Framework was cross-language independence and integration. To achieve this goal, all languages that support the Common Language Specification (CLS) compile into the same intermediate language, support the same set of basic data types, and expose the same set of code-accessibility methods. As a result, not only can classes developed in the different CLS-compliant languages communicate seamlessly with one another, but you can also implement OOP constructs across languages. For example, you could develop a class written in C# that inherits from a class written using Visual Basic (VB). Microsoft has developed several languages that support the .NET Framework. Along with C#, the languages are VB.NET, managed C++, JScript, and F#. In addition to these languages, many third-party vendors have developed versions of other popular languages designed to run under the .NET Framework, such as Pascal and Python.

Easier Deployment

Microsoft needed a way to simplify application deployment. Before the development of the .NET Framework, when components were deployed, component information had to be recorded in the system registry. Many of these components, especially system components, were used by several different client applications. When a client application made a call to the component, the registry was searched to determine the metadata needed to work with the component. If a newer version of the component was deployed, it replaced the registry information of the old component. Often, the new components were incompatible with the old version and caused existing clients to fail. You have probably experienced this problem after installing a service pack that ended up causing more problems than it fixed!

The .NET Framework combats this problem by storing the metadata for working with the component in a file called a manifest, which is packaged in the assembly containing the component code. An assembly is a package containing the code, resources, and metadata needed to run an application. By default, an assembly is marked as private and placed in the same directory as the client assembly. This ensures that the component assembly is not inadvertently replaced or modified and also allows for a simpler deployment because there is no need to work with the registry. If a component needs to be shared, its assembly is deployed to a special directory referred to as the Global Assembly Cache (GAC). The manifest of the assembly contains versioning information, so newer versions of the

component can be deployed side by side with the older versions in the GAC. By default, client assemblies continue to request and use the versions of the components they were intended to use. Older client assemblies will no longer fail when newer versions of the component are installed.

Improved Memory Management

A common problem of programs developed for the Windows platform has been memory management. Often, these programs have caused memory leaks. A memory leak occurs when a program allocates memory from the operating system but fails to release the memory after it is finished working with the memory. This memory is no longer available for other applications or the operating system, causing the computer to run slowly or even to stop responding. This problem is compounded when the program is intended to run for a long time, such as a service that runs in the background. To combat this problem, the .NET Framework uses nondeterministic finalization. Instead of relying on the applications to deallocate the unused memory, the framework uses a garbage collection object. The garbage collector periodically scans for unused memory blocks and returns them to the operating system.

Improved Security Model

Implementing security in today's highly distributed, Internet-based applications is an extremely important issue. In the past, security has focused on the user of the application. Security identities were checked when users logged in to an application, and their identities were passed along as the application made calls to remote servers and databases. This type of security model has proven to be inefficient and complicated to implement for today's enterprise-level, loosely coupled systems. In an effort to make security easier to implement and more robust, the .NET Framework uses the concept of code identity and code access.

When an assembly is created, it is given a unique identity. When a server assembly is created, you can grant access permissions and rights. When a client assembly calls a server assembly, the runtime will check the permissions and rights of the client, and then grant or deny access to the server code accordingly. Because each assembly has an identity, you can also restrict access to the assembly through the operating system. If a user downloads a component from the Web, for example, you can restrict the component's ability to read and write files on the user's system.

Components of the .NET Framework

Now that you have seen some of the major goals of the .NET Framework, let's take a look at the components it contains.

Common Language Runtime

The fundamental component of the .NET Framework is the CLR. The CLR manages the code being executed and provides for a layer of abstraction between the code and the operating system. Built into the CLR are mechanisms for the following:

- loading code into memory and preparing it for execution
- converting the code from the intermediate language to native code
- managing code execution
- managing code and user-level security
- automating deallocation and release of memory
- debugging and tracing code execution
- providing structured exception handling

Framework Base Class Library

Built on top of the CLR is the .NET Framework base class library. Included in this class library are reference types and value types that encapsulate access to the system functionality. *Types* are data structures. A reference type is a complex type—for example, classes and interfaces. A value type is simple type—for example, integer or Boolean. Programmers use these base classes and interfaces as the foundation on which they build applications, components, and controls. The base class library includes types that encapsulate data structures, perform basic input/output operations, invoke security management, manage network communication, and perform many other functions.

Data Classes

Built on top of the base classes are classes that support data management. This set of classes is commonly referred to as ADO.NET. Using the ADO.NET object model, programmers can access and manage data stored in a variety of data storage structures through managed providers. Microsoft has written and tuned the ADO.NET classes and object model to work efficiently in a loosely coupled, disconnected, multitiered environment. ADO.NET not only exposes the data from the database, but also exposes the metadata associated with the data. Data is exposed as a sort of mini-relational database. This means that you can get the data and work with it while disconnected from the data source, and later synchronize the data with the data source.

Microsoft has provided support for several data providers. Data stored in Microsoft SQL Server can be accessed through the native SQL data provider. OLEDB and Open Database Connectivity-managed (ODBC) providers are two generic providers for systems currently exposed through the OLEDB or ODBC standard APIs. Because these managed data providers do not interface directly with the database engine, but rather talk to the unmanaged provider, which then talks to the database engine, using nonnative data providers is less efficient and less robust than using a native provider. Because of the extensibility of the .NET Framework and Microsoft's commitment to open-based standards, many data storage vendors now supply native data providers for their systems.

Built on top of the ADO.NET provider model is the ADO.NET Entity Framework. The Entity Framework bridges the gap between the relation data structure of the database and the object-oriented structure of the programming language. It provides an Object/Relational Mapping (ORM) framework that eliminates the need for programmers to write most of the plumbing code for data access. The framework provides services such as change tracking, identity resolution, and query translation. Programmers retrieve data using Language Integrated Query (LINQ) and manipulate data as strongly typed objects. Chapter 10 takes a detailed look at ADO.NET and data access.

Windows Applications

Prior to the .NET Framework, developing Windows Graphical User Interfaces (GUIs) was dramatically different depending on whether you were developing using C++ or Visual Basic. Although developing GUIs in VB was easy and could be accomplished very quickly, VB developers were isolated from and not fully exposed to the underlying features of the Windows API. On the other hand, although exposed to the full features of the Windows API, developing GUIs in C++ was very tedious and time-consuming. With the .NET Framework, Microsoft has incorporated a set of base classes exposing advanced Windows GUI functionality equally among the .NET-compliant languages. This has allowed Windows GUI development to become consistent across the various .NET-enabled programming languages, combining the ease of development with the full features of the API.

Along with traditional Windows forms and controls, which have been around since the beginning of the Windows Operating System, the .NET Framework includes a set of classes collectively referred to as the Windows Presentation Foundation (WPF). WPF integrates a rendering engine that is built to take advantage of modern graphics hardware and feature rich operating systems like Windows 7. It also includes application development features such as controls, data binding, layout, graphics, and animation. With the WPF set of classes, programmers can create applications that provide a very compelling user experience. You will look more closely at building WPF-based applications in Chapter 11.

Web Applications

The .NET Framework exposes a base set of classes that can be used on a web server to create user interfaces and services exposed to web-enabled clients. These classes are collectively referred to as ASP.NET. Using ASP.NET, you can develop one user interface that can dynamically respond to the type of client device making the request. At runtime, the .NET Framework takes care of discovering the type of client making the request (browser type and version) and exposing an appropriate interface. The GUIs for web applications running on a Windows client have become more robust because the .NET Framework exposes much of the API functionality that previously had been exposed only to traditional Windows Forms-based C++ and VB applications. Another improvement in web application development using the .NET Framework is that server-side code can be written in any .NET-compliant language. Prior to .NET, server-side code had to be written in a scripting language such as VBScript or JScript.

Visual Studio 2012 and the .NET framework 4.5 support several different models for creating web apps. You can create web apps based on an ASP.NET Web Forms project, an ASP.NET MVC project, or a Silverlight Application project. Chapter 12 covers developing web applications.

Windows Store Applications

A new type of Windows application is available for the Windows 8 operating system: the Windows Store app. Windows Store apps are intended for devices such as tablets and phones that take advantage of touch screens for user input, and continuous Internet connectivity. There are two options for building Windows Store apps. You can use one of the .NET languages (C#/VB/C++) in combination with XAML or you can use JavaScript in combination with HTML5. Windows Store apps are built using the WinRT API which is similar to programming against the .NET Framework. Microsoft has even added a .NET Framework API as a wrapper around the WinRT API. This means that you can use the .NET Framework to create great Windows Store apps using the programming language and design experience you are already familiar with. In fact, as you will see in Chapter 13, creating Windows Store applications using C# and XAML is very similar to building WPF applications.

Application Services

Included in the .NET Framework are base class and interface support for exposing services that can be consumed by other applications. Previous to the .NET Framework, applications developed in C++ and VB used Distributed Component Object Model (DCOM) technology. DCOM is technology for communication among software components distributed across networked computers. Because DCOM was based on binary standards, application-to-application communication through firewalls and across the Internet was not easy to implement. The proprietary nature of the DCOM also limited the types of clients that could effectively use and interact with applications exposing services through COM.

Microsoft has addressed these limitations by exposing services through Internet standards. Included in the .NET Framework is a set of classes collectively referred to as the Windows Communication Foundation (WCF). Using WCF, you can send data as messages from one application to another. The message transport and content can be easily changed depending on the consumer and environment. For example, if the service is exposed over the Web, a text-based message over HTTP can be used. On the other hand, if the client is on the same corporate network, a binary message over TCP can be used.

With .NET 4.5, Microsoft has also introduced a new framework for creating HTTP services called the ASP.NET Web API. Using this API you can build services that can reach a broad range of clients, such as web browsers and mobile devices. While you could create these services using WCF, the new Web API makes it easier to implement and includes features such as content negotiation, query composition, and flexible hosting. In Chapter 14, you will create and consume application services using WCF and the ASP.NET Web API.

Working with the .NET Framework

To work with the .NET Framework, you should understand how it is structured and how managed code is compiled and executed. .NET applications are organized and packaged into assemblies. All code executed by the .NET runtime must be contained in an assembly.

Understanding Assemblies and Manifests

The assembly contains the code, resources, and a manifest (metadata about the assembly) needed to run the application. Assemblies can be organized into a single file where all this information is incorporated into a single dynamic link library (DLL) file or executable (EXE) file, or multiple files where the information is incorporated into separate DLL files, graphics files, and a manifest file. One of the main functions of an assembly is to form a boundary for types, references, and security. Another important function of the assembly is to form a unit for deployment.

One of the most crucial portions of an assembly is the manifest; in fact, every assembly must contain a manifest. The purpose of the manifest is to describe the assembly. It contains such things as the identity of the assembly, a description of the classes and other data types the assembly exposes to clients, any other assemblies this assembly needs to reference, and security details needed to run the assembly.

By default, when an assembly is created, it is marked as private. A copy of the assembly must be placed in the same directory or a bin subdirectory of any client assembly that uses it. If the assembly must be shared among multiple client assemblies, it is placed in the Global Assembly Cache (GAC), a special Windows folder. To convert a private assembly into a shared assembly, you must run a utility program to create encryption keys, and you must sign the assembly with the keys. After signing the assembly, you must use another utility to add the shared assembly into the GAC. By mandating such stringent requirements for creating and exposing shared assemblies, Microsoft is trying to ensure that malicious tampering of shared assemblies will not occur. This also ensures a new version of the same assembly can exist without breaking current applications using the old assembly.

Referencing Assemblies and Namespaces

To make the .NET Framework more manageable, Microsoft has given it a hierarchical structure. This hierarchical structure is organized into what are referred to as namespaces. By organizing the framework into namespaces, the chances of naming collisions are greatly reduced. Organizing related functionality of the framework into namespaces also greatly enhances its usability for developers. For example, if you want to build a window's GUI, it is a pretty good bet the functionality you need exists in the `System.Windows` namespace.

All of the .NET Framework classes reside in the `System` namespace. The `System` namespace is further subdivided by functionality. The functionality required to work with a database is contained in the `System.Data` namespace. Some namespaces run several levels deep; for example, the functionality used to connect to a SQL Server database is contained in the `System.Data.SqlClient` namespace.

An assembly may be organized into a single namespace or multiple namespaces. Several assemblies may also be organized into the same namespace.

To gain access to the classes in the .NET Framework, you need to reference the assembly that contains the namespace in your code. Then you can access classes in the assembly by providing their fully qualified names. For example, if you want to add a text box to a form, you create an instance of the `System.Windows.Controls.TextBox` class, like so:

```
private System.Windows.Controls.TextBox newTextBox;
```

Fortunately, in C#, you can use the `using` statement at the top of the code file so that you do not need to continually reference the fully qualified name in the code:

```
using System.Windows.Controls;
private TextBox newTextBox;
```

Compiling and Executing Managed Code

When .NET code is compiled, it is converted into a .NET portable executable (PE) file. The compiler translates the source code into Microsoft intermediate language (MSIL) format. MSIL is CPU-independent code, which means it needs to be further converted into CPU-specific native code before executing.

Along with the MSIL code, the PE file includes the metadata information contained within the manifest. The incorporation of the metadata in the PE file makes the code self-describing. There is no need for additional type library or Interface Definition Language (IDL) files.

Because the source code for the various .NET-compliant languages is compiled into the same MSIL and metadata format based on a common type system, the .NET platform supports language integration. This is a step beyond Microsoft's COM components, where, for example, client code written in VB could instantiate and use the methods of a component written in C++. With .NET language integration, you could write a .NET class in VB that inherits from a class written in C#, and then overrides some of its methods.

Before the MSIL code in the PE file is executed, a .NET Framework just-in-time (JIT) compiler converts it into CPU-specific native code. To improve efficiency, the JIT compiler does not convert all the MSIL code into native code at the same time. MSIL code is converted on an as-needed basis. When a method is executed, the compiler checks to see if the code has already been converted and placed in cache. If it has, the compiled version is used; otherwise, the MSIL code is converted and stored in the cache for future calls.

Because JIT compilers are written to target different CPUs and operating systems, developers are freed from needing to rewrite their applications to target various platforms. It is conceivable that the programs you write for a Windows server platform will also run on a UNIX server. All that is needed is a JIT compiler for the UNIX architecture.

Using the Visual Studio Integrated Development Environment

You can write C# code using a simple text editor and compile it with a command-line compiler. You will find, however, that programming enterprise-level applications using a text editor can be frustrating and inefficient. Most programmers who code for a living find an integrated development environment (IDE) invaluable in terms of ease of use and increased productivity. Microsoft has developed an exceptional IDE in Visual Studio (VS). Integrated into VS are many features that make programming for the .NET Framework more intuitive, easier, and more productive. Some of Visual Studio's useful features are:

- editor features such as automatic syntax checking, autocomplete, and color highlighting
- one IDE for all .NET languages
- extensive debugging support, including the ability to set breakpoints, step through code, and view and modify variables
- integrated help documentation
- drag-and-drop GUI development
- XML and HTML editing
- automated deployment tools that integrate with Windows Installer
- the ability to view and manage servers from within the IDE
- a fully customizable and extensible interface

The following activities will introduce you to some of the many features available in the VS IDE. As you work through these steps, don't worry about the coding details. Just concentrate on getting used to working within the VS IDE. You'll learn more about the code in upcoming chapters.

■ **Note** If you do not have Visual Studio 2012 installed, refer to Appendix C for installation instructions.

ACTIVITY 5-1. TOURING VISUAL STUDIO

In this activity, you will become familiar with the following:

- customizing the IDE
- creating a .NET project and setting project properties
- using the various editor windows in the VS IDE
- using the auto-syntax check and autocomplete features of the VS IDE
- compiling assemblies with the VS IDE

Customizing the IDE

To customize the IDE, follow these steps:

1. Start Visual Studio 2012.

■ **Note** If this is the first time you have started VS, you will be asked to choose a default development setting. Choose the Visual C# Development Settings.

2. You will be presented with the Start Page. The Start Page contains several panes, including one that has links to useful documentation posted on the MSDN (Microsoft Developer Network) web site. Clicking one of these links will launch a browser window hosted inside VS, which will open the documentation on the MSDN site. Take some time to investigate the information and the various links exposed to you on the Start Page.
3. Microsoft has taken considerable effort to make VS a customizable design environment. You can customize just about every aspect of the layout, from the various windows and menus down to the color coding used in the code editor. Select Tools ► Options to open the Options dialog box, shown in Figure 5-1, which allows you to customize many aspects of the IDE.

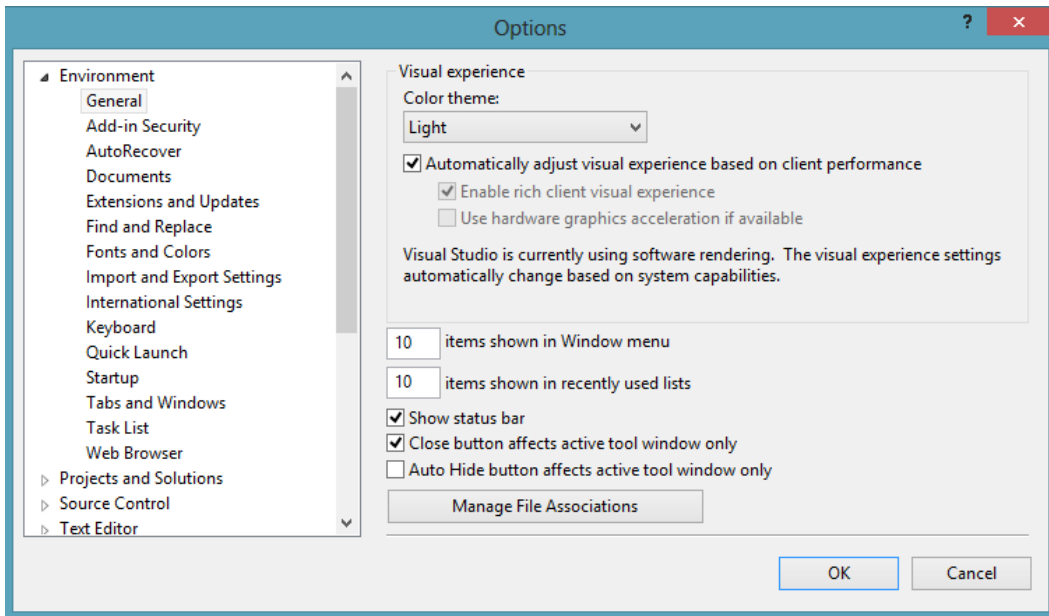


Figure 5-1. VS Options dialog box

4. Click Projects and Solutions in the category list on the left side of the dialog box. You are presented with options to change the default location of projects and what happens when you build and run a project. Select the Show Output Window When Build Starts option.
5. Investigate some of the other customizable options available. Close the Options dialog box when you are finished by clicking the OK button.

Creating a New Project

To create a new project, follow these steps:

1. On the Start Page, click the Create Project link, which launches the New Project dialog box. (You can also choose File ► New ► Project to open this dialog box.)
2. The New Project dialog box allows you to create various projects using built-in templates. There are templates for creating Windows projects, Web projects, WCF projects, as well as many others, depending on what options you chose when installing VS.
3. In the Templates pane, expand the Visual C# node and select the Windows node, as shown in Figure 5-2. Observe the various C# project templates. There are templates for creating various types of Windows applications, including Windows Forms-based applications, class libraries, and console applications.

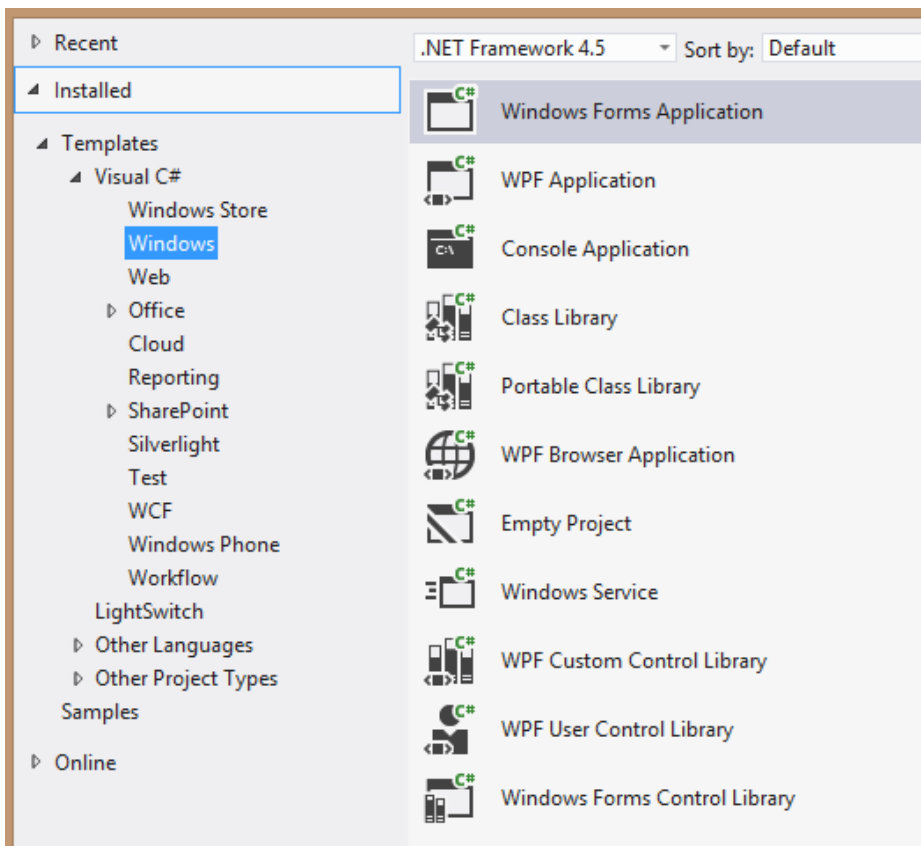


Figure 5-2. VS New Project dialog box

4. Click the Windows Forms Application template. Change the name of the application to DemoChapter5 and click the OK button.

When the project opens, you will be presented with a form designer for a default form (named Form1) that has been added to the project. To the right of this window, you should see the Solution Explorer.

Investigating the Solution Explorer and Class View

The Solution Explorer displays the projects and files that are part of the current solution, as shown in Figure 5-3. By default, when you create a project, a solution is created with the same name as the project. The solution contains some global information, project-linking information, and customization settings, such as a task list and debugging information. A solution may contain more than one related project.




Figure 5-3. Solution Explorer

Under the solution node is the project node. The project node organizes the various files and settings related to a project. The project file organizes this information in an XML document, which contains references to the class files that are part of the project, any external references needed by the project, and compilation options that have been set. Under the Project node is a Properties node, References node, an App.config file, a class file for the Form1 class, and a Program class file.

To practice using the Solution Explorer and some VS features and views, follow these steps:

1. In the Solution Explorer window, right-click the Properties node and select Open. This launches the Project Properties window. Along the left side of the window are several tabs you can use to explore and set various application settings.
2. Select the Application tab, as shown in Figure 5-4. Notice that, by default, the assembly name and default namespace are set to the name of the project.




Figure 5-4. Project Properties Window

3. Explore some of the other tabs in the Project Properties window. Close the window when you are finished by clicking on the x in the tab of the window.
4. In the Solution Explorer window, expand the References node. Under the node are the external assemblies referenced by the application. Notice that several references have been included by default. The default references depend on the type of project. For example, since this is a Windows Application project, a reference to the `System.Windows.Forms` namespace is included by default.
5. The `Form1` class file under the Solution Explorer's project node has a `.cs` extension to indicate it is written in C# code. By default, the name of the file has been set to the same name as the form. Double-click the file in the Solution Explorer window. The form is shown in Design View. Click the View Code button in the toolbar at the top of the Solution Explorer, and the code editor for the `Form1` class will open.
6. Select View ► Class View to launch the Class View window. The top part of the Class View window organizes the project files in terms of the namespace hierarchy. Expanding the `DemoChapter5` root node reveals three sub nodes: a References node, the `DemoChapter5` namespace node, and `DemoChapter5` properties node. A namespace node is designated by the `{ }` symbol to the left of the node name.
7. Listed under the `DemoChapter5` namespace node are the classes that belong to the namespace. Expanding the `Form1` node reveals a Base Types folder. Expanding Base Types shows the classes and interfaces inherited and implemented by the `Form1` class, as shown in Figure 5-5. You can further expand the nodes to show the classes and interfaces inherited and implemented by the `Form` base class.

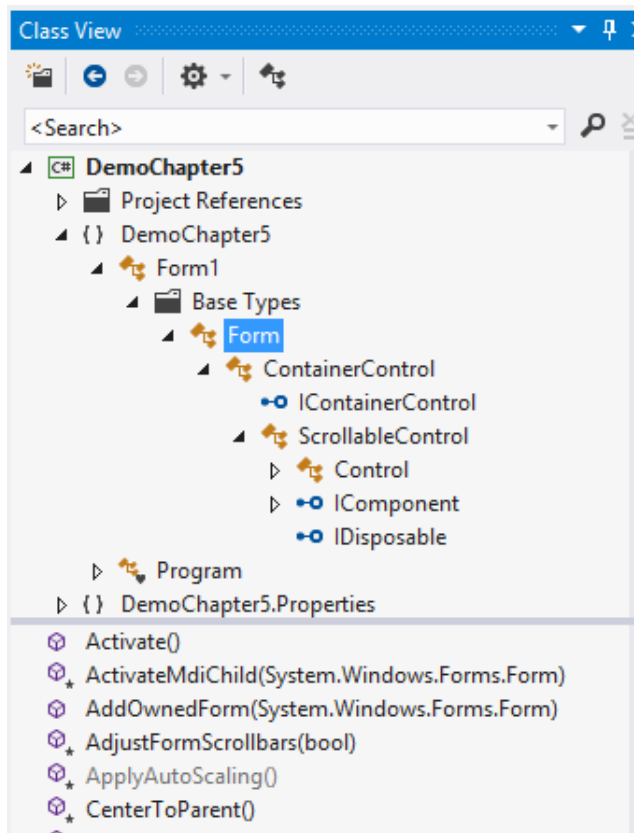


Figure 5-5. Expanded nodes in the Class View

8. The bottom section of the Class View window is a listing of the class's methods, properties, and events. Select the Form node in the top section of the Class View window. Notice the considerable number of methods, properties, and events listed in the bottom section of the window.
9. Right-click the DemoChapter5 project node and select Add ► Class. Name the class DemoClass1 and click the Add button. If the class code is not visible in the code editor, double-click the DemoClass1 node in the Class View window to display it. Wrap the class definition code in a namespace declaration as follows:

```
namespace DemoChapter5
{
    namespace MyDemoNamespace
    {
        class DemoClass1
        {
        }
    }
}
```

10. Notice the updated hierarchy in the Class View. `DemoClass1` now belongs to the `MyDemoNamespace`, which belongs to the `DemoChapter5` namespace. The fully qualified name of `DemoClass1` is now `DemoChapter5.MyDemoNamespace.DemoClass1`.
11. Add the following code to the `DemoClass1` definition. As you add the code, notice the auto-selection drop-down list provided (see Figure 5-6). Pressing the Tab key will select the current item on the list.

```
class DemoClass1: System.Collections.CaseInsensitiveComparer
{
}
```

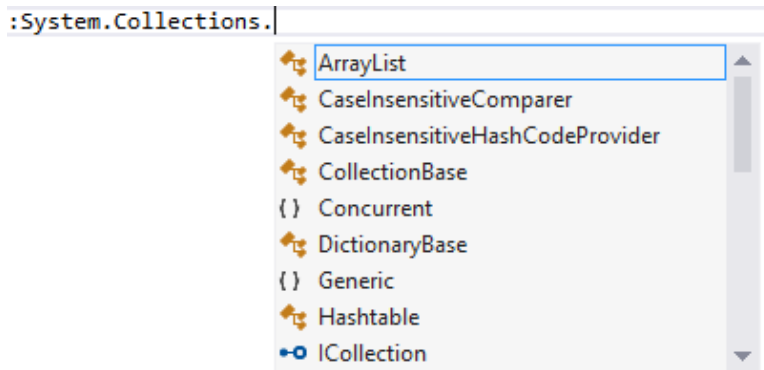


Figure 5-6. Code selection drop-down list

12. Notice the updated hierarchy in the Class View. Expand the Base Types node under the `DemoClass1` node, and you will see the base `CaseInsensitiveComparer` class node. Select this node and you will see the methods and properties of the `CaseInsensitiveComparer` class in the lower section of the Class View window.
13. Right-click the `Compare` method of the `CaseInsensitiveComparer` class node and choose `Browse Definition`. The Object Browser window is opened as a tab in the main window and information about the `Compare` method is displayed. Notice it takes two object arguments, compares them, and returns an integer value based on the result (see Figure 5-7).

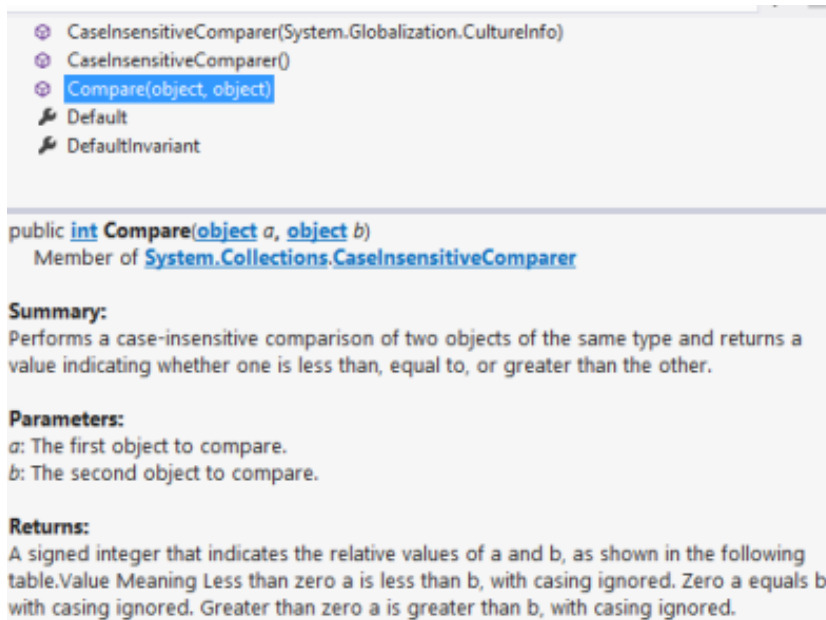


Figure 5-7. Object Browser

14. The Object Browser enables you to explore the object hierarchies and to view information about items and methods within the hierarchy. Take some time to explore the Object Browser. When you are finished, close the Object Browser and close the Class View window.

Exploring the Toolbox and Properties Window

To explore the VS Toolbox and Properties window, follow these steps:

1. In the Solution Explorer window, double-click the `Form1.cs` node. This brings up the `Form1` design tab in the main editing window. Locate the Toolbox tab to the left of the main editing window. Click the tab and the Toolbox window should expand, as shown in Figure 5-8. In the upper-right corner of the Toolbox, you should see the auto-hide icon, which looks like a thumbtack. Click the icon to turn off the auto-hide feature.




Figure 5-8. VS Toolbox

2. Under the All Windows Forms node of the Toolbox are controls that you can drag and drop onto your form to build the GUI. There are also other nodes that contain nongraphical components that help make some common programming tasks easier to create and manage. For example, the Data node contains controls for accessing and managing data stores. Scroll down the Toolbox window and observe the various controls exposed by the designer.
3. Under the All Windows Forms node, select the Label control. Move the cursor over the form; it should change to a crosshairs pointer. Draw a label on the form by clicking, dragging, and then releasing the mouse. In a similar fashion, draw a TextBox control and a Button control on the form. Figure 5-9 shows how the form should look.




Figure 5-9. Sample form layout

4. Turn the auto-hide feature of the Toolbox back on by clicking the Auto Hide (thumbtack) icon in the upper-right corner of the Toolbox window.
5. Locate the Properties window to the right of the main editing window under the Solution Explorer window. (You can also select View ► Properties Window in the menu step to open the Properties window.) The Properties window displays the properties of the currently selected object in the Design View. You can also edit many of the object's properties through this window.
6. In the Form1 design window, click Label1. The Label1 control should be selected in the drop-down list at the top of the Properties window (see Figure 5-10). Locate the Text property and change it to "Enter your password:" (minus the quotes).




Figure 5-10. VS Properties window

■ **Note** You may need to rearrange the controls on the form to see all the text.

7. Set the Password Char property of TextBox1 to *. Change the Text property of Button1 to OK. (Click the control on the form or use the drop-down list at the top of the Properties window to see the control's properties.)
8. Save the project by choosing File ► Save All.

Building and Executing the Assembly

To build and execute the assembly, follow these steps:

1. In the Solution Explorer, double click Form1 to bring up the Design window.
2. In the form designer, double click the Button1 control. The code editor for Form1 will be displayed in the main editing window. A method that handles the button click event is added to the code editor.
3. Add the following code to the method. This code will display the password entered in TextBox1 on the title bar of the form.

```
private void button1_Click(object sender, EventArgs e)
{
    this.Text = "Your password is " + textBox1.Text;
}
```

4. Select Build ► Build Solution. The Output window shows the progress of compiling the assembly (see Figure 5-11). Once the assembly has been compiled, it is ready for execution. (If you can't locate the Output window, select View menu ► Output.)




Figure 5-11. Progress of build displayed in the Output window

5. Select Debug ► Start Debugging. This runs the assembly in debug mode. Once the form loads, enter a password and click the OK button. You should see the message containing the password in the form's title bar. Close the form by clicking the x in the upper right corner.
6. Select File ► Save All, and then exit VS by selecting File ► Exit.

ACTIVITY 5-2. USING THE DEBUGGING FEATURES OF VS

In this activity, you will become familiar with the following:

- setting breakpoints and stepping through the code
- using the various debugging windows in the VS IDE
- locating and fixing build errors using the Error List window

Stepping Through Code

To step through your code, follow these steps:

1. Start VS. Select File ► New ► Project.
2. Under the C# Windows templates, select the Console Application. Rename the project Activity5_2.
3. You will see a Program class file open in the code editor. The class file has a Main method that gets executed first when the application runs. Add the following code to the program class. This code contains a method that loads a list of numbers and displays the contents of the list in the console window.

```
class Program
{
    static List<int> numList = new List<int>();
    static void Main(string[] args)
    {
        LoadList(10);
        foreach (int i in numList)
        {
            Console.WriteLine(i);
        }
        Console.ReadLine();
    }
    static void LoadList(int iMax)
    {
        for (int i = 1; i <= iMax; i++)
        {
            numList.Add(i);
        }
    }
}
```

4. To set a breakpoint (which pauses program execution), place the cursor on the declaration line of the Main method, right-click, and choose Breakpoint ► Insert Breakpoint. A red dot will appear in the left margin to indicate that a breakpoint has been set (see Figure 5-12).

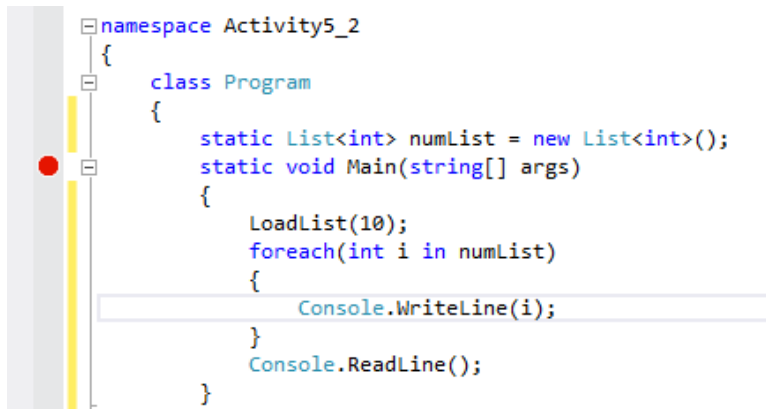


Figure 5-12. Setting a breakpoint in the code editor

5. Select **Debug** ► **Start Debugging**. Program execution will pause at the breakpoint. A yellow arrow indicates the next line of code that will be executed.
6. Select **View** ► **Toolbars** and click the **Debug** toolbar. (A check next to the toolbar name indicates it is visible.) To step through the code one line at a time, select the **Step Into** button on the **Debug** toolbar (see Figure 5-13). (You can also press the F11 key.) Continue stepping through the code until you get to the `LoadList`.




Figure 5-13. Using the **Debug** toolbar

7. Step through the code until the `for` loop has looped a couple of times. At this point, you are probably satisfied that this code is working and you want to step out of this method. On the **Debug** toolbar, click the **Step Out** button. You should return to the `Main` method.
8. Continue stepping through the code until the `for-each` loop has looped a couple of times. At this point, you may want to return to runtime mode. To do this, click the **Continue** button on the **Debug** toolbar. When the **Console** window appears, hit the `enter` key to close the window.
9. Start the application in debug mode again. Step through the code until you get to the method call `LoadList(10);`.
10. On the **Debug** toolbar, choose the **Step Over** button. This will execute the method and reenter break mode after execution returns to the calling code. After stepping over the method, continue stepping through the code for several lines, and then choose the **Stop** button on the **Debug** toolbar. Click the red dot in the left margin to remove the breakpoint.

Setting Conditional Breakpoints

To set conditional breakpoints, follow these steps:

1. In the Program.cs file locate the LoadList method. Set a breakpoint on the following line of code:

```
numList.Add(i);
```

2. Open the Breakpoints window by selecting Debug ► Windows ► Breakpoints. You should see the breakpoint you just set listed in the Breakpoints window (see Figure 5-14).




Figure 5-14. Breakpoints window

3. Right-click the breakpoint in the Breakpoints window and select Condition. You will see the Breakpoint Condition dialog box. Enter `i == 3` as the condition expression and click the OK button (see Figure 5-15).

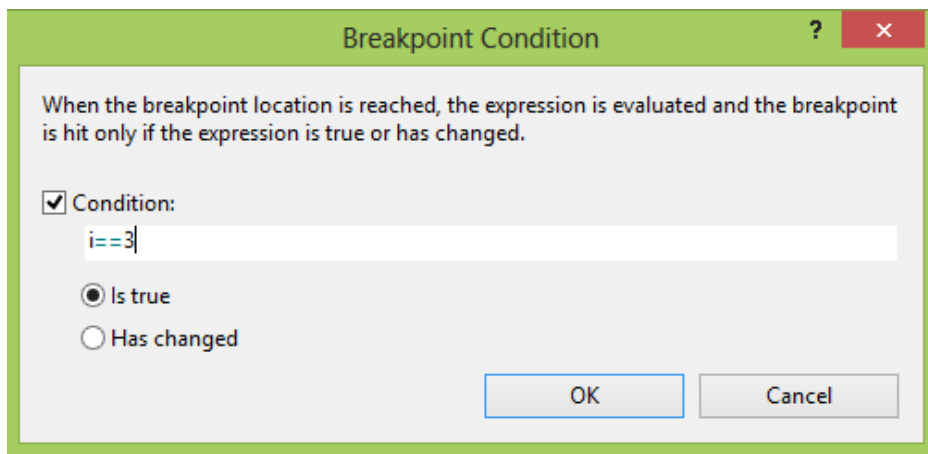
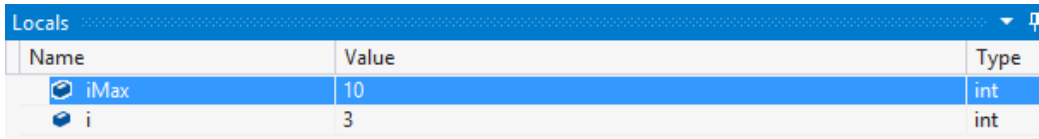


Figure 5-15. Breakpoint Condition dialog box

4. Select Debug ► Start. Program execution will pause, and you will see a yellow arrow indicating the next line that will be executed.

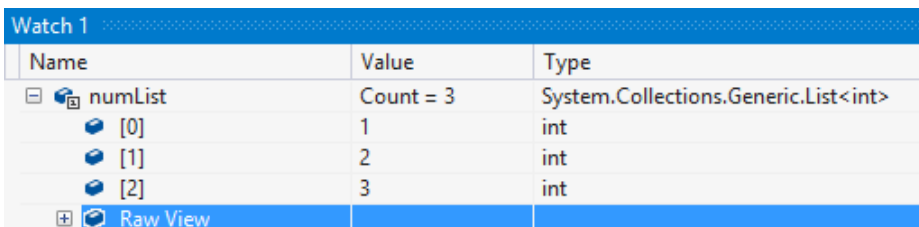
5. Select Debug ► Windows ► Locals. The Locals window is displayed at the bottom of the screen (see Figure 5-16). The value of `i` is displayed in the Locals window. Verify that it is 3. Step through the code using the Debug toolbar and watch the value of `i` change in the Locals window. Click the Stop Debugging button in the Debug toolbar.



Name	Value	Type
iMax	10	int
i	3	int

Figure 5-16. *Locals window*

6. Locate the Breakpoints window at the bottom of your screen. Right-click the breakpoint in the Breakpoints window and select Condition. Clear the current condition by clearing the Condition check box, and then click the OK button.
7. Right-click the breakpoint in the Breakpoints window and select Hit Count. Set the breakpoint to break when the hit count equals 4, and then click OK.
8. Select Debug ► Start. Program execution will pause and the yellow arrow indicates the next line of code that will execute.
9. Right-click the `numList` statement and select Add Watch. A Watch window will be displayed with `numList` in it. Notice that `numList` is a `System.Collections.Generic.List` type. Click the plus sign next to `numList`. Verify that the list contains three items (see Figure 5-17). Step through the code and watch the array fill with items. Click the Stop button in the Debug toolbar.



Name	Value	Type
numList	Count = 3	System.Collections.Generic.List<int>
[0]	1	int
[1]	2	int
[2]	3	int
Raw View		

Figure 5-17. *The Watch window*

Locating and Fixing Build Errors

To locate and fix build errors, follow these steps:

1. In the `Program` class, locate the following line of code and comment it out by placing a two slashes in front of it, as shown here:


```
//static List<int> numList = new List<int>();
```
2. Notice the red squiggly lines under the `numList` in the code. This indicates a build error that must be fixed before the application can run. Hovering over the line reveals more information about the error.
3. Select **Build** ► **Build Solution**. The Error List window will appear at the bottom of the screen, indicating a build error (see Figure 5-18).

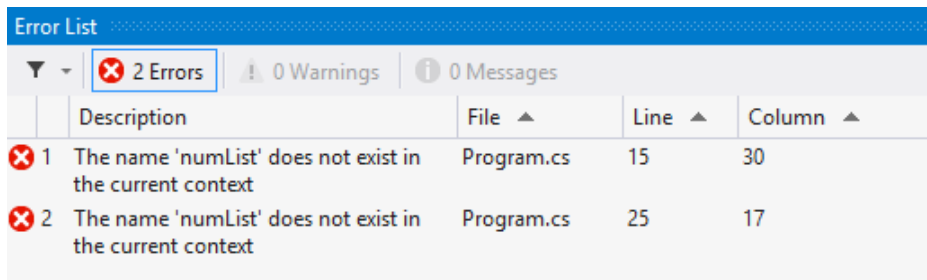


Figure 5-18. Locating build errors with the Error List window

4. Double-click the line containing the build error in the Error List window. The corresponding code will become visible in the code editor.
5. Uncomment the line you commented in step 1 by deleting the slashes. Select **Build** ► **Build Solution**. This time, the Output window is displayed at the bottom of the screen, indicating that there were no build errors.
6. Save the project and exit VS.

Summary

This chapter introduced you to the fundamentals of the .NET Framework. You reviewed some of the underlying goals of the .NET Framework. You also looked at how the .NET Framework is structured and how code is compiled and executed by the CLR. These concepts are relevant and consistent across all .NET-compliant programming languages. In addition, you explored some of the features of the Visual Studio integrated development environment. While you do not need an IDE to develop .NET applications, using one is invaluable when it comes to productivity, debugging, and code management.

There are many types of applications you can develop using C#. These include Windows desktop apps using WPF, web apps using ASP.NET, Windows Store apps using WinRT, and Application Service apps using WCF or ASP.NET Web API. You will look at each of these technologies in future chapters, but first you need to learn how to program using C#. The next chapter is the first in a series that looks at how the OOP concepts—such as class structure, inheritance, and polymorphism—are implemented in C# code.



Creating Classes

In the Chapter 5, you looked at how the .NET Framework was developed and how programs execute under the framework. That chapter introduced you to the Visual Studio IDE, and you gained some familiarity with working in it. You are now ready to start coding! This chapter is the first of a series that will introduce you to how classes are created and used in C#. It covers the basics of creating and using classes. You will create classes, add attributes and methods, and instantiate object instances of the classes in client code.

After reading this chapter, you should be familiar with the following:

- how objects used in OOP depend on class definition files
- the important role encapsulation plays in OOP
- how to define the properties and methods of a class
- the purpose of class constructors
- how to use instances of classes in client code
- the process of overloading class constructors and methods
- how to create and test class definition files with Visual Studio

Introducing Objects and Classes

In OOP, you use objects in your programs to encapsulate the data associated with the entities with which the program is working. For example, a human resources application needs to work with employees. Employees have attributes associated with them that need to be tracked. You may be interested in such things as the employee names, addresses, departments, and so on. Although you track the same attributes for all employees, each employee has unique values for these attributes. In the human resources application, an `Employee` object obtains and modifies the attributes associated with an employee. In OOP, the attributes of an object are referred to as properties.

Along with the properties of the employees, the human resource application also needs an established set of behaviors exposed by the `Employee` object. For example, one employee behavior of interest to the human resources department is the ability to request time off. In OOP, objects expose behaviors through methods. The `Employee` object contains a `RequestTimeOff` method that encapsulates the implementation code.

The properties and methods of the objects used in OOP are defined through classes. A class is a blueprint that defines the attributes and behaviors of the objects that are created as instances of the class. If you have completed the proper analysis and design of the application, you should be able to refer to the UML design documentation to determine which classes need to be constructed and what properties and methods these classes will contain. The UML class diagram contains the initial information you need to construct the classes of the system.

To demonstrate the construction of a class using C#, you will look at the code for a simple `Employee` class. The `Employee` class will have properties and methods that encapsulate and work with employee data as part of a fictitious human resources application.

Defining Classes

Let's examine the source code needed to create a class definition. The first line of code defines the code block as a class definition using the keyword `class` followed by the name of the class. The body of the class definition is enclosed by an open and closing curly bracket. The code block is structured like this:

```
class Employee
{
}
```

Creating Class Properties

After defining the starting and ending point of the class code block, the next step is to define the instance variables (often referred to as fields) contained in the class. These variables hold the data that an instance of your class will manipulate. The `private` keyword ensures that these instance variables can be manipulated only by the code inside the class. Here are the instance variable definitions:

```
private int _empID;
private string _loginName;
private string _password;
private string _department;
private string _name;
```

When a user of the class (client code) needs to query or set the value of these instance variables, public properties are exposed to them. Inside the property block of code are a `Get` block and a `Set` block. The `Get` block returns the value of the private instance variable to the user of the class. This code provides a readable property. The `Set` block provides a write-enabled property; it passes a value sent in by the client code to the corresponding private instance variable. Here is an example of a property block:

```
public string Name
{
    get { return _name; }
    set { _name = value; }
}
```

There may be times when you want to restrict access to a property so that client code can read the property value but not change it. By eliminating the `set` block inside the property block, you create a read-only property. The following code shows how to make the `EmployeeID` property read-only:

```
public int EmployeeID
{
    get { return _empID; }
}
```

■ **Note** The `private` and `public` keywords affect the scope of the code. For more information about scoping, see Appendix A.

Newcomers to OOP often ask why it's necessary to go through so much work to get and set properties. Couldn't you just create public instance variables that the user could read and write to directly? The answer lies in one of the fundamental tenets of OOP: encapsulation. Encapsulation means that the client code does not have direct access to the data. When working with the data, the client code must use clearly defined properties and methods accessed through an instance of the class. The following are some of the benefits of encapsulating the data in this way:

- preventing unauthorized access to the data
- ensuring data integrity through error checking
- creating read-only or write-only properties
- isolating users of the class from changes in the implementation code

For example, you could check to make sure the password is at least six characters long via the following code:

```
public string Password
{
    get { return _password; }
    set
    {
        if (value.Length >= 6)
        {
            _password = value;
        }
        else
        {
            throw new Exception("Password must be at least 6 characters");
        }
    }
}
```

Creating Class Methods

Class methods define the behaviors of the class. For example, the following defines a method for the Employee class that verifies employee logins:

```
public Boolean Login(string loginName, string password)
{
    if (loginName == "Jones" & password == "mj")
    {
        _empID = 1;
        Department = "HR";
        Name = "Mary Jones";
        return true;
    }
    else if (loginName == "Smith" & password == "js")
    {
        _empID = 2;
        Department = "IS";
        Name = "Jerry Smith";
        return true;
    }
}
```

```

    else
    {
        return false;
    }
}

```

When client code calls the `Login` method of the class, the login name and password are passed into the method (these are called input parameters). The parameters are checked. If they match a current employee, the instance of the class is populated with attributes of the employee and a Boolean value of `true` is passed back to the calling code. If the login name and password do not match a current employee, a Boolean value of `false` is passed back to the client code. The `return` keyword is used to return control to the client code with the value indicated.

In the previous method, a value is returned to the client code. This is indicated by the Boolean keyword, which assigns the Boolean type to the return value. The following `AddEmployee` method is another method of the `Employee` class. It's called when an employee needs to be added to the database, and it returns the newly assigned employee ID (as an integer type) to the client. The method also populates the object instance of the `Employee` class with the attributes of the newly added employee.

```

public int AddEmployee(string loginName, string password,
                      string department, string name)
{
    //Data normally saved to database.
    _empID = 3;
    LoginName = loginName;
    Password = password;
    Department = department;
    Name = name;
    return EmployeeID;
}

```

Not all methods return a value to the client. In this case, the method is declared with the `void` keyword to indicate there is no return value. The following method updates a password, but does not return a value to the client.

```

public void UpdatePassword(string password)
{
    //Data normally saved to database.
    Password = password;
}

```

ACTIVITY 6-1. CREATING THE EMPLOYEE CLASS

In this activity, you will become familiar with the following:

- how to create a C# class definition file using Visual Studio
- how to create and use an instance of the class from client code

■ **Note** If you have not already done so, download the starter files. See Appendix C for instructions. This lab is a Windows Form application which is one of the Windows client type applications you can build with Visual Studio.

Defining the Employee Class

To create the Employee class, follow these steps:

1. Start Visual Studio. Select File ► Open ► Project.
2. Navigate to the Activity6_1Starter folder, click the Activity6_1.sln file, and click Open. When the project opens, it will contain a login form. You will use this form later to test the Employee class you create.
3. Select Project ► Add Class. In the Add New Item dialog box, rename the class file to Employee.cs, and then click Add. Visual Studio adds the Employee.cs file to the project and adds the following class definition code to the file (along with the name space definition and some using declarations.):

```
class Employee
{
}
```

4. Enter the following code between the opening and closing brackets to add the private instance variables to the class body in the definition file:

```
private int _empID;
private string _loginName;
private string _password;
private int _securityLevel;
```

5. Next, add the following public properties to access the private instance variables defined in step 4:

```
public int EmployeeID
{
    get { return _empID; }
}
public string LoginName
{
    get { return _loginName; }
    set { _loginName = value; }
}
public string Password
{
    get { return _password; }
    set { _password = value; }
}
public int SecurityLevel
{
    get { return _securityLevel; }
}
```

- After the properties, add the following Login method to the class definition:

```
public Boolean Login(string loginName, string password)
{
    LoginName = loginName;
    Password = password;
    //Data normally retrieved from database.
    //Hard coded for demo only.
    if (loginName == "Smith" & password == "js")
    {
        _empID = 1;
        _securityLevel = 2;
        return true;
    }
    else if (loginName == "Jones" & password == "mj")
    {
        _empID = 2;
        _securityLevel = 4;
        return true;
    }
    else
    {
        return false;
    }
}
```

- Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.

Testing the Employee Class

To test the Employee class, follow these steps:

- Open frmLogin in the code editor and locate the btnLogin click event code.

■ **Tip** Double-clicking the Login button in the form designer will also bring up the event code in the code editor.

- In the body of the btnLogin click event, declare and instantiate a variable of type Employee called oEmployee:

```
Employee oEmployee = new Employee();
```

- Next, call the Login method of the oEmployee object, passing in the values of the login name and the password from the text boxes on the form. Capture the return value in a Boolean variable called bResult:

```
Boolean bResult = oEmployee.Login(txtName.Text,txtPassword.Text);
```

4. After calling the `Login` method, if the result is true, show a message box stating the user's security level, which is retrieved by reading the `SecurityLevel` property of the `oEmployee` object. If the result is false, show a login failed message.

```
If (bResult == true)
{
    MessageBox.Show("Your security level is " + oEmployee.SecurityLevel);
}
else
{
    MessageBox.Show("Login Failed");
}
```

5. Select **Build** ► **Build Solution**. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
6. Select **Debug** ► **Start** to run the project. Test the login form by entering a login name of Smith and a password of js. You should get a message indicating a security level of 2. Try entering your name and a password of pass. You should get a message indicating the login failed.
7. After testing the login procedure, close the form; this will stop the debugger.

■ **Note** The code in this chapter uses exceptions for invalid logins. This is just for demonstration purposes. Throwing errors is an expensive process and should not be used for business processing logic. For more information on proper use of throwing errors see Appendix B.

Using Constructors

In OOP, you use constructors to perform any processing that needs to occur when an object instance of the class becomes instantiated. For example, you could initialize properties of the object instance or establish a database connection. The class constructor method is named the same as the class. When an object instance of a class is instantiated by client code, the constructor method is executed. The following constructor is used in the `Employee` class to initialize the properties of an object instance of the `Employee` class. An employee ID is passed in to the constructor to retrieve the values from data storage, like so:

```
public Employee(int empID)
{
    //Retrieval of data hardcoded for demo
    if (empID == 1)
    {
        _empID = 1;
        LoginName = "Smith";
        Password = "js";
        Department = "IT";
        Name = "Jerry Smith";
    }
}
```

```

else if (empID == 2)
{
    _empID = 2;
    LoginName = "Jones";
    Password = "mj";
    Department = "HR";
    Name = "Mary Jones";
}
else
{
    throw new Exception("Invalid EmployeeID");
}
}

```

Overloading Methods

The ability to overload methods is a useful feature of OOP languages. You overload methods in a class by defining multiple methods that have the same name but contain different signatures. A method signature is a combination of the name of the method and its parameter type list. If you change the parameter type list, you create a different method signature. For example, the parameter type lists can contain a different number of parameters or different parameter types. The compiler will determine which method to execute by examining the parameter type list passed in by the client.

■ **Note** Changing how a parameter is passed (in other words, from `byVal` to `byRef`) does not change the method signature. Altering the return type of the method also does not create a unique method signature. For a more detailed discussion of method signatures and passing arguments, refer to Appendix A.

Suppose you want to provide two methods of the `Employee` class that will allow you to add an employee to the database. The first method assigns a username and password to the employee when the employee is added. The second method adds the employee information but defers the assignment of username and password until later. You can easily accomplish this by overloading the `AddEmployee` method of the `Employee` class, as the following code demonstrates:

```

public int AddEmployee(string loginName, string password,
    string department, string name)
{
    //Data normally saved to database.
    _empID = 3;
    LoginName = loginName;
    Password = password;
    Department = department;
    Name = name;
    return EmployeeID;
}

```

```
public int AddEmployee(string department, string name)
{
    //Data normally saved to database.
    _empID = 3;
    Department = department;
    Name = name;
    return EmployeeID;
}
```

Because the parameter type list of the first method (string, string, string, string) differs from the parameter type list of the second method (string, string), the compiler can determine which method to invoke. A common technique in OOP is to overload the constructor of the class. For example, when an instance of the Employee class is created, one constructor could be used for new employees and another could be used for current employees by passing in the employee ID when the class instance is instantiated by the client. The following code shows the overloading of a class constructor:

```
public Employee()
{
    _empID = -1;
}

public Employee(int empID)
{
    //Retrieval of data hard coded for demo
    if (empID == 1)
    {
        _empID = 1;
        LoginName = "Smith";
        Password = "js";
        Department = "IT";
        Name = "Jerry Smith";
    }
    else if (empID == 2)
    {
        _empID = 2;
        LoginName = "Jones";
        Password = "mj";
        Department = "HR";
        Name = "Mary Jones";
    }
    else
    {
        throw new Exception("Invalid EmployeeID");
    }
}
```


ACTIVITY 6-2. CREATING CONSTRUCTORS AND OVERLOADING METHODS

In this activity, you will become familiar with the following:

- how to create and overload the class constructor method
- how to use overloaded constructors of a class from client code
- how to overload a method of a class
- how to use overloaded methods of a class from client code

Creating and Overloading Class Constructors

To create and overload class constructors, follow these steps:

1. Start Visual Studio. Select File ► Open ► Project.
2. Navigate to the Activity6_2Starter folder, click the Activity6_2.sln file, and then click Open. When the project opens, it will contain a frmEmployeeInfo form that you will use to test the Employee class. The project also includes the Employee.cs file, which contains the Employee class definition code.
3. Open Employee.cs in the code editor and examine the code. The class contains several properties pertaining to employees that need to be maintained.
4. After the property declaration code, add the following private method to the class. This method simulates the generation of a new employee ID.

```
private int GetNextID()
{
    //simulates the retrieval of next
    //available id from database
    return 100;
}
```

5. Create a default class constructor, and add code that calls the GetNextID method and assigns the return value to the private instance variable _empID:

```
public Employee()
{
    _empID = GetNextID();
}
```

6. Overload the default constructor method by adding a second constructor method that takes an integer parameter of empID, like so:

```
public Employee(int empID)
{
    //Constructor for existing employee
}
```

7. Add the following code to the overloaded constructor, which simulates extracting the employee data from a database and assigns the data to the instance properties of the class:

```
//Simulates retrieval from database
if (empID == 1)
{
    _empID = empID;
    LoginName = "smith";
    Password = "js";
    SSN = 123456789;
    Department = "IS";
}
else if (empID == 2)
{
    _empID = empID;
    LoginName = "jones";
    Password = "mj";
    SSN = 987654321;
    Department = "HR";
}
else
{
    throw new Exception("Invalid Employee ID");
}
```

8. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.

Testing the Employee Class Constructors

To test the Employee class constructors, follow these steps:

1. Open the EmployeeInfoForm in the form editor and double click the New Employee button to bring up the click event code in the code editor.
2. In the click event method body, declare and instantiate a variable of type Employee called oEmployee:

```
Employee oEmployee = new Employee();
```

3. Next, update the employeeID text box with the employee ID, disable the employee ID text box, and clear the remaining textboxes:

```
txtEmpID.Text = oEmployee.EmpID.ToString();
txtEmpID.Enabled = false;
txtLoginName.Text = "";
txtPassword.Text = "";
txtSSN.Text = "";
txtDepartment.Text = "";
```

4. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
5. Open the EmployeeInfoForm in the form editor and double click the Existing Employee button to bring up the click event code in the code editor.
6. In the click event method body, declare and instantiate a variable of type Employee called oEmployee. Retrieve the employee ID from the txtEmpID text box and pass it as an argument in the constructor. The int.Parse method converts the text to an integer data type:

```
Employee oEmployee = new Employee(int.Parse(txtEmpID.Text));
```

7. Next, disable the employee ID textbox and fill in the remaining text boxes with the values of the Employee object's properties:

```
txtEmpID.Enabled = false;
txtLoginName.Text = oEmployee.LoginName;
txtPassword.Text = oEmployee.Password;
txtSSN.Text = oEmployee.SSN.ToString();
txtDepartment.Text = oEmployee.Department;
```

8. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
9. Select Debug ► Start to run the project and test the code.
10. When the EmployeeInfo form is displayed, click the New Employee button. You should see that a new employee ID has been generated in the employee ID textbox.
11. Click the Reset button to clear and enable the employee ID text box.
12. Enter a value of 1 for the employee ID and click the Existing Employee button. The information for the employee is displayed on the form.
13. After testing the constructors, close the form, which will stop the debugger.

Overloading a Class Method

To overload a class method, follow these steps:

1. Open the Employee.cs code in the code editor.
2. Add the following Update method to the Employee class. This method simulates the updating of the employee security information to a database:

```
public string Update(string loginName, string password)
{
    LoginName = loginName;
    Password = password;
    return "Security info updated.";
}
```

3. Add a second `Update` method to simulate the updating of the employee human resources data to a database:

```
public string Update(int ssNumber, string department)
{
    SSN = ssNumber;
    Department = department;
    return "HR info updated.";
}
```

4. Select **Build** ► **Build Solution**. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.

Testing the Overloaded Update Method

To test the overloaded `Update` method, follow these steps:

1. Open the `EmployeeInfo` Form in the Form editor and double click the `Update SI` button. You are presented with the click event code in the Code Editor window.
2. In the click event method, declare and instantiate a variable of type `Employee` called `oEmployee`. Retrieve the employee ID from the `txtEmpID` text box and pass it as an argument in the constructor:

```
Employee oEmployee = new Employee(int.Parse(txtEmpID.Text));
```

3. Next, call the `Update` method, passing the values of the login name and password from the text boxes. Show the method return message to the user in a message box:

```
MessageBox.Show(oEmployee.Update(txtLoginName.Text, txtPassword.Text));
```

4. Update the login name and password text boxes with the property values of the `Employee` object:

```
txtLoginName.Text = oEmployee.LoginName;
txtPassword.Text = oEmployee.PassWord;
```

5. Repeat this process to add similar code to the `Update HR` button click event method to simulate updating the human resources information. Add the following code to the click event method:

```
Employee oEmployee = new Employee(int.Parse(txtEmpID.Text));
MessageBox.Show(oEmployee.Update(int.Parse(txtSSN.Text), txtDepartment.Text));
txtSSN.Text = oEmployee.SSN.ToString();
txtDepartment.Text = oEmployee.Department;
```

6. Select **Build** ► **Build Solution**. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
7. Select **Debug** ► **Start** to run the project and test the code.

8. Enter a value of 1 for the employee ID and click the Existing Employee button.
 9. Change the values for the security information and click the Update button.
 10. Change the values for the human resources information and click the Update button.
 11. You should see that the correct Update method is called in accordance with the parameters passed into it. After testing the Update method, close the form.
-

Summary

This chapter gave you a firm foundation in creating and using classes in C# code. Now that you are comfortable constructing and using classes, you are ready to look at implementing some of the more advanced features of OOP. In Chapter 7, you will concentrate on how inheritance and polymorphism are implemented in C# code. As an object-oriented programmer, it is important for you to become familiar with these concepts and learn how to implement them in your programs.

CHAPTER 7



Creating Class Hierarchies

In Chapter 6, you learned how to create classes, add attributes and methods, and instantiate object instances of the classes in client code. This chapter introduces the concepts of inheritance and polymorphism.

Inheritance is one of the most powerful and fundamental features of any OOP language. Using inheritance, you create base classes that encapsulate common functionality. Other classes can be derived from these base classes. The derived classes inherit the properties and methods of the base classes and extend the functionality as needed.

A second fundamental OOP feature is polymorphism. *Polymorphism* lets a base class define methods that must be implemented by any derived classes. The base class defines the message signature that derived classes must adhere to, but the implementation code of the method is left up to the derived class. The power of polymorphism lies in the fact that clients know they can implement methods of classes of the base type in the same fashion. Even though the internal processing of the method may be different, the client knows the inputs and outputs of the methods will be the same.

After reading this chapter, you will learn the following:

- how to create and use base classes
- how to create and use derived classes
- how access modifiers control inheritance
- how to override base class methods
- how to implement interfaces
- how to implement polymorphism through inheritance and through interfaces

Understanding Inheritance

One of the most powerful features of any OOP language is inheritance. Inheritance is the ability to create a base class with properties and methods that can be used in classes derived from the base class.

Creating Base and Derived Classes

The purpose of inheritance is to create a base class that encapsulates properties and methods that can be used by derived classes of the same type. For example, you could create a base class `Account`. A `GetBalance` method is defined in the `Account` class. You can then create two separate classes: `SavingsAccount` and `CheckingAccount`. Because the `SavingsAccount` class and the `CheckingAccount` class use the same logic to retrieve balance information, they inherit the `GetBalance` method from the base class `Account`. This enables you to create one common code base that is easier to maintain and manage.

Derived classes are not limited to the properties and methods of the base class, however. The derived classes may require additional methods and properties that are unique to their needs. For example, the business rules for withdrawing money from a checking account may require that a minimum balance be maintained. A minimum balance, however, may not be required for withdrawals from a savings account. In this scenario, the derived `CheckingAccount` and `SavingsAccount` classes would each need their own unique definition for a `Withdraw` method.

To create a derived class in C#, you enter the name of the class, followed by a colon (`:`) and the name of the base class. The following code demonstrates how to create a `CheckingAccount` class that derives from an `Account` base class:

```
class Account
{
    long _accountNumber;

    public long AccountNumber
    {
        get { return _accountNumber; }
        set { _accountNumber = value; }
    }
    public double GetBalance()
    {
        //code to retrieve account balance from database
        return (double)10000;
    }
}

class CheckingAccount : Account
{
    double _minBalance;

    public double MinBalance
    {
        get { return _minBalance; }
        set { _minBalance = value; }
    }
    public void Withdraw(double amount)
    {
        //code to withdraw from account
    }
}
```

The following code could be implemented by a client creating an object instance of `CheckingAccount`. Notice that the client perceives no distinction between the call to the `GetBalance` method and the call to the `Withdraw` method. In this case, the client has no knowledge of the `Account` class; instead, both methods appear to have been defined by `CheckingAccount`.

```
CheckingAccount oCheckingAccount = new CheckingAccount();
double balance;
oCheckingAccount.AccountNumber = 1000;
balance = oCheckingAccount.GetBalance();
oCheckingAccount.Withdraw(500);
```

Creating a Sealed Class

By default, any C# class can be inherited. When creating classes that can be inherited, you must take care that they are not modified in such a way that derived classes no longer function as intended. If you are not careful, you can create complex inheritance chains that are hard to manage and debug. For example, suppose you create a derived `CheckingAccount` class based on the `Account` class. Another programmer can come along and create a derived class based on the `CheckingAccount` and use it in ways you never intended. (This could easily occur in large programming teams with poor communication and design methods.)

By using the sealed modifier, you can create classes that you know will not be derived from. This type of class is often referred to as a *sealed* or *final* class. By making a class not inheritable, you avoid the complexity and overhead associated with altering the code of base classes. The following code demonstrates the use of the sealed modifier when constructing a class definition:

```
sealed class CheckingAccount : Account
```

Creating an Abstract Class

At this point in the example, a client can access the `GetBalance` method through an instance of the derived `CheckingAccount` class or directly through an instance of the base `Account` class. Sometimes, you may want to have a base class that can't be instantiated by client code. Access to the methods and properties of the class must be through a derived class. In this case, you construct the base class using the abstract modifier. The following code shows the `Account` class definition with the abstract modifier:

```
abstract class Account
```

This makes the `Account` class an abstract class. For clients to gain access to the `GetBalance` method, they must create an instance of the derived `CheckingAccount` class.

Using Access Modifiers in Base Classes

When setting up class hierarchies using inheritance, you must manage how the properties and methods of your classes are accessed. Two access modifiers you have looked at so far are `public` and `private`. If a method or property of the base class is exposed as `public`, it is accessible by both the derived class and any client of the derived class. If you expose the property or method of the base class as `private`, it is not accessible directly by the derived class or the client.

You may want to expose a property or method of the base class to a derived class, but not to a client of the derived class. In this case, you use the protected access modifier. The following code demonstrates the use of the protected access modifier:

```
protected double GetBalance()
{
    //code to retrieve account balance from database
    return (double)10000;
}
```

By defining the `GetBalance` method as protected, it becomes accessible to the derived class `CheckingAccount`, but not to the client code accessing an instance of the `CheckingAccount` class.

ACTIVITY 7-1. IMPLEMENTING INHERITANCE USING BASE AND DERIVED CLASSES

In this activity, you will become familiar with the following:

- creating a base class and derived classes that inherit its methods
- using the protected access modifier to restrict use of base class methods
- creating an abstract base class

Creating a Base Class and Derived Classes

To create the `Account` class, follow these steps:

1. Start Visual Studio. Select **File** ► **Open** ► **Project**.
2. Navigate to the `Activity7_1Starter` folder, click the `Activity7_1.sln` file, and then click **Open**. When the project opens, it will contain a teller form. You will use this form later to test the classes you create.
3. In the Solution Explorer window, right click the **Project** node and select **Add** ► **Class**.
4. In the Add New Item dialog box, rename the class file as `Account.cs` and click **Open**. The `Account.cs` file is added to the project, and the `Account` class definition code is added to the file.
5. Add the following code to the class definition file to create the private instance variable (private is the default modifier for instance variables):

```
int _accountNumber;
```

6. Add the following `GetBalance` method to the class definition:

```
public double GetBalance(int accountNumber)
{
    _accountNumber = accountNumber;
    //Data normally retrieved from database.
    if (_accountNumber == 1)
    {
        return 1000;
    }
    else if (_accountNumber == 2)
    {
        return 2000;
    }
    else
    {
        return -1; //Account number is incorrect
    }
}
```

- After the `Account` class, add the following code to create the `CheckingAccount` and `SavingsAccount` derived classes:

```
class CheckingAccount : Account
{
}
class SavingsAccount : Account
{
}
```

- Select **Build** ► **Build Solution**. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.

Testing the Classes

To test the classes, follow these steps:

- Open the Teller form in the code editor and locate the `btnGetBalance` click event code.
- Inside the event procedure, prior to the try block, declare and instantiate a variable of type `CheckingAccount` called `oCheckingAccount`, a variable of type `SavingsAccount` called `oSavingsAccount`, and a variable of type `Account` called `oAccount`:

```
CheckingAccount oCheckingAccount = new CheckingAccount();
SavingsAccount oSavingsAccount = new SavingsAccount();
Account oAccount = new Account();
```

- Depending on which radio button is selected, call the `GetBalance` method of the appropriate object and pass the account number value from the `AccountNumber` text box. Show the return value in the `Balance` text box. Place the following code in the try block prior to the catch statement:

```
if (rdbChecking.Checked)
{
    txtBalance.Text =
        oCheckingAccount.GetBalance(int.Parse(txtAccountNumber.Text)).ToString();
}
else if (rdbSavings.Checked)
{
    txtBalance.Text =
        oSavingsAccount.GetBalance(int.Parse(txtAccountNumber.Text)).ToString();
}
else if (rdbGeneral.Checked)
{
    txtBalance.Text =
        oAccount.GetBalance(int.Parse(txtAccountNumber.Text)).ToString();
}
```

- Select **Build** ► **Build Solution**. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.

5. Select **Debug** ► Start to run the project. Enter an account number of 1 and click the **Get Balance** button for the **Checking Account** type. You should get a balance of 1,000. Test the other account types. You should get the same result, since all classes are using the same `GetBalance` function defined in the base class.
6. After testing, close the form, which will stop the debugger.

Restricting Use of a Base Class Method to Its Derived Classes

At this point, the `GetBalance` method of the base class is `public`, which means that it can be accessed by derived classes and their clients. Let's alter this so that the `GetBalance` method can be accessed only by the derived classes alone, and not by their clients. To protect the `GetBalance` method in this way, follow these steps:

1. Locate the `GetBalance` method of the `Account` class.
2. Change the access modifier of the `GetBalance` method from `public` to `protected`.
3. Switch to the `frmTeller` code editor and locate the `btnGetBalance` click event code.
4. Hover the cursor over the call to the `GetBalance` method of the `oCheckingAccount` object. You will see a warning stating that it is a protected function and is not accessible in this context.
5. Comment out the code between the `try` and the `catch` statements.
6. Switch to the `Account.cs` code editor.
7. Add the following code to create the following private instance variable to the `SavingsAccount` class definition file:

```
double _dblBalance;
```

8. Add the following `Withdraw` method to the `SavingsAccount` class. This function calls the protected method of the `Account` base class:

```
public double Withdraw(int accountNumber, double amount)
{
    _dblBalance = GetBalance(accountNumber);
    if (_dblBalance >= amount)
    {
        _dblBalance -= amount;
        return _dblBalance;
    }
    else
    {
        Return -1; //Not enough funds
    }
}
```

9. Select **Build** ► **Build Solution**. Make sure there are no build errors in the **Error List** window. If there are, fix them, and then rebuild.

Testing the Protected Base Class Method

To test the `Withdraw` method, follow these steps:

1. Open the `frmTeller` form in the code editor and locate the `btnWithdraw` click event code.
2. Inside the event procedure, prior to the `try` block, declare and instantiate a variable of type `SavingsAccount` called `oSavingsAccount`.

```
SavingsAccount oSavingsAccount = new SavingsAccount();
```

3. Call the `Withdraw` method of the `oSavingsAccount`. Pass the account number value from the `AccountNumber` text box and the withdrawal amount from the `Amount` text box. Show the return value in the `Balance` text box. Place the following code in the `try` block prior to the catch statement:

```
txtBalance.Text = oSavingsAccount.Withdraw  
(int.Parse(txtAccountNumber.Text), double.Parse(txtAmount.Text)).ToString();
```

4. Select **Build** ► **Build Solution**. Make sure there are no build errors in the Error List window. If there are, fix them and then rebuild.
5. Select **Debug** ► **Start** to run the project.
6. Test the `Withdraw` method of the `SavingsAccount` class by entering an account number of 1 and a withdrawal amount of 200. Click the `Withdraw` button. You should get a resulting balance of 800.
7. Enter an account number of 1 and a withdrawal amount of 2000. Click the `Withdraw` button. You should get a -1 indicating there are insufficient funds.
8. After testing the `Withdraw` method, close the form, which will stop the debugger.

Restricting Use of All Members of a Base Class to Its Derived Classes

Because the `Account` base class is public, it can be instantiated by clients of the derived classes. You can alter this by making the `Account` base class an abstract class. An abstract class can be accessed only by its derived classes and can't be instantiated and accessed by their clients. To create and test the accessibility of the abstract class, follow these steps:

1. Locate the `Account` class definition in the `Account.cs` code.
2. Add the `abstract` keyword to the class definition code, like so:

```
abstract class Account
```

3. Select **Build** ► **Build Solution**. You should receive a build error in the Error List window. Find the line of code causing the error.

```
Account oAccount = new Account();
```

4. Comment out the line of code, and select Build ► Build Solution again. It should now build without any errors.
 5. Save and close the project.
-

Overriding the Methods of a Base Class

When a derived class inherits a method from a base class, it inherits the implementation of that method. As the designer of the base class, you may want to let a derived class implement the method in its own unique way. This is known as overriding the base class method.

By default, a derived class can't override the implementation code of its base class. To allow a base class method to be overridden, you must include the keyword `virtual` in the method definition. In the derived class, you define a method with the same method signature and indicate it is overriding a base class method with the `override` keyword. The following code demonstrates the creation of an overridable `Deposit` method in the `Account` base class:

```
public virtual void Deposit(double amount)
{
    //Base class implementation
}
```

To override the `Deposit` method in the derived `CheckingAccount` class, use the following code:

```
public override void Deposit(double amount)
{
    //Derived class implementation
}
```

One scenario to watch for is when a derived class inherits from the base class and a second derived class inherits from the first derived class. When a method overrides a method in the base class, it becomes overridable by default. To limit an overriding method from being overridden further up the inheritance chain, you must include the `sealed` keyword in front of the `override` keyword in the method definition of the derived class. The following code in the `CheckingAccount` class prevents the overriding of the `Deposit` method if the `CheckingAccount` class is derived from:

```
public sealed override void Deposit(double amount)
{
    //Derived class implementation
}
```

When you indicate that a base class method is overridable, derived classes have the option of overriding the method or using the implementation provided by the base class. In some cases, you may want to use a base class method as a template for the derived classes. The base class has no implementation code, but is used to define the method signatures used in the derived classes. This type of class is referred to as an *abstract base class*. You define the class and the methods with the `abstract` keyword. The following code is used to create an abstract `Account` base class with an abstract `Deposit` method:

```
public abstract class Account
{
    public abstract void Deposit(double amount);
}
```

Notice that because there is no implementation code defined in the base class for the `Deposit` method, the body of the method is omitted.

Calling a Derived Class Method from a Base Class

A situation may arise in which you are calling an overridable method in the base class from another method of the base class, and the derived class overrides the method of the base class. When a call is made to the base class method from an instance of the derived class, the base class will call the overridden method of the derived class. The following code shows an example of this situation. A `CheckingAccount` base class contains an overridable `GetMinBalance` method. The `InterestBearingCheckingAccount` class, inheriting from the `CheckingAccount` class, overrides the `GetMinBalance` method.

```
class CheckingAccount
{
    private double _balance = 2000;

    public double Balance
    {
        get { return _balance; }
    }
    public virtual double GetMinBalance()
    {
        return 200;
    }
    public virtual void Withdraw(double amount)
    {
        double minBalance = GetMinBalance();
        if (minBalance < (Balance - amount))
        {
            _balance -= amount;
        }
        else
        {
            throw new Exception("Minimum balance error.");
        }
    }
}
class InterestBearingCheckingAccount : CheckingAccount
{
    public override double GetMinBalance()
    {
        return 1000;
    }
}
```

A client creates an object instance of the `InterestBearingCheckingAccount` class and calls the `Withdraw` method. In this case, the overridden `GetMinimumBalance` method of the `InterestBearingCheckingAccount` class is executed, and a minimum balance of 1,000 is used.

```
InterestBearingCheckingAccount oAccount = new InterestBearingCheckingAccount();
oAccount.Withdraw(500);
```

When the call was made to the `Withdraw` method, you could have prefaced it with the `this` qualifier:

```
double minBalance = this.GetMinBalance();
```

Because the `this` qualifier is the default qualifier if none is used, the code would execute the same way as previously demonstrated. The most derived class implementation (that has been instantiated) of the method is executed. In other words, if a client instantiates an instance of the `InterestBearingCheckingAccount` class, as was demonstrated previously, the base class's call to `GetMinimumBalance` is made to the derived class's implementation. On the other hand, if a client creates an instance of the `CheckingAccount` class, the base class's call to `GetMinimumBalance` is made to its own implementation.

Calling a Base Class Method from a Derived Class

In some cases, you may want to develop a derived class method that still uses the implementation code in the base class but also augments it with its own implementation code. In this case, you create an overriding method in the derived class and call the code in the base class using the base qualifier. The following code demonstrates the use of the base qualifier:

```
public override void Deposit(double amount)
{
    base.Deposit(amount);
    //Derived class implementation.
}
```

Overloading Methods of a Base Class

Methods inherited by the derived class can be overloaded. The method signature of the overloaded class must use the same name as the overloaded method, but the parameter lists must differ. This is the same as when you overload methods of the same class. The following code demonstrates the overloading of a derived method:

```
class CheckingAccount
{
    public void Withdraw(double amount)
    {
    }
}
class InterestBearingCheckingAccount : CheckingAccount
{
    public void Withdraw(double amount, double minBalance)
    {
    }
}
```

Client code instantiating an instance of the `InterestBearingCheckingAccount` has access to both `Withdraw` methods.

```
InterestBearingCheckingAccount oAccount = new InterestBearingCheckingAccount();
oAccount.Withdraw(500);
oAccount.Withdraw(500, 200);
```

Hiding Base Class Methods

If a method in a derived class has the same method signature as that of the base class method, but it is not marked with the `override` keyword, it effectively hides the method of the base class. Although this may be the intended behavior, sometimes it can occur inadvertently. Although the code will still compile, the IDE will issue a warning asking if this is the intended behavior. If you intend to hide a base class method, you should explicitly use the `new` keyword in the definition of the method of the derived class. Using the `new` keyword will indicate to the IDE this is the intended behavior and dismiss the warning. The following code demonstrates hiding a base class method:

```
class CheckingAccount
{
    public virtual void Withdraw(double amount)
    {
    }
}

class InterestBearingCheckingAccount : CheckingAccount
{
    public new void Withdraw(double amount)
    {
    }
    public void Withdraw(double amount, double minBalance)
    {
    }
}
```

ACTIVITY 7-2. OVERRIDING BASE CLASS METHODS

In this activity, you will become familiar with the following:

- overriding methods of a base class
- using the base qualifier in a derived class

Overriding Base Class Methods

To override the `Account` class, follow these steps:

1. Start VS. Select **File** ► **Open** ► **Project**.
2. Navigate to the `Activity7_2Starter` folder, click the `Activity7_2.sln` file, and then click **Open**. When the project opens, it will contain a teller form. You will use this form later to test the classes you will create. The project also contains a `BankClasses.cs` file. This file contains code for the `Account` base class and the derived classes `SavingsAccount` and `CheckingAccount`.
3. Examine the `Withdraw` method defined in the base class `Account`. This method checks to see whether there are sufficient funds in the account and, if there are, updates the balance. You will override this method in the `CheckingAccount` class to ensure that a minimum balance is maintained.

4. Change the `Withdraw` method definition in the `Account` class to indicate it is overridable, like so:

```
public virtual double Withdraw(double amount)
```

5. Add the following `GetMinimumBalance` method to the `CheckingAccount` class definition:

```
public double GetMinimumBalance()
{
    return 200;
}
```

6. Add the following overriding `Withdraw` method to the `CheckingAccount` class definition. This method adds a check to see that the minimum balance is maintained after a withdrawal.

```
public override double Withdraw(double amount)
{
    if (Balance >= amount + GetMinimumBalance())
    {
        _balance -= amount;
        return Balance;
    }
    else
    {
        return -1; //Not enough funds
    }
}
```

7. Select **Build ► Build Solution**. Make sure there are no build errors in the Error List window. If there are, fix them and then rebuild.

Testing the Overwritten Methods

To test the modified `Withdraw` methods you have created, follow these steps:

1. Open the `frmTeller` form in the code editor and locate the `btnWithdraw` click event code.
2. Depending on which radio button is selected, call the `Withdraw` method of the appropriate object and pass the value of the `txtAmount` text box. Add the following code in the `try` block to show the return value in the `txtBalance` text box:

```
if (rdbChecking.Checked)
{
    oCheckingAccount.AccountNumber = int.Parse(txtAccountNumber.Text);
    txtBalance.Text =
        oCheckingAccount.Withdraw(double.Parse(txtAmount.Text)).ToString();
}
```

```

else if (rdbSavings.Checked)
{
    oSavingsAccount.AccountNumber = int.Parse(txtAccountNumber.Text);
    txtBalance.Text =
        oSavingsAccount.Withdraw(double.Parse(txtAmount.Text)).ToString();
}

```

3. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
4. Select Debug ► Start to run the project.
5. Enter an account number of 1, choose the Checking option button, and click the Get Balance button. You should get a balance of 1000.
6. Enter a withdrawal amount of 200 and click the Withdraw button. You should get a resulting balance of 800.
7. Enter a withdrawal amount of 700 and click the Withdraw button. You should get a -1 (indicating insufficient funds) because the resulting balance would be less than the minimum balance of 200.
8. Enter an account number of 1, choose the Savings option button, and click the Get Balance button. You should get a balance of 1000.
9. Enter a withdrawal amount of 600 and click the Withdraw button. You should get a resulting balance of 400.
10. Enter a withdrawal amount of 400 and click the Withdraw button. You should get a resulting balance of 0 because there is no minimum balance for the savings account that uses the Account base class's Withdraw method.
11. After testing, close the form, which will stop the debugger.

Using the Base Qualifier to Call a Base Class Method

At this point, the Withdraw method of the CheckingAccount class overrides the Account class's Withdraw method. None of the code in the base class's method is executed. You will now alter the code so that when the CheckingAccount class's code is executed, it also executes the base class's Withdraw method. Follow these steps:

1. Locate the Withdraw method of the Account class.
2. Change the implementation code so that it decrements the balance by the amount passed to it.

```

public virtual double Withdraw(double amount)
{
    _balance -= amount;
    return Balance;
}

```

3. Change the `Withdraw` method of the `CheckingAccount` class so that after it checks for sufficient funds, it calls the `Withdraw` method of the `Account` base class.

```
public override double Withdraw(double amount)
{
    if (Balance >= amount + GetMinimumBalance())
    {
        return base.Withdraw(amount);
    }
    else
    {
        return -1; //Not enough funds.
    }
}
```

4. Add a `Withdraw` method to the `SavingsAccount` class that is similar to the `Withdraw` method of the `CheckingAccount` class, but does not check for a minimum balance.

```
public override double Withdraw(double amount)
{
    if (Balance >= amount)
    {
        return base.Withdraw(amount);
    }
    else
    {
        return -1; //Not enough funds.
    }
}
```

5. Select **Build** ► **Build Solution**. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.

Testing the Use of the Base Modifier

To test the `Withdraw` method, follow these steps:

1. Select **Debug** ► **Start**.
2. Enter an account number of 1, choose the **Checking** option button, and click the **Get Balance** button. You should get a balance of 1000.
3. Enter a withdrawal amount of 600 and click the **Withdraw** button. You should get a resulting balance of 400.
4. Enter a withdrawal amount of 300 and click the **Withdraw** button. You should get a -1 (insufficient funds) because the resulting balance would be less than the 200 minimum.
5. Enter an account number of 1, choose the **Savings** option button, and click the **Get Balance** button. You should get a balance of 1000.

6. Enter a withdrawal amount of 600 and click the Withdraw button. You should get a resulting balance of 400.
 7. Enter a withdrawal amount of 300 and click the Withdraw button. You should get a resulting balance of 100, because there is no minimum balance for the savings account that uses the Account base class's Withdraw method.
 8. After testing, close the form, which will stop the debugger.
-

Implementing Interfaces

As you saw earlier, you can create an abstract base class that does not contain any implementation code but defines the method signatures that must be used by any class that inherits from the base class. When you use an abstract class, classes that derive from it must implement its inherited methods. You could use another technique to accomplish a similar result. In this case, instead of defining an abstract class, you define an interface that defines the method signatures.

Classes that implement an interface are contractually required to implement the interface signature definition and can't alter it. This technique is useful to ensure that client code using the classes knows which methods are available, how they should be called, and the return values to expect. The following code shows how to declare an interface definition:

```
public interface IAccount
{
    string GetAccountInfo(int accountNumber);
}
```

A class implements the interface by using a semicolon followed by the name of the interface after the class name. When a class implements an interface, it must provide implementation code for all methods defined by the interface. The following code demonstrates how a `CheckingAccount` implements the `IAccount` interface:

```
public class CheckingAccount : IAccount
{
    public string GetAccountInfo(int accountNumber)
    {
        return "Printing checking account info";
    }
}
```

Because implementing an interface and inheriting from an abstract base class are similar, you might ask why you should bother using an interface. One advantage of using interfaces is that a class can implement multiple interfaces. The .NET Framework does not support inheritance from more than one class. As a workaround to multiple inheritance, the ability to implement multiple interfaces is included. Interfaces are also useful to enforce common functionality across disparate types of classes.

Understanding Polymorphism

Polymorphism is the ability of derived classes inheriting from the same base class to respond to the same method call in their own unique way. This simplifies client code because the client code does not need to worry about which class type it is referencing, as long as the class types implement the same method interfaces.

For example, suppose you want all account classes in a banking application to contain a `GetAccountInfo` method with the same interface definition but different implementations based on account type. Client code could loop through a collection of account-type classes, and the compiler would determine at runtime which specific account-type implementation needs to be executed. If you later added a new account type that implements the `GetAccountInfo` method, you would not need to alter existing client code.

You can achieve polymorphism either by using inheritance or by implementing interfaces. The following code demonstrates the use of inheritance. First, you define the base and derived classes.

```
public abstract class Account
{
    public abstract string GetAccountInfo();
}

public class CheckingAccount : Account
{
    public override string GetAccountInfo()
    {
        return "Printing checking account info";
    }
}

public class SavingsAccount : Account
{
    public override string GetAccountInfo()
    {
        return "Printing savings account info";
    }
}
```

You then create a list of type `Account` and add a `CheckingAccount` and a `SavingsAccount`.

```
List<Account> AccountList = new List<Account>();
CheckingAccount oCheckingAccount = new CheckingAccount();
SavingsAccount oSavingsAccount = new SavingsAccount();
AccountList.Add(oCheckingAccount);
AccountList.Add(oSavingsAccount);
```

You then loop through the `List` and call the `GetAccountInfo` method of each `Account`. Each `Account` type will implement its own implementation of the `GetAccountInfo`.

```
foreach (Account a in AccountList)
{
    MessageBox.Show(a.GetAccountInfo());
}
```

You can also achieve a similar result by using interfaces. Instead of inheriting from the base class `Account`, you define and implement an `IAccount` interface.

```
public interface IAccount
{
    string GetAccountInfo();
}
```

```

public class CheckingAccount : IAccount
{
    public string GetAccountInfo()
    {
        return "Printing checking account info";
    }
}
public class SavingsAccount : IAccount
{
    public string GetAccountInfo()
    {
        return "Printing savings account info";
    }
}

```

You then create a list of type `IAccount` and add a `CheckingAccount` and a `SavingsAccount`.

```

List<IAccount>AccountList = new List<IAccount>();
CheckingAccount oCheckingAccount = new CheckingAccount();
SavingsAccount oSavingsAccount = new SavingsAccount();
AccountList.Add(oCheckingAccount);
AccountList.Add(oSavingsAccount);

```

You then loop through the `AccountList` and call the `GetAccountInfo` method of each `Account`. Each `Account` type will implement its own implementation of the `GetAccountInfo`.

```

foreach (IAccount a in AccountList)
{
    MessageBox.Show(a.GetAccountInfo());
}

```

ACTIVITY 7-3. IMPLEMENTING POLYMORPHISM

In this activity, you will become familiar with the following:

- creating polymorphism through inheritance
- creating polymorphism through interfaces

Implementing Polymorphism Using Inheritance

To implement polymorphism using inheritance, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Select the Console Application template under the C# templates. Name the project `Activity7_3`.

3. The project includes a Program.cs file. This file contains a Main method that launches a Windows Console application. Right click the project node in the Solution Explorer Window and select Add ► class. Name the file Account.cs.
4. In the Account.cs file alter the code to create an abstract base Account class. Include an accountNumber property and an abstract method GetAccountInfo that takes no parameters and returns a string.

```
public abstract class Account
{
    private int _accountNumber;

    public int AccountNumber
    {
        get { return _accountNumber; }
        set { _accountNumber = value; }
    }

    public abstract string GetAccountInfo();
}
```

5. Add the following code to create two derived classes: CheckingAccount and SavingsAccount. These classes will override the GetAccountInfo method of the base class.

```
public class CheckingAccount : Account
{
    public override string GetAccountInfo()
    {
        return "Printing checking account info for account number "
            + AccountNumber.ToString();
    }
}
public class SavingsAccount : Account
{
    public override string GetAccountInfo()
    {
        return "Printing savings account info for account number "
            + AccountNumber.ToString();
    }
}
```

6. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.

Testing the Polymorphic Inheritance Method

To test the polymorphic method, follow these steps:

1. Open the Program.cs file in the code editor and locate the Main method.
2. Create an instance of a list of account types.

```
List<Account> AccountList = new List<Account>();
```

3. Create instances of CheckingAccount and SavingsAccount.

```
CheckingAccount oCheckingAccount = new CheckingAccount();
oCheckingAccount.AccountNumber = 100;
SavingsAccount oSavingsAccount = new SavingsAccount();
oSavingsAccount.AccountNumber = 200;
```

4. Add the oCheckingAccount and oSavingsAccount to the list using the Add method of the list.

```
AccountList.Add(oCheckingAccount);
AccountList.Add(oSavingsAccount);
```

5. Loop through the list and call the GetAccountInfo method of each account type in the list and show the results in a console window.

```
foreach (Account a in AccountList)
{
    Console.WriteLine(a.GetAccountInfo());
}
Console.ReadLine();
```

6. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
7. Select Debug ► Start to run the project. You should see a console window with the return string for the GetAccountInfo method of each object in the list.
8. After testing the polymorphism, hit the enter key to close the console window, which will stop the debugger.

Implementing Polymorphism Using an Interface

To implement polymorphism using an interface, follow these steps:

1. View the code for the Account.cs file in the code editor.
2. Comment out the code for the Account, CheckingAccount, and SavingsAccount classes.

3. Define an interface `IAccount` that contains the `GetAccountInfo` method.

```
public interface IAccount
{
    string GetAccountInfo();
}
```

4. Add the following code to create two classes: `CheckingAccount` and `SavingsAccount`. These classes will implement the `IAccount` interface.

```
public class CheckingAccount : IAccount
{
    private int _accountNumber;
    public int AccountNumber
    {
        get { return _accountNumber; }
        set { _accountNumber = value; }
    }
    public string GetAccountInfo()
    {
        return "Printing checking account info for account number "
            + AccountNumber.ToString();
    }
}

public class SavingsAccount : IAccount
{
    private int _accountNumber;
    public int AccountNumber
    {
        get { return _accountNumber; }
        set { _accountNumber = value; }
    }
    public string GetAccountInfo()
    {
        return "Printing savings account info for account number "
            + AccountNumber.ToString();
    }
}
```

5. Select **Build** ► **Build Solution**. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.

Testing the Polymorphic Interface Method

To test the polymorphic method, follow these steps:

1. Open the Program.cs file in the code editor and locate the Main method.
2. Change the code to create an instance of a list of IAccount types.

```
List<IAccount>AccountList = new List<IAccount>();
```

3. Change the code for each loop to loop through the list and call the GetAccountInfo method of each IAccount type in the list.

```
foreach (IAccount a in AccountList)
{
    Console.WriteLine(a.GetAccountInfo());
}
Console.ReadLine();
```

4. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
 5. Select Debug ► Start to run the project. You should see a console window with the return string for the GetAccountInfo method of each object in the list.
 6. After testing the polymorphism, hit the enter key to close the console window, which will stop the debugger.
-

Summary

This chapter introduced you to two of OOP's most powerful features: inheritance and polymorphism. Knowing how to implement these features is fundamental to becoming a successful object-oriented programmer, regardless of the language you use.

In Chapter 8, you will take a closer look at how the objects in your applications collaborate. The topics covered include how objects pass messages to one another, how events drive your programs, how data is shared among instances of a class, and how exceptions are handled.



Implementing Object Collaboration

In Chapter 7, you learned how to create and use class hierarchies in C#. That chapter also introduced the concepts of inheritance, polymorphism, and interfaces. In this chapter, you'll learn how to get the objects of an application to work together to perform tasks. You will see how objects communicate through messaging and how events initiate application processing. You'll also learn how the objects respond and communicate exceptions that may occur as they carry out their assigned tasks.

After reading this chapter, you should be familiar with the following:

- the process of object communication through messaging
- the different types of messaging that can occur
- how to use delegation in C# applications
- how objects can respond to events and publish their own events
- the process of issuing and responding to exceptions
- how to create shared data and procedures among several instances of the same class
- how to issue message calls asynchronously

Communicating Through Messaging

One of the advantages of OOP is that many aspects of it mimic the real world. Previously, you saw how you can use the analogy of employees in a company representing instances of an `Employee` class. We can extend this analogy to how instances of classes (objects) interact in an application. For example, in large companies, the employees perform specialized functions. One person is in charge of accounts payable processing, and another is responsible for the accounts receivable operations. When an employee needs to request a service—paid time off (PTO), for example—the employee sends a message to her manager. We can think of this interaction as a service request between the requester (employee) and a service provider (manager). This request can involve just a single object (a self-served request), two objects, or it can be a complex chain of requester/service provider requests. For example, the employee requests the PTO from her manager, who, in turn, checks with the human resources (HR) department to see if the employee has enough accumulated time. In this case, the manager is both a service provider to the employee and a service requester to the HR department.

Defining Method Signatures

When a message passes between a requester and a service provider, the requester may or may not expect a response. For example, when an employee requests PTO, she expects a response indicating approval or denial. However,

when the accounting department issues paychecks, the staff members do not expect everyone in the company to issue a response e-mail thanking them!

A common requirement when a message is issued is to include the information necessary to carry out the request. When an employee requests PTO, her manager expects her to provide him with the dates off she is requesting. In OOP terminology, you refer to the name of the method (requested service) and the input parameters (requester-supplied information) as the method signature.

The following code demonstrates how methods are defined in C#. The access modifier is first followed by the return type (void is used if no return value is returned) and then the name of the method. Parameter types and names are listed in parenthesis separated by commas. The body of the method is contained in opening and closing curly brackets.

```
public int AddEmployee(string firstName,string lastName)
{
    //Code to save data to database
}
public void LogMessage(string message)
{
    //Code to write to log file.
}
```

Passing Parameters

When you define a method in the class, you also must indicate how the parameters are passed. Parameters may be passed by value or by reference.

If you choose to pass the parameters by value, a copy of the parameter data is passed from the calling routine to the requested method. The requested method works with the copy and, if changes are made to the data, the requested method must pass the copy back to the calling routine so that the caller can choose to discard the changes or replicate them. Returning to the company analogy, think about the process of updating your employee file. The HR department does not give you direct access to the file; instead, it sends you a copy of the values in the file. You make changes to the copy, and then you send it back to the HR department. The HR department then decides whether to replicate these changes to the actual employee file. In C#, passing parameters by value is the default, so no keyword is used. In the following method, the parameter is passed by value:

```
public int AddEmployee(string firstName)
{
    //Code to save data to database
}
```

Another way you can pass parameters is by reference. In this case, the requesting code does not pass in a copy of the data, but instead passes a reference to where the data is located. Using the previous example, instead of sending you a copy of the data in your employee file when you want to make updates, the HR department informs you where the file is located, and tells you to go to it to make the changes. In this case, clearly, it would be better to pass the parameters by reference. In C# code, when passing parameters by reference, the `ref` keyword is used. The following code shows how to define the method to pass values by reference:

```
public int AddEmployee(ref string firstName)
{
    //Code to save data to database
}
```

When applications are designed as components that communicate across processing boundaries, and that are even hosted on different computers, it is advantageous to pass parameters by value instead of by reference. Passing parameters by reference can cause increased overhead, because when the service-providing object must work with parameter information, it needs to make calls across the processing boundaries and the network. A single processing request can result in many calls back and forth between the requester and the server provider. Passing values by reference can also cause problems when maintaining data integrity. The requester is opening a channel for the data to be manipulated without the server's knowledge or control.

On the other hand, passing values by reference may be the better choice when the calling code and the requested method are in the same processing space (they occupy the same "cubicle," so to speak) and have a clearly established trust relationship. In this situation, allowing direct access to the memory storage location and passing the parameters by reference may offer a performance advantage over passing the parameters by value.

The other situation in which passing parameters by reference may be advantageous is if the object is a complex data type, such as another object. In this case, the overhead of copying the data structure and passing it across process and network boundaries outweighs the overhead of making repeated calls across the network.

■ **Note** The .NET Framework addresses the problem of complex data types by allowing you to efficiently copy and pass those types by serializing and deserializing them in an XML structure.

Understanding Event-Driven Programming

So far, you have been looking at messaging between the objects in which a direct request initiates the message interaction. If you think about how you interact with objects in real life, you often receive messages in response to an event that has occurred. To continue the office analogy, for example, when the sandwich vendor comes into the building, a message is issued over the intercom informing employees that the lunch truck has arrived. This type of messaging is referred to as broadcast messaging. A message is issued, and the receiver decides whether to ignore or respond to the message.

Another way this event message could be issued is by the receptionist sending an e-mail to a list of employees who are interested in knowing when the sandwich vendor shows up. In this case, the interested employees would subscribe with the receptionist to receive the event message. This type of messaging is often referred to as subscription-based messaging.

Applications built with the .NET Framework are object-oriented, event-driven programs. If you trace the request/service provider processing chains that occur in your applications, you can identify the event that kicked off the processing. In the case of Windows applications, the user interacting with a GUI usually initiates the event. For example, a user might initiate the process of saving data to a database by clicking a button. Classes in applications can also initiate events. A security class could broadcast an event message when an invalid login is detected. You can also subscribe to external events. You could create a Web service that would issue an event notification when a change occurs in a stock that you are tracking in the stock market. You could write an application that subscribes to the service and responds to the event notification.

Understanding Delegation

In order to implement event-based programming in C#, you must first understand delegation. Delegation is when you request a service by making a method call to a service provider. The service provider then reroutes this service request to another method, which services the request. The delegate class can examine the service request and dynamically determines at runtime where to route the request. Returning to the company analogy, when a manager receives a service request, she often delegates it to a member of her department. (In fact, many would argue that a common trait among successful managers is the ability to know when and how to delegate responsibilities.)

When you create a delegated method, you first define the delegated method's signature. Because the delegated method does not actually service the request, it does not contain any implementation code. The following code shows a delegated method used to compare integer values:

```
public delegate Boolean CompareInt(int I1, int I2);
```

Once the delegated method's signature is defined, you can then create the methods to which to delegate. These methods must have the same parameters and return types as the delegated method. The following code shows two methods to which the delegated method will delegate:

```
private Boolean AscendOrder(int I1, int I2)
{
    if (I1 < I2)
        { return true;}
    else
        { return false; }
}
private Boolean DescendOrder(int I1, int I2)
{
    if (I1 > I2)
        { return true; }
    else
        { return false; }
}
```

Once the delegate and its delegating methods have been defined, you are ready to use the delegate. The following code shows a portion of a sorting routine that determines which delegated method to call depending on a `SortType` passed in as a parameter:

```
public void SortIntegers(SortType sortDirection, int[] intArray)
{
    CompareInt CheckOrder;
    if (sortDirection == SortType.Ascending)
        { CheckOrder = new CompareInt(AscendOrder); }
    else
        { CheckOrder = new CompareInt(DescendOrder); }
    // Code continues ...
}
```

Implementing Events

In C#, when you want to issue event messages, first you declare a delegate type for the event. The delegate type defines the set of arguments that will be passed to the method that handles the event.

```
public delegate void DataUpdateEventHandler(string msg);
```

Once the delegate is declared, an event of the delegate type is declared.

```
public event DataUpdateEventHandler DataUpdate;
```

When you want to raise the event, you call the event passing in the appropriate arguments.

```
public void SaveInfo()
{
    try
    {
        DataUpdate("Data has been updated");
    }
    catch
    {
        DataUpdate("Data could not be updated");
    }
}
```

Responding To Events

To consume an event in client code, an event handling method is declared that executes program logic in response to the event. This event handler must have the same method signature as the event delegate declared in the class issuing the event.

```
void odata_DataUpdate(string msg)
{
    MessageBox.Show(msg);
}
```

This event handler is registered with the event source using the += operator. This process is referred to as event wiring. The following code wires up the event handler for the DataUpdate event declared previously:

```
Data odata = new Data();
odata.DataUpdate += new DataUpdateEventHandler(odata_DataUpdate);
odata.SaveInfo();
```

Windows Control Event Handling

Windows Forms also implement event handlers by using the += operator to wire up the event handler to the event. The following code wires up a button to a click event and a textbox to a mouse down event:

```
this.button1.Click += new System.EventHandler(this.button1_Click);
this.textBox1.MouseDown += new System.Windows.Forms.MouseEventHandler(this.textBox1_MouseDown);
```

The event handler methods for control events take two parameters: the first parameter, sender, provides a reference to the object that raised the event. The second parameter passes an object containing information specific to the event that is being handled. The following code shows an event handler method for a button click event and an event handler for the textbox mouse down event. Notice how e is used to determine if the left button was clicked.

```
private void button1_Click(object sender, EventArgs e)
{
}

private void textBox1_MouseDown(object sender, MouseEventArgs e)
```

```

{
    if (e.Button == System.Windows.Forms.MouseButtons.Left)
    {
        //code goes here.
    }
}

```

ACTIVITY 8-1. ISSUING AND RESPONDING TO EVENT MESSAGES

In this activity, you will learn to do the following:

- create and raise events from a server class
- handle events from client classes
- handle GUI events

Adding and Raising Event Messaging in the Class Definition

To add and raise event messaging in a class definition file, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose a Windows Forms Application project. Name the project Activity8_1.
3. A default form is included in the project. Add controls to the form and change the property values, as listed in Table 8-1. Your completed form should look similar to Figure 8-1.

Table 8-1. Login Form and Control Properties

Object	Property	Value
Form1	Name	frmLogin
	Text	Login
Label1	Name	lblName
	Text	Name:
Label2	Name	lblPassword
	Text	Password:
Textbox1	Name	txtName
	Text	(empty)
Textbox2	Name	txtPassword
	Text	(empty)
	PasswordChar	*
Button1	Name	btnLogin
	Text	Login
Button2	Name	btnClose
	Text	Close

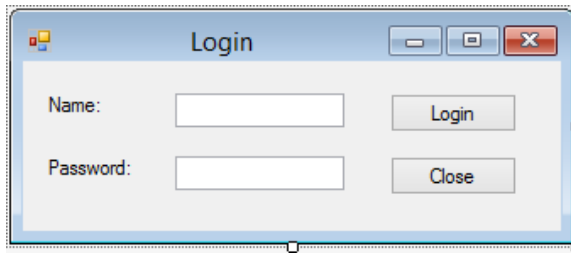


Figure 8-1. The completed login form

4. Select Project ► Add Class. Name the class Employee. Open the Employee class code in the code editor.
5. Above the class declaration, add the following line of code to define the login event handler delegate. You will use this event to track employee logins to your application.

```
public delegate void LoginEventHandler(string loginName, Boolean status);
```

6. Inside the class declaration, add the following line of code to define the LoginEvent as the delegate type:

```
public event LoginEventHandler LoginEvent;
```

7. Add the following Login method to the class, which will raise the LoginEvent:

```
public void Login(string loginName, string password)
{
    //Data normally retrieved from database.
    if (loginName == "Smith" && password == "js")
    {
        LoginEvent(loginName, true);
    }
    else
    {
        LoginEvent(loginName, false);
    }
}
```

8. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.

Receiving Events in the Client Class

To receive events in the client class, follow these steps:

1. Open frmLogin in the design window.
2. Double-click the Login button to view the Login button click event handler.

3. Add the following code to wire up the Employee class's LoginEvent with an event handler in the form class:

```
private void btnLogin_Click(object sender, EventArgs e)
{
    Employee oEmployee = new Employee();
    oEmployee.LoginEvent += new LoginEventHandler(oEmployee_LoginEvent);
    oEmployee.Login(txtName.Text, txtPassword.Text);
}
```

4. Add the following event handler method to the form that gets called when the Employee class issues a LoginEvent:

```
void oEmployee_LoginEvent(string loginName, bool status)
{
    MessageBox.Show("Login status :" + status);
}
```

5. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
6. Select Debug ► Start to run the project.
7. To test to make sure the Login event is raised, enter a login name of Smith and a password of js. This should trigger a login status of true.
8. After testing the Login event, close the form, which will stop the debugger.

Handling Multiple Events with One Method

To handle multiple events with one method, follow these steps:

1. Open frmLogin in the form designer by right-clicking the frmLogin node in the Solution Explorer and choosing View Designer.
2. From the Toolbox, add a MenuStrip control to the form. Click where it says “Type Here” and enter File for the top-level menu and Exit for its submenu, as shown in Figure 8-2.




Figure 8-2. Adding the MenuStrip control

3. Add the following method to handle the click event of the menu and the Close button:

```
private void FormClose(object sender, EventArgs e)
{
    this.Close();
}
```

4. Open frmLogin in the designer window. In the properties window, select the exitToolStripMenuItem. Select the event button (lightning bolt) at the top of the properties window to show the events of the control. In the click event drop-down, select the FormClose method (see Figure 8-3).

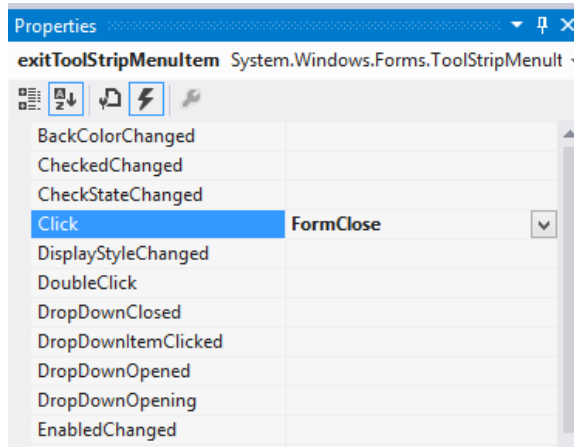


Figure 8-3. Wiring up an event handler

5. Repeat step 4 to wire up the btnClose button click event to the FormClose method.
6. Expand the Form1.cs node in the Solution window. Right click on the Form1.Designer.cs node and select View Code.
7. In the code editor, expand the Windows Form Designer generated code region. Search for the code listed below. This code was generated by the form designer to wire up the events to the FormClose method.

```
this.btnClose.Click += new System.EventHandler(this.FormClose);
this.exitToolStripMenuItem.Click += new System.EventHandler(this.FormClose);
```

8. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
 9. Select Debug ► Start to run the project. Test the Exit menu and the Close button.
 10. After testing, save the project, and then exit Visual Studio.
-

Handling Exceptions in the .NET Framework

When objects collaborate, things can go wrong. Exceptions are things that you do not expect to occur during normal processing. For example, you may be trying to save data to a database over the network when the connection fails, or you may be trying to save to a drive without a disk in the drive. Your applications should be able to gracefully handle any exceptions that occur during application processing.

The .NET Framework uses a structured exception handling mechanism. The following are some of the benefits of this structured exception handling:

- common support and structure across all .NET languages
- support for the creation of protected blocks of code
- the ability to filter exceptions to create efficient robust error handling
- support of termination handlers to guarantee that cleanup tasks are completed, regardless of any exceptions that may be encountered

The .NET Framework also provides an extensive number of exception classes used to handle common exceptions that might occur. For example, the `FileNotFoundException` class encapsulates information such as the file name, error message, and the source for an exception that is thrown when there is an attempt to access a file that does not exist. In addition, the .NET Framework allows the creation of application-specific exception classes you can write to handle common exceptions that are unique to your application.

Using the Try-Catch Block

When creating code that could end up causing an exception, you should place it in a try block. Code placed inside the try block is considered protected. If an exception occurs while the protected code is executing, code processing is transferred to the catch block, where it is handled. The following code shows a method of a class that tries to read from a file that does not exist. When the exception is thrown, it is caught in the catch block.

```
public string ReadText(string filePath)
{
    StreamReader sr;
    try
    {
        sr = File.OpenText(filePath);
        string fileText = sr.ReadToEnd();
        sr.Close();
        return fileText;
    }
    catch(Exception ex)
    {
        return ex.Message;
    }
}
```

All try blocks require at least one nested catch block. You can use the catch block to catch all exceptions that may occur in the try block, or you can use it to filter exceptions based on the type of exception. This enables you to dynamically respond to different exceptions based on the exception type. The following code demonstrates filtering exceptions based on the different exceptions that could occur when trying to read a text file from disk:

```
public string ReadText(string filePath)
{
    StreamReader sr;
    try
    {
        sr = File.OpenText(filePath);
        string fileText = sr.ReadToEnd();
        sr.Close();
        return fileText;
    }
    catch (DirectoryNotFoundException ex)
    {
        return ex.Message;
    }
    catch (FileNotFoundException ex)
    {
        return ex.Message;
    }
    catch(Exception ex)
    {
        return ex.Message;
    }
}
```

Adding a Finally Block

Additionally, you can nest a finally block at the end of the try block. Unlike the catch block, the use of the finally block is optional. The finally block is for any cleanup code that needs to occur, even if an exception is encountered. For example, you may need to close a database connection or release a file. When the code of the try block is executed and an exception occurs, processing will evaluate each catch block until it finds the appropriate catch condition. After the catch block executes, the finally block will execute. If the try block executes and no exception is encountered, the catch blocks don't execute, but the finally block will still get processed. The following code shows a finally block being used to close and dispose a StreamReader:

```
public string ReadText(string filePath)
{
    StreamReader sr = null;
    try
    {
        sr = File.OpenText(filePath);
        string fileText = sr.ReadToEnd();
        return fileText;
    }
}
```

```

    catch (DirectoryNotFoundException ex)
    {
        return ex.Message;
    }
    catch (FileNotFoundException ex)
    {
        return ex.Message;
    }
    catch (Exception ex)
    {
        return ex.Message;
    }
    finally
    {
        if (sr != null)
        {
            sr.Close();
            sr.Dispose();
        }
    }
}

```

Throwing Exceptions

During code execution, when an inappropriate call is made, for example, a parameter to a method has an invalid value or an parameter passed to a method causes an exception, you can throw an exception to inform the calling code of the violation. In the following code, if the `orderDate` parameter is greater than the current date, an `ArgumentOutOfRangeException` is thrown back to the calling code informing them of the violation.

```

if (orderDate > DateTime.Now)
{
    throw new ArgumentOutOfRangeException ("Order date can not be in the future.");
}
//Processing code...

```

■ **Note** If you cannot find an appropriate predefined exception class in the .NET Framework, you can create a custom exception class that derives from the `System.Exception` class.

Nesting Exception Handling

In some cases, you may be able to correct an exception that occurred and continue processing the rest of the code in the `try` block. For example, a division-by-zero error may occur, and it would be acceptable to assign the result a value of zero and continue processing. In this case, a `try-catch` block could be nested around the line of code that would

cause the exception. After the exception is handled, processing would return to the line of code in the outer try-catch block immediately after the nested try block. The following code demonstrates nesting one try block within another:

```
try
{
    try
    {
        Y = X1 / X2;
    }
    catch (DivideByZeroException ex)
    {
        Y = 0;
    }
    //Rest of processing code.
}
catch (Exception ex)
{
    //Outer exception processing
}
```

■ **Note** For more information about handling exceptions and the .NET Framework exception classes, refer to Appendix B.

Static Properties and Methods

When you declare an object instance of a class, the object instantiates its own instances of the properties and methods of the class it implements. For example, if you were to write a counting routine that increments a counter, then instantiated two object instances of the class, the counters of each object would be independent of each other; when you incremented one counter, the other would not be affected. Normally, this object independence is the behavior you want. However, sometimes you may want different object instances of a class to access the same, shared variables. For example, you might want to build in a counter that logs how many of the object instances have been instantiated. In this case, you would create a static property value in the class definition. The following code demonstrates how you create a static `TaxRate` property in a class definition:

```
public class AccountingUtilities
{
    private static double _taxRate = 0.06;

    public static double TaxRate
    {
        get { return _taxRate; }
    }
}
```

To access the static property, you don't create an object instance of the class; instead, you refer to the class directly. The following code shows a client accessing the static `TaxRate` property defined previously:

```
public class Purchase
{
    public double CalculateTax(double purchasePrice)
    {
        return purchasePrice * AccountingUtilities.TaxRate;
    }
}
```

Static methods are useful if you have utility functions that clients need to access, but you don't want the overhead of creating an object instance of a class to gain access to the method. Note that static methods can access only static properties. The following code shows a static method used to count the number of users currently logged in to an application:

```
public class UserLog
{
    private static int _userCount;
    public static void IncrementUserCount()
    {
        _userCount += 1;
    }
    public static void DecrementUserCount()
    {
        _userCount -= 1;
    }
}
```

When client code accesses a static method, it does so by referencing the class directly. The following code demonstrates accessing the static method defined previously:

```
public class User
{
    //other code ...
    public void Login(string userName, string password)
    {
        //code to check credentials
        //if successful
        UserLog.IncrementUserCount();
    }
}
```

Although you may not use static properties and methods often when creating the classes in your applications, they are useful when creating base class libraries and are used throughout the .NET Framework system classes. The following code demonstrates the use of the `Compare` method of the `System.String` class. This is a static method that compares two strings alphabetically. It returns a positive value if the first string is greater, a negative value if the second string is greater, or zero if the strings are equal.

```
public Boolean CheckStringOrder(string string1, string string2)
{
    if (string.Compare(string1, string2) >= 0)
```



```

{
    return true;
}
else
{
    return false;
}
}

```

ACTIVITY 8-2. IMPLEMENTING EXCEPTION HANDLING AND STATIC METHODS

In this activity, you will learn how to do the following:

- create and call static methods of a class
- use structured exception handling

Creating Static Methods

To create the static methods, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose a Windows Application project. Name the project Activity8_2.
3. Visual Studio creates a default form for the project which you'll use to create a login form named Logger. Add controls to the form and change the property values, as listed in Table 8-2. Your completed form should look similar to Figure 8-4.

Table 8-2. *Logger Form and Control Properties*

Object	Property	Value
Form1	Name	frmLogger
	Text	Logger
Textbox1	Name	txtLogPath
	Text	c:\Test\LogTest.txt
Textbox2	Name	txtLogInfo
	Text	Test Message
Button1	Name	btnLogInfo
	Text	Log Info

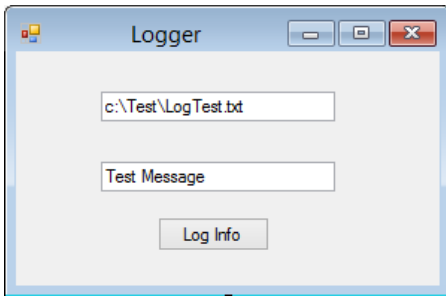


Figure 8-4. The completed logger form

4. Select Project ► Add Class. Name the class `Logger`.
5. Because you will be using the `System.IO` class within the `Logger` class, add a using statement to the top of the class file:

```
using System.IO;
```

6. Add a static `LogWrite` method to the class. This method will write information to a log file. To open the file, create a `FileStream` object. Then create a `StreamWriter` object to write the information to the file. Notice the use of the `using` blocks to properly dispose the `FileStream` and `StreamWriter` objects and release the resources.

```
public static string LogWrite(string logPath, string logInfo)
{
    using (FileStream oFileStream = new FileStream(logPath, FileMode.Open,
        FileAccess.Write))
    {
        using (StreamWriter oStreamWriter = new StreamWriter(oFileStream))
        {
            oFileStream.Seek(0, SeekOrigin.End);
            oStreamWriter.WriteLine(DateTime.Now);
            oStreamWriter.WriteLine(logInfo);
            oStreamWriter.WriteLine();
        }
    }
    return "Info Logged";
}
```

7. Open `frmLogger` in the visual design editor. Double click the `btnLogInfo` button to bring up the `btnLogInfo_Click` event method in the code editor. Add the following code, which runs the `LogWrite` method of the `Logger` class and displays the results in the form's text property. Note that because you designated the `LogWrite` method as static (in step 6), the client does not need to create an object instance of the `Logger` class. Static methods are accessed directly through a class reference.

```
private void btnLogInfo_Click(object sender, EventArgs e)
{
    this.Text = Logger.LogWrite(txtLogPath.Text, txtLogInfo.Text);
}
```

8. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
9. Select Debug ► Run. When the form launches, click the Log Info button. You should get an unhandled exception message of type `System.IO.DirectoryNotFoundException`. Stop the debugger.

Creating the Structured Exception Handler

To create the structured exception handler, follow these steps:

1. Open the `Logger` class code in the code editor.
2. Locate the `LogWrite` method and add a try-catch block around the current code. In the catch block, return a string stating the logging failed.

```
try
{
    using (FileStream oFileStream = new FileStream(logPath, FileMode.Open,
        FileAccess.Write))
    {
        using (StreamWriter oStreamWriter = new StreamWriter(oFileStream))
        {
            oFileStream.Seek(0, SeekOrigin.End);
            oStreamWriter.WriteLine(DateTime.Now);
            oStreamWriter.WriteLine(logInfo);
            oStreamWriter.WriteLine();
        }
    }
    return "Info Logged";}
catch
{
    return "Logging Failed";
}
```

3. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
4. Select Debug ► Run. When the form launches, click the Log Info button. This time, you should not get the exception message because it was handled by the `LogWrite` method. You should see the message “Logging Failed” in the form’s caption. Close the form.

Filtering Exceptions

To filter exceptions, follow these steps:

1. Alter the catch block to return different messages, depending on which exception is thrown.

```
catch (FileNotFoundException ex)
{
    return ex.Message;
}
catch (IOException ex)
{
    return ex.Message;
}
catch
{
    return "Logging Failed";
}
```

2. Set a break point on the `LogWrite` method of the `Logger` class.
 3. Select **Debug** ► Start to run the project. Test the catch block by clicking the **Log Info** button. Execution will stop at the break point. Step through the code and notice it gets caught by the `IOException` block.
 4. After testing, close the form.
 5. Using Notepad, create the `LogTest.txt` file in a `Test` folder on the C drive and close the file. Make sure the file and folder are not marked as read only.
 6. Select **Debug** ► Start to run the project. Test the `WriteLog` method by clicking the **Log Info** button. This time, the form's caption should indicate that the log write was successful.
 7. Stop the debugger.
 8. Open the `LogTest.txt` file using Notepad and verify that the information was logged.
 9. Save the project, and then exit Visual Studio.
-

Using Asynchronous Messaging

When objects interact by passing messages back and forth, they can pass the message synchronously or asynchronously.

When a client object makes a synchronous message call to a server object, the client suspends processing and waits for a response back from the server before continuing. Synchronous messaging is the easiest to implement and is the default type of messaging implemented in the .NET Framework. However, sometimes this is an inefficient way of passing messages. For example, the synchronous messaging model is not well suited for long-running file reading and writing, making service calls across slow networks, or message queuing in disconnected client scenarios. To more effectively handle these types of situations, the .NET Framework provides the plumbing needed to pass messages between objects asynchronously.

When a client object passes a message asynchronously, the client can continue processing. After the server completes the message request, the response information will be sent back to the client.

If you think about it, you interact with objects in the real world both synchronously and asynchronously. A good example of synchronous messaging is when you are in the checkout line at the grocery store. When the clerk can't determine the price of one of the items, he calls the manager for a price check and suspends the checkout process until a result is returned. An example of an asynchronous message call is when the clerk notices that he is running low on change. He alerts the manager that he will need change soon, but he can continue to process his customer's items until the change arrives.

In an effort to make asynchronous programming easier to implement, the .NET Framework 4.5 has introduced the Task-Based Asynchronous Pattern (TAP). When conforming to TAP, asynchronous methods return a task object. This task object represents the ongoing operation and allows communication back to the caller. When calling the asynchronous method, the client simply uses the `await` modifier and the compiler takes care of all the necessary plumbing code.

In the .NET Framework, when you want to create a method that can be called asynchronously, you use the `async` modifier. An `async` method provides a way to perform potentially long-running processes without blocking the caller. This is very useful if the main thread containing the User Interface (UI) needs to call a long running process. If the process is called synchronously, then the UI will freeze until the process completes.

An `async` method should contain at least one `await` statement. When the `await` statement is encountered, processing is suspended and control is returned to the caller (UI in this case). So the user can continue to interact with the UI. Once the `async` task is complete, processing is returned back to the `await` statement and the caller can be alerted that the processing is finished. In order to alert the caller that the process is complete, you create a `Task` return type. If the `async` method needs to pass information back to the caller a `Task<TResult>` is used. The following code shows a method that reads a text file asynchronously. Notice the `async` modifier and the `await` keyword used to call the `ReadToEndAsync` method of the `StreamReader` class. The method passes back a `Task` object to the caller with a string result.

```
public static async Task<string> LogReadAsync(string filePath)
{
    StreamReader oStreamReader;
    string fileText;
    try
    {
        oStreamReader = File.OpenText(filePath);
        fileText = await oStreamReader.ReadToEndAsync();
        oStreamReader.Close();
        return fileText;
    }
    catch (FileNotFoundException ex)
    {
        return ex.Message;
    }
    catch (IOException ex)
    {
        return ex.Message;
    }
    catch
    {
        return "Logging Failed";
    }
}
```

ACTIVITY 8-3. CALLING METHODS ASYNCHRONOUSLY

In this activity, you will learn how to do the following:

- call methods synchronously
- call methods asynchronously

Creating a Method and Calling It Synchronously

To create the method and call it synchronously, follow these steps:

1. Start Visual Studio. Select File ► Open ► Project.
2. Open the solution file you completed in Activity8_2.
3. Add the buttons shown in Table 8-3 to the frmLogger form. Figure 8-5 shows the completed form.

Table 8-3. Additional Buttons for the Logger Form

Object	Property	Value
Button1	Name	btnSyncRead
	Text	Sync Read
Button2	Name	btnAsyncRead
	Text	Async Read
Button3	Name	btnMessage
	Text	Message

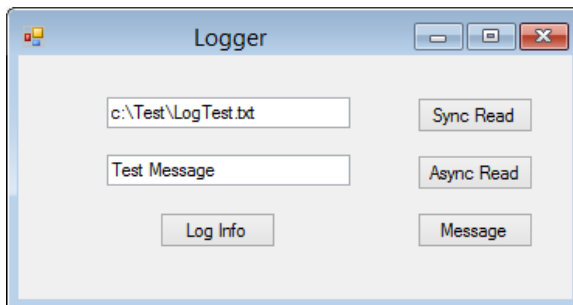


Figure 8-5. The completed logger form for synchronous and asynchronous reading

4. Open the Logger class in the code editor.
5. Recall that because you are using the System.IO namespace within the Logger class, you added a using statement to the top of the file. You are also going to use System.Threading namespace, so add a using statement to include this namespace.

```
using System.Threading;
```

6. Add a static `LogRead` function to the class. This function will read information from a log file. To open the file, create a `FileStream` object. Then create `StreamReader` object to read the information from the file. You are also using the `Thread` class to suspend processing for five seconds to simulate a long call across a slow network.

```
public static string LogRead(string filePath)
{
    StreamReader oStreamReader;
    string fileText;
    try
    {
        oStreamReader = File.OpenText(filePath);
        fileText = oStreamReader.ReadToEnd();
        oStreamReader.Close();
        Thread.Sleep(5000);
        return fileText;
    }
    catch (FileNotFoundException ex)
    {
        return ex.Message;
    }
    catch (IOException ex)
    {
        return ex.Message;
    }
    catch
    {
        return "Logging Failed";
    }
}
```

7. Open `frmLogger` in the visual design editor. Double click the `btnMessage` button to bring up the `btnMessage_Click` event method in the code editor. Add code to display a message box.

```
private void btnMessage_Click(object sender, EventArgs e)
{
    MessageBox.Show("Hello");
}
```

8. Open `frmLogger` in the visual design editor. Double-click the `btnSyncRead` button to bring up the `btnSyncRead_Click` event method in the code editor. Add code that calls the `LogRead` method of the `Logger` class and displays the results in a message box.

```
private void btnSyncRead_Click(object sender, EventArgs e)
{
    MessageBox.Show(Logger.LogRead(txtLogPath.Text));
}
```

9. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
10. Select Debug ► Run. When the form launches, click the Sync Read button. After clicking the Sync Read button, try clicking the Message button. You should not get a response when clicking the Message button because you called the ReadLog method synchronously. After the ReadLog method returns a result, the Message button will respond when clicked.
11. When you have finished testing, close the form.

Creating and Calling an Asynchronous Method

To create an asynchronous method, follow these steps:

1. Open the Logger class code in the code editor.
2. Check for the following using statement at the top of the file. This namespace exposes the Task class and other types that are used to implement asynchronous programming.

```
using System.Threading.Tasks;
```

3. Create an asynchronous method that reads the text file. The use of the Task's Delay method is to simulate a long running process.

```
public static async Task<string> LogReadAsync(string filePath)
{
    string fileText;
    try
    {
        using (StreamReader oStreamReader = File.OpenText(filePath))
        {
            fileText = await oStreamReader.ReadToEndAsync();
        }
        await Task.Delay(10000);
        return fileText;
    }
    catch (FileNotFoundException ex)
    {
        return ex.Message;
    }
    catch (IOException ex)
    {
        return ex.Message;
    }
    catch
    {
        return "Logging Failed";
    }
}
```


4. Open frmLogger in the visual design editor. Double-click the btnAsyncRead button to bring up the btnAsyncRead_Click event method in the code editor. Alter the method so that it is asynchronous.

```
private async void btnAsyncRead_Click(object sender, EventArgs e)
```

5. Add code to call the LogReadAsync method of the Logger class and display the results in a message box.

```
btnAsyncRead.Enabled = false;
string s = await Logger.LogReadAsync(txtLogPath.Text);
MessageBox.Show(s);
btnAsyncRead.Enabled = true;
```

6. Select Build ► Build Solution. Make sure there are no build errors in the Error List window. If there are, fix them, and then rebuild.
 7. Select Debug ► Run. When the form launches, click the Async Read button. After clicking the Async Read button, click the Message button. This time, you should get a response because you called the ReadLog method asynchronously. After five seconds you should see a message box containing the results of the Logger.LogReadAsync method.
 8. When you have finished testing, close the form.
 9. Save the project, and then exit Visual Studio.
-

Summary

This chapter described how the objects in your applications collaborate. You saw how objects pass messages to one another, how events drive your programs, how instances of a class share data, and how to handle exceptions.

In Chapter 9, we look at collections and arrays. Collections and arrays organize similar objects into a group. Working with collections is one of the most common programming constructs you will need to apply in your applications. You will examine some of the basic types of collections available in the .NET Framework and learn how to employ collections in your code.



Working with Collections

In the previous chapter, you looked at how objects collaborate and communicate in object-oriented programs. That chapter introduced the concepts of messaging, events, delegation, exception handling, and asynchronous programming. In this chapter, you will look at how collections of objects are organized and processed. The .NET Framework contains an extensive set of classes and interfaces for creating and managing collections of objects. You will look at the various types of collection structures .NET provides and learn what they are designed for and when to use each. You will also look at how to use generics to create highly reusable, efficient collections.

In this chapter, you will learn the following:

- the various types of collections exposed by the .NET Framework
- how to work with arrays and array lists
- how to create generic collections
- how to implement queues and stacks

Introducing the .NET Framework Collection Types

Programmers frequently need to work with collections of types. For example, if you are working with employee time records in a payroll system, you need to group the records by employee, loop through the records, and add up the hours for each.

All collections need a basic set of functions, such as adding objects, removing objects, and iterating through their objects. In addition to the basic set, some collections need additional specialized functions. For example, a collection of help desk e-mail requests may need to implement a first-in, first-out functionality when adding and removing items from the collection. On the other hand, if the requests were prioritized by severity, the collection would need the ability to sort its items by priority.

The .NET Framework provides a variety of basic and specialized collection classes for you to use. The `System.Collections` namespace contains interfaces and classes that define various types of collections, such as lists, queues, hash tables, and dictionaries. Table 9-1 lists and describes some of the commonly used collection classes. If you do not find a collection class with the functionality you need, you can extend a .NET Framework class to create your own.

Table 9-1. *Commonly Used Collection Classes*

Class	Description
Array	Provides the base class for language implementations that support strongly typed arrays.
ArrayList	Represents a weakly typed list of objects using an array whose size is dynamically increased as required.
SortedList	Represents a collection of key/value pairs that are sorted by the keys and are accessible by key and by index.
Queue	Represents a first-in, first-out (FIFO) collection of objects.
Stack	Represents a simple last-in, first-out (LIFO), nongeneric collection of objects.
Hashtable	Represents a collection of key/value pairs that are organized based on the hash code of the key.
CollectionBase	Provides the abstract base class for a strongly typed collection.
DictionaryBase	Provides the abstract base class for a strongly typed collection of key/value pairs.

Table 9-2 describes some of the interfaces implemented by these collection classes.

Table 9-2. *Collection Class Interfaces*

Interface	Description
ICollection	Defines size, enumerators, and synchronization methods for all nongeneric collections.
IComparer	Exposes a method that compares two objects.
IDictionary	Represents a nongeneric collection of key/value pairs.
IDictionaryEnumerator	Enumerates the elements of a nongeneric dictionary.
IEnumerable	Exposes the enumerator, which supports a simple iteration over a nongeneric collection.
IEnumerator	Supports a simple iteration over a nongeneric collection.
IList	Represents a nongeneric collection of objects that can be individually accessed by index.

In this chapter, you will work with some of the commonly used collection classes, beginning with the Array and ArrayList classes.

Working with Arrays and Array Lists

An array is one of the most common data structures in computer programming. An array holds data elements of the same data type. For example, you can create an array of integers, strings, or dates. Arrays are often used to pass values to methods as parameters. For example, when you use a console application, it's common to provide command line switches. The following DOS command is used to copy a file on your computer:

```
copy win.ini c:\windows /y
```

The source file, destination path, and overwrite indicator are passed into the copy program as an array of strings.

You access the elements of an array through its index. The index is an integer representing the position of the element in the array. For example, an array of strings representing the days of the week has the following index values:

Index	Value
0	Sunday
1	Monday
2	Tuesday
3	Wednesday
4	Thursday
5	Friday
6	Saturday

This days-of-the-week example is a one-dimensional array, which means the index is represented by a single integer. Arrays can also be multidimensional. The index of an element of a multidimensional array is a set of integers equal to the number of dimensions. Figure 9-1 shows a seating chart that represents a two-dimensional array where a student's name (value) is referenced by the ordered pair of row number, seat number (index).

	Row 0	Row 1	Row 2
Seat 0	Mary (0,0)	Jim (1,0)	Jane (2,0)
Seat 1	Bob (0,1)	Noah (1,1)	Cindy (2,1)
Seat 2	Amy (0,2)	Morgan (1,2)	Greg (2,2)

Figure 9-1. A two-dimensional array

You implement array functionality when you declare its type. The common types implemented as arrays are numeric types such as integers or double types, as well as the character and string types. When declaring a type as an array, you use square brackets ([]) after the type, followed by the name of the array. The elements of the array are designated by a comma-separated list enclosed by curly brackets ({}). For example, the following code declares an array of type `int` and fills it with five values:

```
int[] intArray = { 1, 2, 3, 4, 5 };
```

Once a type is declared as an array, the properties and methods of the `Array` class are exposed. Some of the functionality includes querying for the upper and lower bounds of the array, updating the elements of the array, and copying the elements of the array. The `Array` class contains many static methods used to work with arrays, such as methods for clearing, reversing, and sorting its elements.

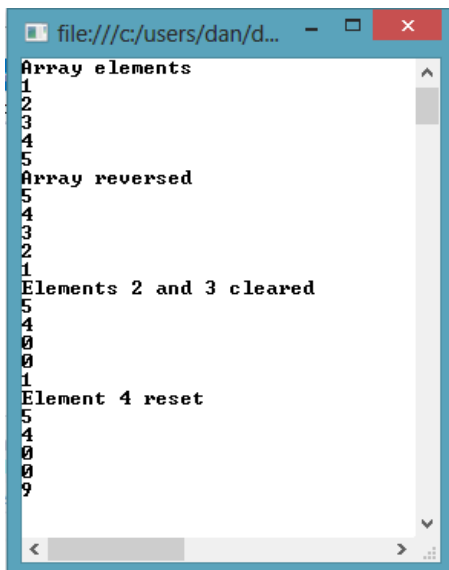
The following code demonstrates declaring and working with an array of integers. It also uses several static methods exposed by the `Array` class. Notice the `foreach` loop used to list the values of the array. The `foreach` loop provides a way to iterate through the elements of the array.

```

int[] intArray = { 1, 2, 3, 4, 5 };
Console.WriteLine("Upper Bound");
Console.WriteLine(intArray.GetUpperBound(0));
Console.WriteLine("Array elements");
foreach (int item in intArray)
{
    Console.WriteLine(item);
}
Array.Reverse(intArray);
Console.WriteLine("Array reversed");
foreach (int item in intArray)
{
    Console.WriteLine(item);
}
Array.Clear(intArray, 2, 2);
Console.WriteLine("Elements 2 and 3 cleared");
foreach (int item in intArray)
{
    Console.WriteLine(item);
}
intArray[4] = 9;
Console.WriteLine("Element 4 reset");
foreach (int item in intArray)
{
    Console.WriteLine(item);
}
Console.ReadLine();

```

Figure 9-2 shows the output of this code in the console window.



```

file:///c:/users/dan/d... - □ ×
Array elements
1
2
3
4
5
Array reversed
5
4
3
2
1
Elements 2 and 3 cleared
5
4
0
0
1
Element 4 reset
5
4
0
0
9

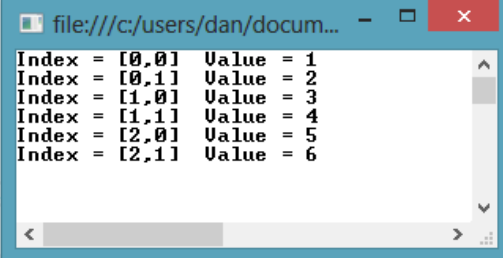
```

Figure 9-2. One-dimensional array output

Although one-dimensional arrays are the most common type you will run into, you should understand how to work with the occasional multidimensional array. Two-dimensional arrays are used to store (in active memory) and process data that fits in the rows and columns of a table. For example, you may need to process a series of measurements (temperature or radiation level) taken at hourly intervals over several days. To create a multidimensional array, you place one or more commas inside the square brackets to indicate the number of dimensions. One comma indicates two dimensions; two commas indicate three dimensions, and so forth. When filling a multidimensional array, curly brackets within curly brackets define the elements. The following code declares and fills a two-dimensional array:

```
int[,] twoDArray = { { 1, 2 }, { 3, 4 }, { 5, 6 } };
//Print the index and value of the elements
for (int i = 0; i <= twoDArray.GetUpperBound(0); i++)
{
    for (int x = 0; x <= twoDArray.GetUpperBound(1); x++)
    {
        Console.WriteLine("Index = [{0},{1}] Value = {2}", i, x, twoDArray[i, x]);
    }
}
```

Figure 9-3 shows the output of this code in the console window.



```
file:///c:/users/dan/docum...
Index = [0,0] Value = 1
Index = [0,1] Value = 2
Index = [1,0] Value = 3
Index = [1,1] Value = 4
Index = [2,0] Value = 5
Index = [2,1] Value = 6
```

Figure 9-3. Two-dimensional array output

When you work with collections, you often do not know the number of items they need to contain until runtime. This is where the `ArrayList` class fits in. The capacity of an array list automatically expands as required, with the memory reallocation and copying of elements performed automatically. The `ArrayList` class also provides methods and properties for working with the array elements that the `Array` class does not provide. The following code demonstrates some of these properties and methods. Notice that the capacity of the list expands dynamically as more names are added.

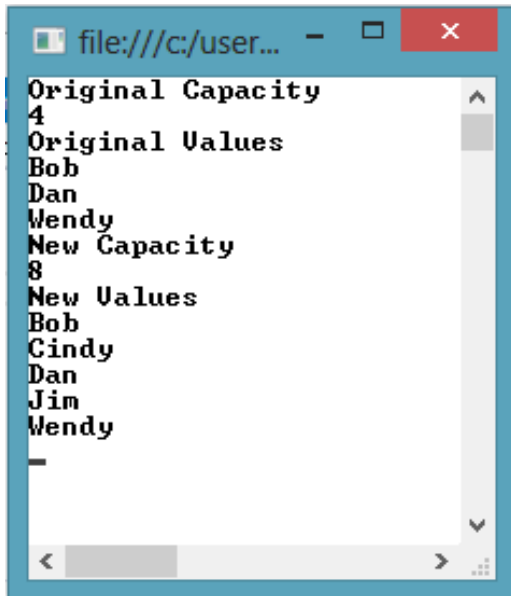
```
ArrayList nameList = new ArrayList();
nameList.Add("Bob");
nameList.Add("Dan");
nameList.Add("Wendy");
Console.WriteLine("Original Capacity");
Console.WriteLine(nameList.Capacity);
Console.WriteLine("Original Values");
foreach (object name in nameList)
{
    Console.WriteLine(name);
}
```

```

nameList.Insert(nameList.IndexOf("Dan"), "Cindy");
nameList.Insert(nameList.IndexOf("Wendy"), "Jim");
Console.WriteLine("New Capacity");
Console.WriteLine(nameList.Capacity);
Console.WriteLine("New Values");
foreach (object name in nameList)
{
    Console.WriteLine(name);
}

```

Figure 9-4 shows the output in the console window.



```

Original Capacity
4
Original Values
Bob
Dan
Wendy
New Capacity
8
New Values
Bob
Cindy
Dan
Jim
Wendy

```

Figure 9-4. The `ArrayList` output

Although it's often easier to work with an array list than with an array, an array list can have only one dimension. Also, an array of a specific type offers better performance than an array list, because the elements of `ArrayList` are of type `Object`. When types are added to the array list, they are cast to a generic `Object` type. When the items are retrieved from the list, they must be cast once again to the specific type.

ACTIVITY 9-1. WORKING WITH ARRAYS AND ARRAY LISTS

In this activity, you will become familiar with the following:

- creating and using arrays
- working with multidimensional arrays
- working with array lists

Creating and Using Arrays

To create and populate an array, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose a console application project. Name the project Activity9_1. The console application contains a class called Program with a Main method. The Main method is the first method that is accessed when the application is launched.
3. Notice that the Main method accepts an input parameter of a string array called args. The args array contains any command line arguments passed in when the console application is launched. The members of the args array are separated by a space when passed in.

```
static void Main(string[] args)
{
}
```

4. Add the following code to the Main method to display the command line arguments passed in:

```
Console.WriteLine("parameter count = {0}", args.Length);

for (int i = 0; i < args.Length; i++)
{
    Console.WriteLine("Arg[{0}] = [{1}]", i, args[i]);
}
Console.ReadLine();
```

5. In Solution Explorer, right-click the project node and choose Properties. In the project properties window, select the Debug tab. In the command line arguments field, enter “C# coding is fun” (see Figure 9-5).

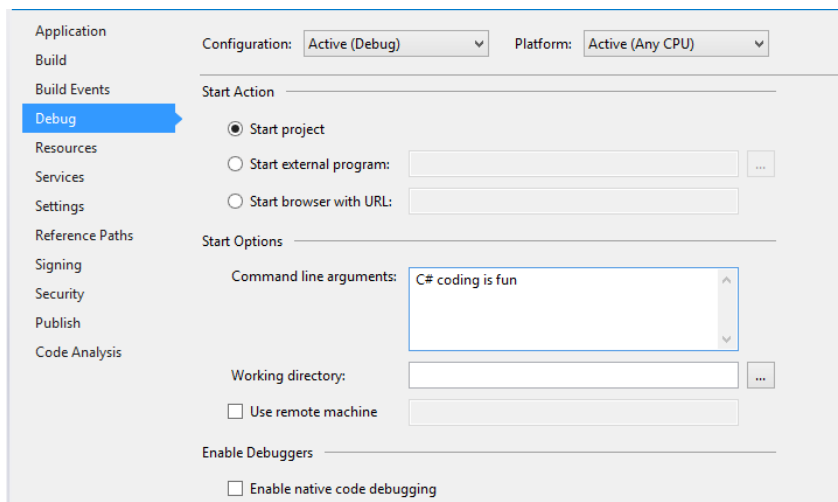


Figure 9-5. Adding command line arguments

6. Select Debug ► Start to run the project. The console window should launch with the output shown in Figure 9-6. After viewing the output, stop the debugger.

```

parameter count = 4
Arg[0] = [C#]
Arg[1] = [coding]
Arg[2] = [is]
Arg[3] = [fun]

```

Figure 9-6. The console output for the array

7. Add the following code before the `Console.ReadLine` method in the `Main` method. This code clears the value of the array at index 1 and sets the value at index 3 to “great”.

```

Array.Clear(args, 1, 1);
args[3] = "great";
for (int i = 0; i < args.Length; i++)
{
    Console.WriteLine("Arg[{0}] = [{1}]", i, args[i]);
}

```

8. Select Debug ► Start to run the project. The console window should launch with the additional output shown in Figure 9-7. After viewing the output, stop the debugger.

```

parameter count = 4
Arg[0] = [C#]
Arg[1] = [coding]
Arg[2] = [is]
Arg[3] = [fun]
Arg[0] = [C#]
Arg[1] = []
Arg[2] = [is]
Arg[3] = [great]

```

Figure 9-7. The console output for the updated array

Working with Multidimensional Arrays

To create and populate a multidimensional array, follow these steps:

1. Comment out the code in the Main method.
2. Add the following code to the Main method to create and populate a two-dimensional array:

```
string[,] seatingChart = new string[2,2];
seatingChart[0, 0] = "Mary";
seatingChart[0, 1] = "Jim";
seatingChart[1, 0] = "Bob";
seatingChart[1, 1] = "Jane";
```

3. Add the following code to loop through the array and print the names to the console window:

```
for (int row = 0; row < 2; row++)
{
    for (int seat = 0; seat < 2; seat++)
    {
        Console.WriteLine("Row: {0} Seat: {1} Student: {2}",
            (row + 1),(seat + 1),seatingChart[row, seat]);
    }
}
Console.ReadLine();
```

4. Select Debug ► Start to run the project. The console window should launch with the output that shows the seating chart of the students (see Figure 9-8).

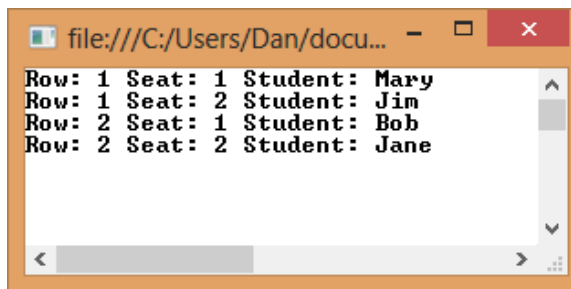


Figure 9-8. The console output for the two-dimensional array

5. After viewing the output, stop the debugger.

Working with Array Lists

Although the two dimensional array you just created works, it may be more intuitive to store the information about each student's seating assignment in a seating assignment class and then to organize these objects into an array list structure. To create and populate an array list of seating assignments, follow these steps:

1. Add a class file to the project named `SeatingAssignment.cs`.
2. Add the following code to create the `SeatingAssignment` class. This class contains a `Row`, `Seat`, and `Student` property. It also contains an overloaded constructor to set these properties.

```
public class SeatingAssignment
{
    int _row;
    int _seat;
    string _student;
    public int Row
    {
        get { return _row; }
        set { _row = value; }
    }
    public int Seat
    {
        get { return _seat; }
        set { _seat = value; }
    }
    public string Student
    {
        get { return _student; }
        set { _student = value; }
    }
    public SeatingAssignment(int row, int seat, string student)
    {
        this.Row = row;
        this.Seat = seat;
        this.Student = student;
    }
}
```

3. In the `Main` method of the `Program` class, comment out the previous code and add the following using statement to the top of the file:

```
using System.Collections;
```

4. Add the following code to create an `ArrayList` of seating assignments:

```
ArrayList seatingChart = new ArrayList();
seatingChart.Add(new SeatingAssignment(0, 0, "Mary"));
seatingChart.Add(new SeatingAssignment(0, 1, "Jim"));
seatingChart.Add(new SeatingAssignment(1, 0, "Bob"));
seatingChart.Add(new SeatingAssignment(1, 1, "Jane"));
```

5. After the `ArrayList` is populated, add the following code to write the `SeatingAssignment` information to the console window.

```
foreach (SeatingAssignment sa in seatingChart)
{
    Console.WriteLine("Row: {0} Seat: {1} Student: {2}",
        (sa.Row + 1), (sa.Seat + 1), sa.Student);
}
Console.ReadLine();
```

6. Select **Debug** ► **Start** to run the project. The console window should launch with the same output as shown in Figure 9-8 (the seating chart of the students).
7. One of the advantages of the `ArrayList` class is the ability to add and remove items dynamically. Add the following code after the code in step 4 to add two more students to the seating chart:

```
seatingChart.Add(new SeatingAssignment(2, 0, "Bill"));
seatingChart.Add(new SeatingAssignment(2, 1, "Judy"));
```

8. Select **Debug** ► **Start** to run the project. The console window should launch with the output showing the new students.
 9. When finished, stop the debugger, and close Visual Studio.
-

Using Generic Collections

Working with collections is a common requirement of application programming. Most of the data we work with needs to be organized in a collection. For example, you may need to retrieve customers from a database and load them into a drop-down list in the UI (User Interface). The customer information is represented by a customer class, and the customers are organized into a customer collection. The collection can then be sorted, filtered, and looped through for processing.

With the exception of a few of the specialized collections strongly typed to hold strings, the collections provided by the .NET Framework are weakly typed. The items held by the collections are of type `Object`, and so they can be of any type, since all types derive from the `Object` type.

Weakly typed collections can cause performance and maintenance problems for your application. One problem is that there are no inherent safeguards for limiting the types of objects stored in the collection. The same collection can hold any type of item, including dates, integers, or a custom type such as an employee object. If you build and expose a collection of integers, and that collection inadvertently gets passed a date, the chances are high that the code will fail at some point.

Fortunately, C# supports generics, and the .NET Framework provides generic-based collections in the `System.Collections.Generic` namespace. Generics let you define a class without specifying its type. The type is specified when the class is instantiated. Using a generic collection provides the advantages of type safety and the performance of a strongly typed collection while also providing the code reuse associated with weakly typed collections.

The following code shows how to create a strongly typed collection of customers using the `Generic.List` class. The list type (in this case, `Customer`) is placed between the angle brackets (`<>`). Customer objects are added to the collection, and then the customers in the collection are retrieved, and the customer information is written out to the console. (You will look at binding collections to UI controls in Chapter 11.)

```

List<Customer> customerList = new List<Customer>();
customerList.Add(new Customer
    ("WHITC", "White Clover Markets", "Karl Jablonski"));
customerList.Add(new Customer("RANCH", "Rancho grande", "Sergio Gutiérrez"));
customerList.Add(new Customer("ALFKI", "Alfreds Futterkiste", "Maria Anders"));
customerList.Add
    (new Customer("FRANR", "France restauration", "Carine Schmitt"));
foreach (Customer c in customerList)
{
    Console.WriteLine("Id: {0} Company: {1} Contact: {2}",
        c.CompanyId, c.CompanyName, c.ContactName);
}

```

There may be times when you need to extend the functionality of the collection provided by the .NET Framework. For example, you may need the ability to sort the collection of customers by either the `CompanyId` or the `CompanyName`. To implement sorting, you need to define a sorting class that implements the `IComparer` interface. The `IComparer` interface ensures the sorting class implements a `Compare` method with the appropriate signature. (Interfaces were covered in Chapter 7.) The following `CustomerSorter` class sorts a list of `Customer` by `CompanyName`. Note that since the `CompanyName` property is a string, you can use the string comparer to sort them.

```

public class CustomerSorter : IComparer<Customer>
{
    public int Compare(Customer customer1, Customer customer2)
    {
        return customer1.CompanyName.CompareTo(customer2.CompanyName);
    }
}

```

Now you can sort the customers by `CompanyName` and then display them.

```
customerList.Sort(new CustomerSorter());
```

The output is shown in Figure 9-9.



```

file:///c:/users/dan/documents/visual studio 2012/Projects/Chap9... - □ ×
Id: ALFKI Company: Alfreds Futterkiste Contact: Maria Anders
Id: FRANR Company: France restauration Contact: Carine Schmitt
Id: RANCH Company: Rancho grande Contact: Sergio Gutiérrez
Id: WHITC Company: White Clover Markets Contact: Karl Jablonski

```

Figure 9-9. The console output for the sorted list of `Customer`

ACTIVITY 9-2. IMPLEMENTING AND EXTENDING GENERIC COLLECTIONS

In this activity, you will become familiar with the following:

- implementing a generic collection
- extending a generic collection to implement sorting

To create and populate a generic list, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose a console application project. Name the project Activity9_2.
3. Select Project ► Add Class. Name the class file Request.cs.
4. Add the following properties to the Request class:

```
public class Request
{
    string _requestor;
    int _priority;
    DateTime _date;
    public string Requestor
    {
        get { return _requestor; }
        set { _requestor = value; }
    }
    public int Priority
    {
        get { return _priority; }
        set { _priority = value; }
    }
    public DateTime Date
    {
        get { return _date; }
        set { _date = value; }
    }
}
```

5. Overload the constructor of the Request class to set the properties in the constructor.

```
public Request(string requestor, int priority, DateTime date)
{
    this.Requestor = requestor;
    this.Priority = priority;
    this.Date = date;
}
```

6. Add a method to override the ToString method of the base Object class. This will return the request information as a string when the method is called.

```
public override string ToString()
{
    return String.Format("{0}, {1}, {2}",this.Requestor,
        this.Priority.ToString(), this.Date);
}
```

- Open the Program class in the code editor and add the following code to the Main method. This code populates a generic list of type Request and displays the values in the console window.

```
static void Main(string[] args)
{
    List<Request> reqList = new List<Request>();
    reqList.Add(new Request("Dan",2 ,new DateTime(2011,4,2)));
    reqList.Add(new Request("Alice", 5, new DateTime(2011, 2, 5)));
    reqList.Add(new Request("Bill", 3, new DateTime(2011, 6, 19)));
    foreach (Request req in reqList)
    {
        Console.WriteLine(req.ToString());
    }
    Console.ReadLine();
}
```

- Select Debug ► Start to run the project. The console window should launch with the request items listed in the order they were added to the reqList.
- Select Project ► Add Class. Name the class DateSorter.
- Add the following code to the DateSorter class. This class implements the IComparer interface and is used to enable sorting requests by date.

```
public class DateSorter:IComparer<Request>
{
    public int Compare(Request R1, Request R2)
    {
        return R1.Date.CompareTo(R2.Date);
    }
}
```

- Add the following code in the Main method of the Program class prior to the Console.ReadLine method. This code sorts the reqList by date and displays the values in the console window.

```
Console.WriteLine("Sorted by date.");
reqList.Sort(new DateSorter());
foreach (Request req in reqList)
{
    Console.WriteLine(req.ToString());
}
```

12. Select Debug ► Start to run the project. The console window should launch with the output shown in Figure 9-10. After viewing the output, stop the debugger and exit Visual Studio.

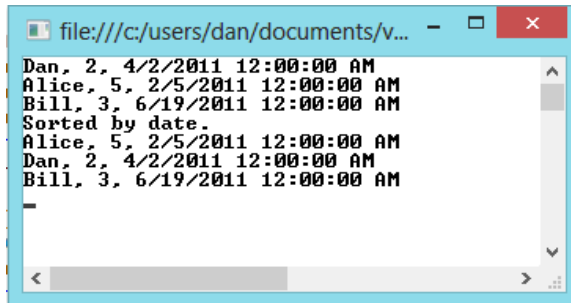


Figure 9-10. Generic collection unsorted and sorted by date

Programming with Stacks and Queues

Two special types of collections often used in programming are the stack and the queue. A stack is a last-in, first-out collection of objects. A queue represents a first-in, first-out collection of objects.

A stack is a good way to maintain a list of moves made in a chess game. When a user wants to undo his moves, he begins with his most recent move, which is the last one added to the list and also the first one retrieved. Another example of using a stack occurs when a program executes a series of method calls. A stack maintains the addresses of the methods, and execution returns to the methods in the reverse order in which they were called. When placing items in a stack, you use the push method. The pop method removes items from the stack. The peek method returns the object at the top of the stack without removing it. The following code demonstrates adding and removing items from a stack. In this case, you're using generics to implement a stack of `ChessMove` objects. The `RecordMove` method adds the most recent move to the stack. The `GetLastMove` method returns the most recent move on the stack.

```
Stack<ChessMove> moveStack = new Stack<ChessMove>();
void RecordMove(ChessMove move)
{
    moveStack.Push(move);
}
ChessMove GetLastMove()
{
    return moveStack.Pop();
}
```

An application that services help desk requests is a good example of when to use a queue. A collection maintains a list of help desk requests sent to the application. When requests are retrieved from the collection for processing, the first ones in should be the first ones retrieved. The `Queue` class uses the `enqueue` and `dequeue` methods to add and remove items. It also implements the `peek` method to return the item at the beginning of the queue without removing the item. The following code demonstrates adding and removing items from a `PaymentRequest` queue. The `AddRequest` method adds a request to the queue and the `GetNextRequest` method removes a request from the queue.


```

Queue<PaymentRequest> payRequest = new Queue<PaymentRequest>();
void AddRequest(PaymentRequest request)
{
    payRequest.Enqueue(request);
}
PaymentRequest GetNextRequest()
{
    return payRequest.Dequeue();
}

```

ACTIVITY 9-3. IMPLEMENTING STACKS AND QUEUES

In this activity, you will become familiar with the following:

- implementing a stack collection
- implementing a queue collection

To create and populate a generic list, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose a console application project. Name the project `Activity9_3`.
3. Add the following code to the `Main` method of the `Program` class. This code creates a stack of strings and loads it with strings representing moves in a game. It then uses the `Peek` method to write out the move stored at the top of the stack to the console window.

```

Stack<string> moveStack = new Stack<string>();
Console.WriteLine("Loading Stack");
for (int i = 1; i <= 5; i++)
{
    moveStack.Push("Move " + i.ToString());
    Console.WriteLine(moveStack.Peek().ToString());
}

```

4. Add the following code to the `Main` method of the `Program` class after the code in step 3. This code removes the moves from the stack using the `Pop` method and writes them to the console window.

```

Console.WriteLine("Press the enter key to unload the stack.");
Console.ReadLine();
for (int i = 1; i <= 5; i++)
{
    Console.WriteLine(moveStack.Pop().ToString());
}
Console.ReadLine();

```

5. Select Debug ► Start to run the project. The console window should launch with the output shown in Figure 9-11. Notice the last-in, first-out pattern of the stack. After viewing the output, stop the debugger.

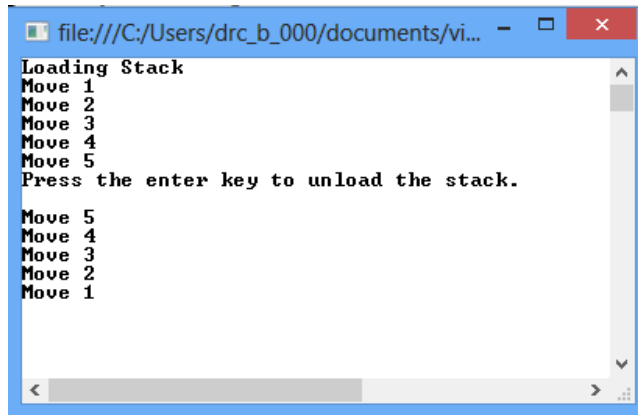


Figure 9-11. The last-in, first-out pattern of the stack

6. Comment out the code entered in steps 3 and 4. Add the following code to the `Main` method of the `Program` class. This code creates a queue of strings and loads it with strings representing requests to a consumer help line. It then uses the `Peek` method to write out the request stored at the beginning of the queue of to the console window.

```
Queue<string> requestQueue = new Queue<string>();
Console.WriteLine("Loading Queue");
for (int i = 1; i <= 5; i++)
{
    requestQueue.Enqueue("Request " + i.ToString());
    Console.WriteLine(requestQueue.Peek().ToString());
}
```

7. Add the following code to the `Main` method of the `Program` class after the code in step 6. This code removes the requests from the queue using the `Dequeue` method and writes them to the console window.

```
Console.WriteLine("Press the enter key to unload the queue.");
Console.ReadLine();
for (int i = 1; i <= 5; i++)
{
    Console.WriteLine(requestQueue.Dequeue().ToString());
}
Console.ReadLine();
```

8. Select **Debug** ► **Start** to run the project. The console window should launch with the output shown in Figure 9-12. Notice that as the queue is loaded the first request stays at the top of the queue. Also notice the first-in, first-out pattern of the queue. After viewing the output, stop the debugger and exit Visual Studio.

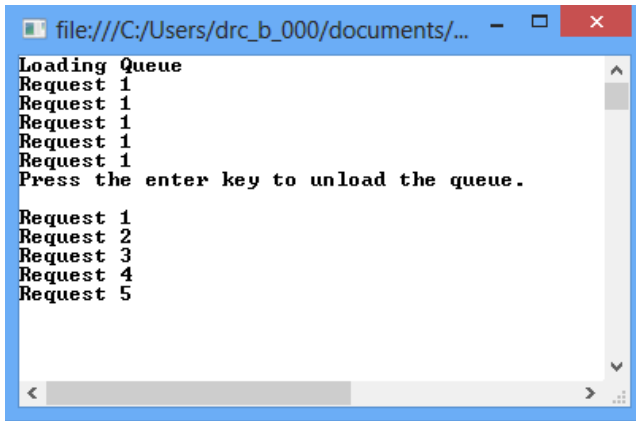


Figure 9-12. The first-in, first-out pattern of the queue

Summary

In this chapter, you examined the various types of collections exposed by the .NET Framework. You learned how to work with arrays, array lists, queues, stacks, and generic collections.

This chapter is the final one in a series that introduced you to the various OOP constructs such as classes, inheritance, and polymorphism. You should have a firm understanding of how class structures, object collaboration, and collections are implemented in C#. You have been introduced to the Visual Studio IDE and you've practiced using it. You are now ready to put the pieces together and develop a working application.

Chapter 10 is the first in a series in which you will develop .NET applications. In the process, you will investigate data access using ADO.NET and the Entity Framework; create a Windows-based GUI using the Windows Presentation Foundation (WPF) and the new Windows 8 app store; create a web-based GUI using Web Forms and ASP.NET; and create web services using the Windows Communication Foundation (WCF) and the ASP.NET Web API.



Implementing the Data Access Layer

In the past several chapters, you have looked at various object-oriented programming constructs such as classes, inheritance, and polymorphism as they are implemented in C# code. You have been introduced to and practiced using the Visual Studio integrated development environment. You should also have a firm understanding of how class structures and object collaboration are implemented.

You are now ready to put the pieces together and develop a working application. Because most business applications involve working with and updating data in a back-end relational database, you will look at how the .NET Framework provides the functionality to work with relational data.

After reading this chapter, you will understand the following:

- how to establish a connection to a database using the `Connection` object
- how to use a `Command` object to execute SQL queries
- how to use a `Command` object to execute stored procedures
- how to retrieve records with the `DataReader` object
- how to populate `DataTables` and `DataSets`
- how to establish relationships between tables in a `DataSet`
- how to edit and update data in a `DataSet`
- how to create an Entity Data Model
- how to use Language-Integrated Query (LINQ) to Entity Framework (EF) to query data
- how to use the Entity Framework to update data

Introducing ADO.NET

A majority of applications developed for businesses need to interact with a data storage device. Data storage can occur in many different forms: for example, in a flat file system, as is the case with many traditional mainframe systems, or in a relational database management system, such as SQL Server, Oracle, or MySQL. You can also maintain data in a hierarchical textual file structure, as is the case with XML. To access and work with data in a consistent way across these various data stores, the .NET Framework provides a set of classes organized into the `System.Data` namespace. This collection of classes is known as ADO.NET.

Looking at the history of Microsoft's data access technologies reveals an evolution from a connected model to a disconnected one. When developing the traditional two-tier client-server applications prevalent in the 1980s and early 1990s, it was often more efficient to open a connection with the database, work with the data implementing server-side cursors, and close the connection when finished working with the data. The problem with this approach became apparent in the late 1990s as companies tried to evolve their data-driven applications from traditional

two-tier client-server applications to multi-tiered web-based models: opening and holding a connection open until processing was complete was not scalable. Scalability is the ability of an application to handle an increasing number of simultaneous clients without a noticeable degradation of performance. Microsoft designed ADO.NET to be highly scalable. To achieve scalability, Microsoft designed ADO.NET around a disconnected model. A connection is made to the database, the data and metadata are retrieved and cached locally, and the connection is closed.

Another problem with the traditional data access technologies developed during this time was the lack of interoperability. Systems with a high degree of interoperability can easily exchange data back and forth regardless of the implementation technologies of the various systems. Traditional data access technologies rely on proprietary methods of data exchange. Using these techniques, it is hard for a system built using Microsoft technologies such as ADO (pre-.NET) and DCOM to exchange data with a system built using Java technologies such as JDBC and CORBA. The industry as a whole realized it was in the best interest of all parties to develop open standards for exchanging data between disparate systems. Microsoft has embraced these standards and has incorporated support of the standards into the .NET Framework.

Working with Data Providers

To establish a connection to a data source, such as a SQL Server database, and work with its data, you must use the appropriate .NET provider classes. The SQL Server provider classes are located in the `System.Data.SqlClient` namespace. Other data providers exist, such as the OleDb data provider for Oracle classes located in the `System.Data.OleDb` namespace. Each of these providers implements a similar class structure, which you can use to interact with its intended data source. Table 10-1 summarizes the main classes of the `System.Data.SqlClient` provider namespace.

Table 10-1. *Classes in the System.Data.SqlClient Namespace*

Class	Responsibility
<code>SqlConnection</code>	Establishes a connection and a unique session with a database.
<code>SqlCommand</code>	Represents a Transact-SQL statement or stored procedure to execute at the database.
<code>SqlDataReader</code>	Provides a means of reading a forward-only stream of rows from the database.
<code>SqlDataAdapter</code>	Fills a <code>DataSet</code> and updates changes back to the database.
<code>SqlParameter</code>	Represents a parameter used to pass information to and from stored procedures.
<code>SqlTransaction</code>	Represents a Transact-SQL transaction to be made in the database.
<code>SqlError</code>	Collects information relevant to a warning or error returned by the database server.
<code>SqlException</code>	Defines the exception that is thrown when a warning or error is returned by the database server.

A similar set of classes exists in the `System.Data.OleDb` provider namespace. For example, instead of the `SqlConnection` class, you have an `OleDbConnection` class.

Establishing a Connection

The first step to retrieving data from a database is to establish a connection, which is done using a `Connection` object based on the type of provider being used. To establish a connection to SQL Server, you instantiate a `Connection` object of type `SqlConnection`. You also need to provide the `Connection` object with a `ConnectionString`. The `ConnectionString` consists of a series of semicolon-delimited name-value pairs that provide information needed

to connect to the database server. Some of the information commonly passed by the `ConnectionString` is the name of the target server, the name of the database, and security information. The following code demonstrates a `ConnectionString` used to connect to a SQL Server database:

```
"Data Source=TestServer;Initial Catalog=Pubs;User ID=Dan;Password=training"
```

The attributes you need to provide through the `ConnectionString` are dependent on the data provider you are using. The following code demonstrates a `ConnectionString` used to connect to an Access database using the `OleDb` provider for Access:

```
"Provider=Microsoft.Jet.OleDb.4.0;Data Source=D:\Data\Northwind.mdb"
```

The next step is to invoke the `Open` method of the `Connection` object. This will result in the `Connection` object loading the appropriate driver and opening a connection to the data source. Once the connection is open, you can work with the data. After you are done interacting with the database, it is important you invoke the `Dispose` method of the `Connection` object, because when a `Connection` object falls out of scope or is garbage-collected, the connection is not implicitly released. The following code demonstrates the process of opening a connection to the Pubs database in SQL Server, working with the data, and disposing the connection:

```
SqlConnection pubConnection = new SqlConnection();
string connString;
try
{
    connString = "Data Source=drcsv01;Initial Catalog=pubs;Integrated Security=True";
    pubConnection.ConnectionString = connString;
    pubConnection.Open();
    //work with data
}
catch (SqlException ex)
{
    throw ex;
}
finally
{
    pubConnection.Dispose();
}
```

C# provides a `using` statement to aid in ensuring connections get closed and resources are disposed of properly. When the execution leaves the scope of the `using` statement, the connection is automatically closed and the connection object is disposed of efficiently. The previous code can be rewritten to take advantage of the `using` statement as follows:

```
string connString = "Data Source=drcsv01;Initial Catalog=pubs;Integrated Security=True";
using(SqlConnection pubConnection = new SqlConnection())
{
    try
    {
        pubConnection.ConnectionString = connString;
        pubConnection.Open();
        //work with data
    }
}
```

```

    catch (SqlException ex)
    {
        throw ex;
    }
}

```

Executing a Command

Once your application has established and opened a connection to a database, you can execute SQL statements against it. A `Command` object stores and executes command statements against the database. You can use the `Command` object to execute any valid SQL statement understood by the data store. In the case of SQL Server, these can be Data Manipulation Language statements (Select, Insert, Update, and Delete), Data Definition Language statements (Create, Alter, and Drop), or Data Control Language statements (Grant, Deny, and Revoke). The `CommandText` property of the `Command` object holds the SQL statement that will be submitted. The `Command` object contains three methods for submitting the `CommandText` to the database depending on what is returned. If records are returned, as is the case when a `Select` statement is executed, then you can use the `ExecuteReader`. If a single value is returned—for example, the results of a `Select Count` aggregate function—you should use the `ExecuteScalar` method. When no records are returned from a query—for example, from an `Insert` statement—you should use the `ExecuteNonQuery` method. The following code demonstrates using a `Command` object to execute a SQL statement against the Pubs database that returns the number of employees. Notice the use of the nested `using` clause, one for the connection and one for the command.

```

public int GetEmployeeCount()
{
    string connString = "Data Source=drcsv01;Initial Catalog=pubs;Integrated Security=True";

    using (SqlConnection pubConnection = new SqlConnection())
    {
        using (SqlCommand pubCommand = new SqlCommand())
        {
            try
            {
                pubConnection.ConnectionString = connString;
                pubConnection.Open();
                pubCommand.Connection = pubConnection;
                pubCommand.CommandText = "Select Count(emp_id) from employee";
                return (int)pubCommand.ExecuteScalar();
            }
            catch (SqlException ex)
            {
                throw ex;
            }
        }
    }
}

```

Using Stored Procedures

In many application designs, instead of executing a SQL statement directly, clients must execute stored procedures. Stored procedures are an excellent way to encapsulate the database logic, increase scalability, and enhance the security of multi-tiered applications. To execute a stored procedure, you use a `Command` object, setting its `CommandType` property to `StoredProcedure` and its `CommandText` property to the name of the stored procedure. The following code executes a stored procedure that returns the number of employees in the Pubs database:

```
public int GetEmployeeCount()
{
    string connString = "Data Source=drcsv01;Initial Catalog=pubs;Integrated Security=True";

    using (SqlConnection pubConnection = new SqlConnection())
    {
        using (SqlCommand pubCommand = new SqlCommand())
        {
            try
            {
                pubConnection.ConnectionString = connString;
                pubConnection.Open();
                pubCommand.Connection = pubConnection;
                pubCommand.CommandType = CommandType.StoredProcedure;
                pubCommand.CommandText = "GetEmployeeCount";
                return (int)pubCommand.ExecuteScalar();
            }
            catch (SqlException ex)
            {
                throw ex;
            }
        }
    }
}
```

When executing a stored procedure, you often must supply input parameters. You may also need to retrieve the results of the stored procedure through output parameters. To work with parameters, you need to instantiate a parameter object of type `SqlParameter`, and then add it to the `Parameters` collection of the `Command` object. When constructing the parameter, you supply the name of the parameter and the SQL Server data type. For some data types, you also supply the size. If the parameter is an output, input-output, or return parameter, then you must indicate the parameter direction. The following example calls a stored procedure that accepts an input parameter of a letter. The procedure passes back a count of the employees whose last name starts with the given letter. The count is returned in the form of an output parameter.

```
public int GetEmployeeCount(string lastInitial)
{
    string connString = "Data Source=drcsv01;Initial Catalog=pubs;Integrated Security=True";
    using (SqlConnection pubConnection = new SqlConnection(connString))
    {
        using (SqlCommand pubCommand = new SqlCommand())
```



```

    {
        try
        {
            pubConnection.Open();
            pubCommand.Connection = pubConnection;
            pubCommand.CommandText = "GetEmployeeCountByLastInitial";
            SqlParameter inputParameter = pubCommand.Parameters.Add
                (@LastInitial, SqlDbType.NChar, 1);
            inputParameter.Value = lastInitial.ToCharArray()[0];
            SqlParameter outputParameter = pubCommand.Parameters.Add
                (@EmployeeCount, SqlDbType.Int);
            outputParameter.Direction = ParameterDirection.Output;
            pubCommand.CommandType = CommandType.StoredProcedure;
            pubCommand.ExecuteNonQuery();
            return (int)outputParameter.Value;
        }
        catch (SqlException ex)
        {
            throw ex;
        }
    }
}

```

Using the DataReader Object to Retrieve Data

A `DataReader` object accesses data through a forward-only, read-only stream. Oftentimes, you will want to loop through a set of records and process the results sequentially without the overhead of maintaining the data in a cache. A good example of this would be loading a list or array with the values returned from the database. After declaring an object of type `SqlDataReader`, you instantiate it by invoking the `ExecuteReader` method of a `Command` object. The `Read` method of the `DataReader` object accesses the records returned. The `Close` method of the `DataReader` object is called after the records have been processed. The following code demonstrates the use of a `DataReader` object to retrieve a list of names from a SQL Server database and return it to the client:

```

public ArrayList ListNames()
{
    string connString = "Data Source=localhost;Initial Catalog=pubs;Integrated Security=True";
    using (SqlConnection pubConnection = new SqlConnection(connString))
    {
        using (SqlCommand pubCommand = new SqlCommand())
        {
            try
            {
                pubConnection.ConnectionString = connString;
                pubConnection.Open();
                pubCommand.Connection = pubConnection;
                pubCommand.CommandText =
                    "Select lname from employee";
                using (SqlDataReader employeeDataReader = pubCommand.ExecuteReader())

```

```

        {
            ArrayList nameArray = new ArrayList();
            while (employeeDataReader.Read())
            {
                nameArray.Add(employeeDataReader["lname"]);
            }
            return nameArray;
        }
    }
}
catch (SQLException ex)
{
    throw ex;
}
}
}
}

```

Using the DataAdapter to Retrieve Data

In many cases, you need to retrieve a set of data from a database, work with the data, and return any updates to the data back to the database. In that case, you use a `DataAdapter` as a bridge between the data source and the in-memory cache of the data. This in-memory cache of data is contained in either a stand alone `DataTable` or a `DataSet`, which is a collection of `DataTables`.

■ **Note** The `DataTable` and `DataSet` objects are discussed in greater detail in the “Working with `DataTables` and `DataSets`” section.

To retrieve a set of data from a database, first you instantiate a `DataAdapter` object. You then set the `SelectCommand` property of the `DataAdapter` to an existing `Command` object. Finally, you execute the `Fill` method, passing the name of a `DataSet` object to fill. If the connection object is not open when the fill method is called it is opened to retrieve the data and then closed. If it is open when the fill method is called, it remains open after the data is retrieved. Here you see how to use a `DataAdapter` to fill a `DataTable` and pass the `DataTable` back to the client:

```

public DataTable GetEmployees()
{
    string connString = "Data Source=localhost;Initial Catalog=pubs;Integrated Security=True";
    using (SqlConnection pubConnection = new SqlConnection(connString))
    {
        using (SqlCommand pubCommand = new SqlCommand())
        {
            pubCommand.Connection = pubConnection;
            pubCommand.CommandText = "Select emp_id, lname, Hire_Date from employee";
            using (SqlDataAdapter employeeAdapter = new SqlDataAdapter())

```

```

        {
            employeeAdapter.SelectCommand = pubCommand;
            DataTable employeeDataTable = new DataTable();
            employeeAdapter.Fill(employeeDataTable);
            return employeeDataTable;
        }
    }
}

```

You may find that you need to retrieve a set of data by executing a stored procedure as opposed to passing in a SQL statement. The following code demonstrates executing a stored procedure that accepts an input parameter and returns a set of records. The records are loaded into a `DataSet` object and returned to the client.

```

public DataSet GetEmployees(string lastInitial)
{
    string connString = "Data Source=localhost;Initial Catalog=pubs;Integrated Security=True";
    using (SqlConnection pubConnection = new SqlConnection(connString))
    {
        using (SqlCommand pubCommand = new SqlCommand())
        {
            pubCommand.Connection = pubConnection;
            pubCommand.CommandText = "Select emp_id, lname, Hire_Date from employee";
            pubCommand.CommandText = "GetEmployeesByLastInitial";
            SqlParameter inputParameter = pubCommand.Parameters.Add
                ("@LastInitial", SqlDbType.NChar, 1);
            inputParameter.Value = lastInitial.ToCharArray()[0];
            pubCommand.CommandType = CommandType.StoredProcedure;
            using (SqlDataAdapter employeeAdapter = new SqlDataAdapter())
            {
                employeeAdapter.SelectCommand = pubCommand;
                DataSet employeeDataSet = new DataSet();
                employeeAdapter.Fill(employeeDataSet);
                return employeeDataSet;
            }
        }
    }
}

```

ACTIVITY 10-1. RETRIEVING DATA FROM A SQL SERVER DATABASE

In this activity, you will become familiar with the following:

- establishing a connection to a SQL Server database
- executing queries through a `Command` object
- retrieving data with a `DataReader` object
- executing a stored procedure using a `Command` object

■ **Note** For the activities in this chapter to work, you must have access to a SQL Server 2005 or higher database server with the sample Microsoft Pubs and Northwind databases installed. You must be logged on under a Windows account that has been given the appropriate rights to these databases. You may have to alter the `ConnectionString` depending on your settings. For more information, refer to the “Software Requirements” section in the Introduction and Appendix C.

Creating a Connection and Executing SQL Queries

To create a connection and execute SQL queries, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose a Console Application project. Name the project `Activity10_1`.
3. After the project opens, add a new class to the project named `Author`.
4. Open the `Author` class code in the code editor. Add the following using statements at the top of the file:

```
using System.Data;
using System.Data.SqlClient;
```

5. Add code to declare a private class-level variable containing the connection string:

```
class Author
{
    string _connString = "Data Source=localhost;Initial Catalog=pubs;Integrated
Security=True";
```

6. Add a method to the class that will use a `Command` object to execute a query to count the number of authors in the `Authors` table. Because you are only returning a single value, you will use the `ExecuteScalar` method of the `Command` object.

```
public int CountAuthors()
{
    using (SqlConnection pubConnection = new SqlConnection(_connString))
    {
        using (SqlCommand pubCommand = new SqlCommand())
        {
            pubCommand.Connection = pubConnection;
            pubCommand.CommandText = "Select Count(au_id) from authors";
            pubConnection.Open();
            return (int)pubCommand.ExecuteScalar();
        }
    }
}
```

7. Add the following code to the `Main` method of the `Program` class, which will execute the `GetAuthorCount` method defined in the `Author` class:

```
static void Main(string[] args)
{
    try
    {
        Author author = new Author();
        Console.WriteLine(author.CountAuthors());
        Console.ReadLine();
    }
    catch (Exception ex)
    {
        Console.WriteLine(ex.Message);
        Console.ReadLine();
    }
}
```

8. Select **Debug** ► Start to run the project. The Console window should launch with the number of authors displayed. After viewing the output, stop the debugger.

Using the `DataReader` Object to Retrieve Records

To use the `DataReader` object to retrieve records, follow these steps:

1. Open the `Author` class code in the code editor.
2. Add a public method to the class definition called `GetAuthorList` that returns a generic `List` of strings.

```
public List<string> GetAuthorList()
{
}
```

3. Add the following code, which executes a SQL `Select` statement to retrieve the authors' last names. A `DataReader` object then loops through the records and creates a list of names that gets returned to the client.

```
List<string> nameList = new List<string>();
using (SqlConnection pubConnection = new SqlConnection(_connString))
{
    using (SqlCommand authorsCommand = new SqlCommand())
    {
        authorsCommand.Connection = pubConnection;
        authorsCommand.CommandText = "Select au_lname from authors";
        pubConnection.Open();
        using (SqlDataReader authorDataReader = authorsCommand.ExecuteReader())
        {
            while (authorDataReader.Read() == true)
            {
                nameList.Add(authorDataReader.GetString(0));
            }
        }
    }
}
```

```

        return nameList;
    }
}

```

4. Change the code in the Main method of the Program class to show the list of names in the console window.

```

static void Main(string[] args)
{
    try
    {
        Author author = new Author();
        foreach (string name in author.GetAuthorList())
        {
            Console.WriteLine(name);
        }
        Console.ReadLine();
    }
    catch (Exception ex)
    {
        Console.WriteLine(ex.Message);
        Console.ReadLine();
    }
}

```

5. Select Debug ► Start to run the project. The Console window should launch with the names of the authors displayed. After viewing the output, stop the debugger.

Executing a Stored Procedure Using a Command Object

To execute a stored procedure using a Command object, follow these steps:

1. Open the Author class code in the code editor.
2. Add a public method that overloads the GetAuthorList method by accepting an integer parameter named Royalty. This function will call the stored procedure by royalty in the Pubs database. The procedure takes an integer input of the royalty percentage and returns a list of author IDs with the percentage.

```

public List<string> GetAuthorList(int royalty)
{
    List<string> nameList = new List<string>();
    using (SqlConnection pubConnection = new SqlConnection(_connString))
    {
        using (SqlCommand authorsCommand = new SqlCommand())
        {
            authorsCommand.Connection = pubConnection;
            authorsCommand.CommandType = CommandType.StoredProcedure;
            authorsCommand.CommandText = "byroyalty";
            SqlParameter inputParameter = new SqlParameter();

```

```

        inputParameter.ParameterName = "@percentage";
        inputParameter.Direction = ParameterDirection.Input;
        inputParameter.SqlDbType = SqlDbType.Int;
        inputParameter.Value = royalty;
        authorsCommand.Parameters.Add(inputParameter);
        pubConnection.Open();
        using (SqlDataReader authorDataReader = authorsCommand.ExecuteReader())
        {
            while (authorDataReader.Read() == true)
            {
                nameList.Add(authorDataReader.GetString(0));
            }
        }
        return nameList;
    }
}
}

```

3. In the Main method of the Program class, supply an input parameter of 25 to the GetAuthorList method.

```
foreach (string name in author.GetAuthorList(25))
```

4. Select Debug ► Start to run the project. The Console window should launch with the IDs of the authors displayed. After viewing the output, stop the debugger.
 5. When finished testing, exit Visual Studio.
-

Working with DataTables and DataSets

DataTables and DataSets are in-memory caches of data that provide a consistent relational programming model for working with data regardless of the data source. A DataTable represents one table of relational data and consists of columns, rows, and constraints. You can think of a DataSet as a minirelational database, which includes the data tables and the relational integrity constraints between them. If you are retrieving data from a single table, you can populate and use the DataTable directly without the overhead of creating a DataSet first. There are several ways to create a DataTable or DataSet. The most obvious method is to populate a DataTable or DataSet from an existing relational database management system (RDBMS) such as a SQL Server database. As mentioned previously, a DataAdapter object provides the bridge between the RDBMS and the DataTable or DataSet. By using a DataAdapter object, the DataTable or DataSet is totally independent from the data source. Although you need to use a specific set of provider classes to load either type of object, you use the same set of .NET Framework classes to work with a DataTable or DataSet, regardless of how it was created and populated. The System.Data namespace contains the framework classes for working with DataTable or DataSet objects. Table 10-2 lists some of the main classes contained in the System.Data namespace.

Table 10-2. *The Main Members of the System.Data Namespace*

Class	Description
DataSet	Represents a collection of DataTable and DataRelation objects. Organizes an in-memory cache of relational data.
DataTable	Represents a collection of DataColumn, DataRow, and Constraint objects. Organizes records and fields related to a data entity.
DataColumn	Represents the schema of a column in a DataTable.
DataRow	Represents a row of data in a DataTable.
Constraint	Represents a constraint that can be enforced on DataColumn objects.
ForeignKeyConstraint	Enforces referential integrity of a parent/child relationship between two DataTable objects.
UniqueConstraint	Enforces uniqueness of a DataColumn or set of DataColumns. This is required to enforce referential integrity in a parent/child relationship.
DataRelation	Represents a parent/child relation between two DataTable objects.

Populating a DataTable from a SQL Server Database

To retrieve data from a database, you set up a connection with the database using a Connection object. After a connection is established, you create a Command object to retrieve the data from the database. As stated earlier, if you are retrieving data from a single table or result set, you can populate and work with a DataTable directly without creating a DataSet object. The Load method of the DataTable fills the table with the contents of a DataReader object. The following code fills a DataTable with data from the publishers table of the Pubs database:

```
public DataTable GetPublishers()
{
    string connString = "Data Source=drcsrv01;" +
        "Initial Catalog=pubs;Integrated Security=True";
    DataTable pubTable;
    using (SqlConnection pubConnection = new SqlConnection(connString))
    {
        using (SqlCommand pubCommand = new SqlCommand())
        {
            pubCommand.Connection = pubConnection;
            pubCommand.CommandText =
                "Select pub_id, pub_name, city from publishers";
            pubConnection.Open();
            using (SqlDataReader pubDataReader = pubCommand.ExecuteReader())
            {
                pubTable = new DataTable();
                pubTable.Load(pubDataReader);
                return pubTable;
            }
        }
    }
}
```


Populating a DataSet from a SQL Server Database

When you need to load data into multiple tables and maintain the referential integrity between the tables, you need to use the `DataSet` object as a container for the `DataTables`. To retrieve data from a database and fill the `DataSet`, you set up a connection with the database using a `Connection` object. After a connection is established, you create a `Command` object to retrieve the data from the database, and then create a `DataAdapter` to fill the `DataSet`, setting the previously created `Command` object to the `SelectCommand` property of the `DataAdapter`. Create a separate `DataAdapter` for each `DataTable`. The final step is to fill the `DataSet` with the data by executing the `Fill` method of the `DataAdapter`. The following code demonstrates filling a `DataSet` with data from the publishers table and the titles table of the Pubs database:

```
public DataSet GetPubsDS()
{
    string connString = "Data Source=localhost;" +
        "Initial Catalog=pubs;Integrated Security=True";
    DataSet bookInfoDataSet= new DataSet();
    using (SqlConnection pubConnection = new SqlConnection(connString))
    {
        //Fill pub table
        using (SqlCommand pubCommand = new SqlCommand())
        {
            pubCommand.Connection = pubConnection;
            pubCommand.CommandText =
                "Select pub_id, pub_name, city from publishers";
            using (SqlDataAdapter pubDataAdapter = new SqlDataAdapter())
            {
                pubDataAdapter.SelectCommand = pubCommand;
                pubDataAdapter.Fill(bookInfoDataSet, "Publishers");
            }
        }
        //Fill title table
        using (SqlCommand titleCommand = new SqlCommand())
        {
            titleCommand.Connection = pubConnection;
            titleCommand.CommandText =
                "Select pub_id, title, ytd_sales from titles";
            using (SqlDataAdapter titleDataAdapter = new SqlDataAdapter())
            {
                titleDataAdapter.SelectCommand = titleCommand;
                titleDataAdapter.Fill(bookInfoDataSet, "Titles");
            }
        }
    }
    return bookInfoDataSet;
}
```

Establishing Relationships between Tables in a DataSet

In an RDBMS system, referential integrity between tables is enforced through a primary key and foreign key relationship. Using a `DataRelation` object, you can enforce data referential integrity between the tables in the `DataSet`. This object contains an array of `DataColumn` objects that define the common field(s) between the parent

table and the child table used to establish the relation. Essentially, the field identified in the parent table is the primary key, and the field identified in the child table is the foreign key. When establishing a relationship, create two `DataColumn` objects for the common column in each table. Next, create a `DataRelation` object, pass a name for the `DataRelation`, and pass the `DataColumn` objects to the constructor of the `DataRelation` object. The final step is to add the `DataRelation` to the `Relations` collection of the `DataSet` object. The following code establishes a relationship between the publishers and the titles tables of the `bookInfoDataSet` created in the previous section:

```
//Create relationship between tables
DataRelation Pub_TitleRelation;
DataColumn Pub_PubIdColumn;
DataColumn Title_PubIdColumn;
Pub_PubIdColumn = bookInfoDataSet.Tables["Publishers"].Columns["pub_id"];
Title_PubIdColumn = bookInfoDataSet.Tables["Titles"].Columns["pub_id"];
Pub_TitleRelation = new DataRelation("PubsToTitles", Pub_PubIdColumn, Title_PubIdColumn);
bookInfoDataSet.Relations.Add(Pub_TitleRelation);
```

Editing Data in the DataSet

Clients often need to be able to update a `DataSet`. They may need to add records, delete records, or update an existing record. Because `DataSet` objects are disconnected by design, the changes made to the `DataSet` are not automatically propagated back to the database. They are held locally until the client is ready to replicate the changes back to the database. To replicate the changes, you invoke the `Update` method of the `DataAdapter`, which determines what changes have been made to the records and implements the appropriate SQL command (`Update`, `Insert`, or `Delete`) that has been defined to replicate the changes back to the database.

When you create an `Update Command`, the command text references parameters in the command's `Parameters` collection that correspond to the fields in the `DataTable`.

```
SqlCommand updateCommand = new SqlCommand();
string updateSQL = "Update publishers set pub_name = @pub_name, " +
    " city = @city where pub_id = @pub_id";
updateCommand = new SqlCommand(updateSQL, pubConnection);
updateCommand.CommandType = CommandType.Text;
```

A `Parameter` object is added to the `Command` object's `Parameter` collection for each parameter in the `Update` statement. The `Add` method of the `Parameters` collection is passed information on the name of the parameter, the SQL data type, size, and the source column of the `DataSet`.

```
updateCommand.Parameters.Add("@pub_id", SqlDbType.Char, 4, "pub_id");
updateCommand.Parameters.Add("@city", SqlDbType.VarChar, 20, "city");
updateCommand.Parameters.Add("@pub_name", SqlDbType.VarChar, 40, "pub_name");
```

Once the parameters are created, the `updateCommand` object is set to the `UpdateCommand` property of the `DataAdapter` object.

```
pubDataAdapter.UpdateCommand = updateCommand;
```

Now that the `SqlDataAdapter` has been set up, you call the `Update` method of the `SqlDataAdapter` passing in the `DataSet` and the name of the table to update. The `SqlDataAdapter` examines the `RowState` property of the rows in the table and executes the update statement iteratively for each row to be updated.

```
pubDataAdapter.Update(bookInfoDataSet, "Publishers");
```

In a similar fashion, you could implement the `InsertCommand` and the `DeleteCommand` properties of the `DataAdapter` to allow clients to insert new records or delete records in the database.

■ **Note** For simple updates to a single table in the data source, the .NET Framework provides a `CommandBuilder` class to automate the creation of the `InsertCommand`, `UpdateCommand`, and `DeleteCommand` properties of the `DataAdapter`.

ACTIVITY 10-2. WORKING WITH DATASET OBJECTS

In this activity, you will become familiar with the following:

- populating a `DataSet` from a SQL Server database
- editing data in a `DataSet`
- updating changes from the `DataSet` to the database
- establishing relationships between tables in a `DataSet`

Populating a `DataSet` from a SQL Server Database

To populate a `DataSet` from a SQL Server database, follow these steps:

1. Start Visual Studio. Select **File** ► **New** ► **Project**.
2. Choose **Windows Forms Application**. Rename the project to `Activity10_2` and click the **OK** button.
3. After the project opens, add a new class to the project named `Author`.
4. Open the `Author` class code in the code editor. Add the following `using` statements at the top of the file:

```
using System.Data;
using System.Data.SqlClient;
```

5. Add the following code to declare a private class level variable for the connection string.

```
class Author
{
    string _connString = "Data Source=localhost;" +
        "Initial Catalog=pubs;Integrated Security=True";
```

6. Create a method of the `Author` class called `GetData` that will use a `DataAdapter` object to fill the `DataSet` and return it to the client.

```
public DataSet GetData()
{
    DataSet authorDataSet;
    using (SqlConnection pubConnection = new SqlConnection())
```

```

    {
        pubConnection.ConnectionString = _connString;
        using(SqlCommand selectCommand = new SqlCommand())
        {
            selectCommand.CommandText =
                "Select au_id, au_lname,au_fname from authors";
            selectCommand.Connection = pubConnection;
            using (SqlDataAdapter pubDataAdapter = new SqlDataAdapter())
            {
                pubDataAdapter.SelectCommand = selectCommand;
                authorDataSet = new DataSet();
                pubDataAdapter.Fill(authorDataSet, "Author");
            }
        }
        return authorDataSet;
    }
}

```

7. Build the project and fix any errors.
8. Add the controls listed in Table 10-3 to Form1 and set the properties as shown.

Table 10-3. Form1 Controls

Control	Property	Value
DataGridView	Name	dgvAuthors
	AllowUserToAddRows	False
	AllowUserToDeleteRows	False
	ReadOnly	False
Button	Name	btnGetData
	Text	Get Data
Button	Name	btnUpdate
	Text	Update

9. Open the Form1 class code file in the code editor. Declare a class-level DataSet object after the class declaration.

```

public partial class Form1 : Form
{
    private DataSet _pubDataSet;
}

```

10. Open Form1 in the Form Designer. Double-click on the Get Data button to open the button click event method in the code editor.
11. Add the following code to the btnGetData click event procedure, which will execute the GetData method defined in the Author class. This dataset is then loaded into the grid using the DataSource property.

```
private void btnGetData_Click(object sender, EventArgs e)
{
    Author author = new Author();
    _pubDataSet = author.GetData();
    dgvAuthors.DataSource = _pubDataSet.Tables["Author"];
}
```

12. Build the project and fix any errors. Once the project builds, run the project in debug mode and test the `GetData` method. You should see the grid filled with author information. After testing, stop the debugger.

Editing and Updating Data in a DataSet

To edit and update data in a `DataSet`, follow these steps:

1. Open the `Author` class code in the code editor.
2. Create a method of the `Author` class called `UpdateData` that will use the `Update` method of the `DataAdapter` object to pass updates made to the `DataSet` to the Pubs database.

```
public void UpdateData(DataSet changedData)
{
    using (SqlConnection pubConnection = new SqlConnection())
    {
        pubConnection.ConnectionString = _connString;
        using (SqlCommand updateCommand = new SqlCommand())
        {
            updateCommand.CommandText = "Update authors set au_lname = @au_lname, " +
                "au_fname = @au_fname where au_id = @au_id";
            updateCommand.Parameters.Add ("@au_id", SqlDbType.VarChar, 11, "au_id");
            updateCommand.Parameters.Add
                ("@au_lname", SqlDbType.VarChar, 40, "au_lname");
            updateCommand.Parameters.Add
                ("@au_fname", SqlDbType.VarChar, 40, "au_fname");
            updateCommand.Connection = pubConnection;
            using (SqlDataAdapter pubDataAdapter = new SqlDataAdapter())
            {
                pubDataAdapter.UpdateCommand = updateCommand;
                pubDataAdapter.Update(changedData, "Author");
            }
        }
    }
}
```

3. Build the project and fix any errors.
4. Open `Form1` in the Form Designer. Double-click on the Update Data button to open the button click event method in the code editor.
5. Add the following code to the `btnUpdate` click event procedure, which will execute the `UpdateData` method defined in the `Author` class. By using the `GetChanges` method of the `DataSet` object, only data that has changed is passed for updating.

```
private void btnUpdate_Click(object sender, EventArgs e)
{
    Author author = new Author();
    author.UpdateData(_pubDataSet.GetChanges());
}

```

- Build the project and fix any errors. Once the project builds, run the project in debug mode and test the Update method. First, click the Get Data button. Change the last name of several authors and click the Update button. Click the Get Data button again to retrieve the changed values back from the database. After testing, stop the debugger.

Establishing Relationships between Tables in a DataSet

To establish relationships between tables in a DataSet, follow these steps:

- Add a new class named StoreSales to the project.
- Open the StoreSales class code in the code editor. Add the following using statements at the top of the file:

```
using System.Data;
using System.Data.SqlClient;

```

- Add the following code to declare private class level variables for the connection string and DataSet.

```
class StoreSales
{
    string _connString = "Data Source=localhost;" +
        "Initial Catalog=pubs;Integrated Security=True";
}

```

- Create a method of the StoreSales class called GetData that will select store information and sales information and establish a relationship between them. This information is used to fill a DataSet and return it to the client.

```
public DataSet GetData()
{
    DataSet storeSalesDataSet;
    storeSalesDataSet = new DataSet();
    using (SqlConnection pubConnection = new SqlConnection(_connString))
    {
        //Fill store table
        using (SqlCommand storeCommand = new SqlCommand())
        {
            storeCommand.Connection = pubConnection;
            storeCommand.CommandText =
                "SELECT [stor_id],[stor_name],[city],[state] FROM [stores]";
            using (SqlDataAdapter storeDataAdapter = new SqlDataAdapter())
            {
                storeDataAdapter.SelectCommand = storeCommand;
                storeDataAdapter.Fill(storeSalesDataSet, "Stores");
            }
        }
    }
}

```

```

    }
    //Fill sales table command
    using (SqlCommand salesCommand = new SqlCommand())
    {
        salesCommand.Connection = pubConnection;
        salesCommand.CommandText =
            "SELECT [stor_id],[ord_num],[ord_date],[qty] FROM [sales]";
        using (SqlDataAdapter salesDataAdapter = new SqlDataAdapter())
        {
            salesDataAdapter.SelectCommand = salesCommand;
            salesDataAdapter.Fill(storeSalesDataSet, "Sales");
        }
    }
    //Create relationship between tables
    DataColumn Store_StoreIdColumn =
        storeSalesDataSet.Tables["Stores"].Columns["stor_id"];
    DataColumn Sales_StoreIdColumn =
        storeSalesDataSet.Tables["Sales"].Columns["stor_id"];
    DataRelation StoreSalesRelation = new DataRelation
        ("StoresToSales", Store_StoreIdColumn, Sales_StoreIdColumn);
    storeSalesDataSet.Relations.Add(StoreSalesRelation);

    return storeSalesDataSet;
}
}
}

```

5. Build the project and fix any errors.
6. Add a second form to the project. Add the controls listed in Table 10-4 to Form2 and set the properties as shown.

Table 10-4. Form2 Controls

Control	Property	Value
DataGridView	Name	dgvStores
DataGridView	Name	dgvSales
Button	Name	btnGetData
	Text	Get Data

7. Open the Form2 class code file in the code editor. Declare a class-level DataSet object after the class declaration.

```

public partial class Form2 : Form
{
    DataSet StoreSalesDataSet;
}

```

8. Open Form2 in the Form Designer. Double-click on the Get Data button to open the button click event method in the code editor.
9. Add the following code to the btnGetData click event procedure, which will execute the GetData method defined in the StoreSales class. This Stores table is then loaded into the Stores grid using the DataSource property. Setting the DataMember property of the Sales grid loads it with the sales data of the store selected in the Stores grid.

```
private void btnGetData_Click(object sender, EventArgs e)
{
    StoreSales storeSales = new StoreSales();
    StoreSalesDataSet = storeSales.GetData();
    dgvStores.DataSource = StoreSalesDataSet.Tables["Stores"];
    dgvSales.DataSource = StoreSalesDataSet.Tables["Stores"];
    dgvSales.DataMember = "StoresToSales";
}
```

10. Open the Program class in the code editor. Change the code to launch Form2 when the form loads.

```
Application.Run(new Form2());
```

11. When the form loads, click the Get Data button to load the grids. Selecting a new row in the Stores grid should update the Sales grid to show the store's sales. When you are finished testing, stop the debugger and exit Visual Studio.

Working with the Entity Framework

The Entity Framework (EF) is an Object-Relational Mapping (ORM) technology built into ADO.NET. EF eliminates the mismatch between the object-oriented programming constructs of the .NET language and the relational data constructs of the database system. For example, to load and work with a customer object, a developer has to send a SQL string to the database engine. This requires the developer to be familiar with the relational schema of the data. In addition, the SQL is hardcoded into the application and the application is not shielded from changes in the underlying schema. Another disadvantage is that since the application sends the SQL statements as a string to the database engine for processing, Visual Studio can't implement syntax checking and issue warnings and build errors to help with programmer productivity.

The Entity Framework provides the mapping schema that allows programmers to work at a higher level of abstraction. They can write code using object-oriented constructs to query and load the entities (objects defined by classes). The mapping schema translates the queries against the entities into the required database-specific language needed to perform CRUD (create, read, update, and delete) operations against the data.

In order to use the Entity Framework in your application, you must first add an ADO.NET Entity Data Model to your application. This step launches the Entity Data Model Wizard, which allows you to develop your model from scratch or generate it from an existing database. Choosing to generate it from an existing database allows you to create a connection to the database and select the tables, views, and stored procedures you want to include in the model. The output of the wizard is an .edmx file. This file is an XML-based file that has three sections. The first consists of

store schema definition language (SSDL); this describes the tables and relationships where the data is stored. The following code shows a portion of the SSDL for a data model generated from the Pubs database:

```
<EntityContainer Name="pubsModelStoreContainer">
  <EntitySet Name="sales" EntityType="pubsModel.Store.sales"
    store:Type="Tables" Schema="dbo" />
  <EntitySet Name="stores" EntityType="pubsModel.Store.stores"
    store:Type="Tables" Schema="dbo" />
  <AssociationSet Name="FK_sales_stor_id_1273C1CD"
    Association="pubsModel.Store.FK_sales_stor_id_1273C1CD">
    <End Role="stores" EntitySet="stores" />
    <End Role="sales" EntitySet="sales" />
  </AssociationSet>
</EntityContainer>
<EntityType Name="sales">
  <Key>
    <PropertyRef Name="stor_id" />
    <PropertyRef Name="ord_num" />
    <PropertyRef Name="title_id" />
  </Key>
  <Property Name="stor_id" Type="char" Nullable="false" MaxLength="4" />
  <Property Name="ord_num" Type="varchar" Nullable="false" MaxLength="20" />
  <Property Name="ord_date" Type="datetime" Nullable="false" />
  <Property Name="qty" Type="smallint" Nullable="false" />
  <Property Name="payterms" Type="varchar" Nullable="false" MaxLength="12" />
  <Property Name="title_id" Type="varchar" Nullable="false" MaxLength="6" />
</EntityType>
```

The second section consists of conceptual schema definition language (CSDL); it specifies the entities and relationships between them. These entities are used to work with data in the application. The following code comes from the CSDL section of a data model generated from the Pubs database:

```
<EntityContainer Name="pubsEntities" annotation:LazyLoadingEnabled="true">
  <EntitySet Name="sales" EntityType="pubsModel.sale" />
  <EntitySet Name="stores" EntityType="pubsModel.store" />
  <AssociationSet Name="FK_sales_stor_id_1273C1CD"
    Association="pubsModel.FK_sales_stor_id_1273C1CD">
    <End Role="stores" EntitySet="stores" />
    <End Role="sales" EntitySet="sales" />
  </AssociationSet>
</EntityContainer>
<EntityType Name="sale">
  <Key>
    <PropertyRef Name="stor_id" />
    <PropertyRef Name="ord_num" />
    <PropertyRef Name="title_id" />
  </Key>
```

```

<Property Name="stor_id" Type="String" Nullable="false"
    MaxLength="4" Unicode="false" FixedLength="true" />
<Property Name="ord_num" Type="String" Nullable="false"
    MaxLength="20" Unicode="false" FixedLength="false" />
<Property Name="ord_date" Type="DateTime" Nullable="false" />
<Property Name="qty" Type="Int16" Nullable="false" />
<Property Name="payterms" Type="String" Nullable="false"
    MaxLength="12" Unicode="false" FixedLength="false" />
<Property Name="title_id" Type="String" Nullable="false"
    MaxLength="6" Unicode="false" FixedLength="false" />
<NavigationProperty Name="store"
    Relationship="pubsModel.FK_sales_stor_id_1273C1CD"
    FromRole="sales" ToRole="stores" />
</EntityType>

```

The final section of the .edmx file consists of code written in the mapping specification language (MSL). The MSL maps the conceptual model to the storage model. The following code shows a portion of the MSL section of a data model generated from the Pubs database:

```

<EntityContainerMapping StorageEntityContainer="pubsModelStoreContainer"
    CdmEntityContainer="pubsEntities">
  <EntitySetMapping Name="sales"><EntityTypeMapping TypeName="pubsModel.sale">
    <MappingFragment StoreEntitySet="sales">
      <ScalarProperty Name="stor_id" ColumnName="stor_id" />
      <ScalarProperty Name="ord_num" ColumnName="ord_num" />
      <ScalarProperty Name="ord_date" ColumnName="ord_date" />
      <ScalarProperty Name="qty" ColumnName="qty" />
      <ScalarProperty Name="payterms" ColumnName="payterms" />
      <ScalarProperty Name="title_id" ColumnName="title_id" />
    </MappingFragment></EntityTypeMapping></EntitySetMapping>

```

Visual Studio provides a visual designer to work with the entity model. Using this designer you can update the entities, add associations, and implement inheritance. Changes made in the visual model are translated back to the .edmx file which is updated accordingly. Figure 10-1 shows entities in the visual designer that were generated from the pubs database.

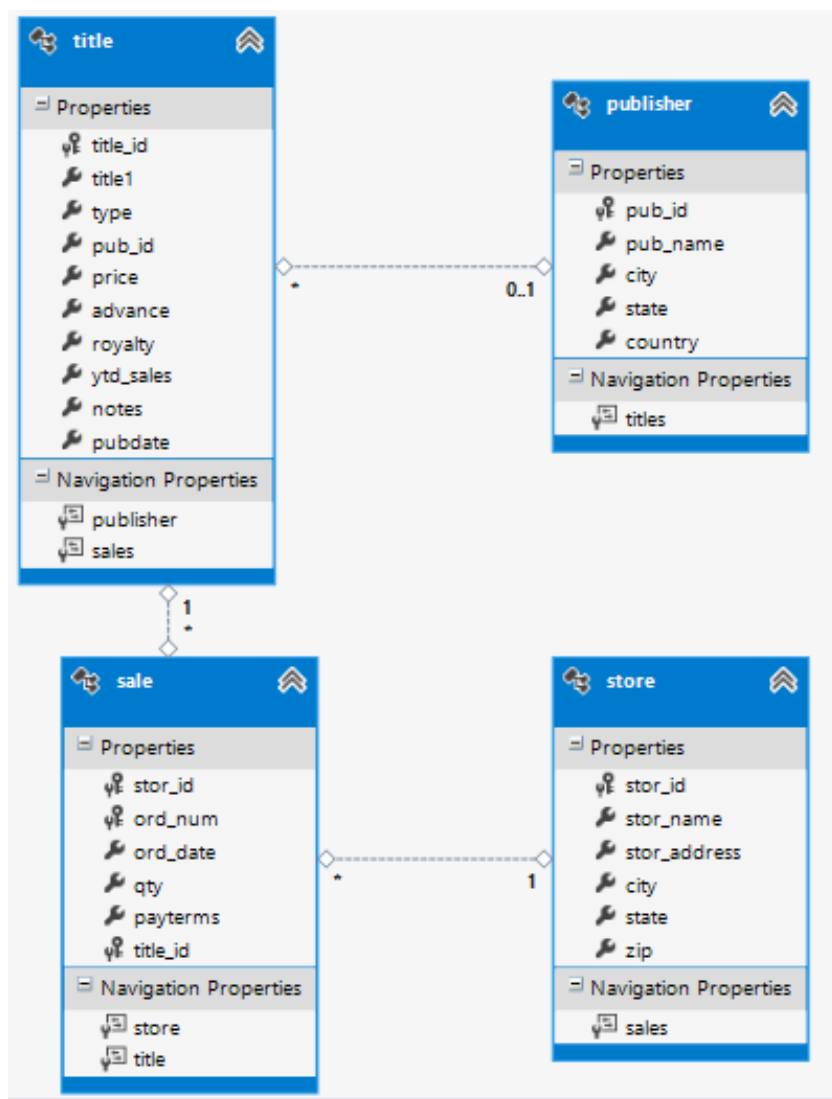


Figure 10-1. Entities in the entity model designer

Querying Entities with LINQ to EF

When creating the ADO.NET entity data model using the Entity Data Model Wizard, an `ObjectContext` class is created that represents the entity container defined in the model. The `ObjectContext` class supports CRUD-based queries against the entity model. Queries written against the `ObjectContext` class are written using LINQ to EF. As noted previously, LINQ stands for Language-Integrated Query. LINQ allows developers to write queries in C# syntax, which, when executed, are converted to the query syntax of the data provider. Once the query is executed and data is returned, the Entity Framework converts the results back to the entity object model.

The following code uses the `Select` method to return all the rows from the `Stores` table and return the results as a list of `Store` entities. The store names are then written to the console window.

```
var context = new pubsEntities();
var query = from s in context.stores
            select s;
var stores = query.ToList();
foreach (store s in stores)
{
    Console.WriteLine(s.stor_name);
}
Console.ReadLine();
```

LINQ to EF provides a rich set of query operations, including filtering, ordering, and grouping operations. The following code demonstrates filtering stores by state:

```
var context = new pubsEntities();
var query = from s in context.stores
            where s.state == "WA"
            select s;
var stores = query.ToList();
```

The following code selects sales entities that have ordered more than 25 objects and then orders them by descending date:

```
var context = new pubsEntities();
var query = from s in context.sales
            where s.qty > 25
            orderby s.ord_date descending
            select s;
var sales = query.ToList();
```

Since the Entity Framework includes navigation properties between entities, you can easily build complex queries based on related entities. The following query selects stores with more than five sales orders:

```
var context = new pubsEntities();
var query = from s in context.stores
            where s.sales.Count > 5
            select s;
var stores = query.ToList();
```

■ **Note** For more information on the LINQ query language, refer to the MSDN library at <http://msdn.microsoft.com>.

Updating Entities with the Entity Framework

The Entity Framework tracks changes made to the entity types represented in the Context object. You can add, update or delete entity objects. When you are ready to persist the changes back to the database, you call the `SaveChanges` method of the context object. The EF creates and executes the insert, update, or delete statements against the database. You can also explicitly map stored procedures to implement the database commands. The following code selects a store using the store ID, updates the store name, and sends it back to the database:

```
var context = new pubsEntities();
var store = (from s in context.stores
            where s.stor_id == storeId
            select s).First();
store.stor_name = "DRC Books";
context.SaveChanges();
```

ACTIVITY 10-3. RETRIEVING DATA WITH THE ENTITY FRAMEWORK

In this activity, you will become familiar with the following:

- creating an Entity Data Model
- executing queries using LINQ to EF

Creating an Entity Data Model

To create an entity data model, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose Console Application. Rename the project to Activity10_3 and click the OK button.
3. Right-click on the project node in solution explorer and select Add ► New Item.
4. Under the Data node in the Add New Item window, select an ADO.NET Entity Data Model. Name the model Pubs.edmx and click Add.
5. In the Choose Model Contents screen, select the Generate from database and click Next.
6. In the Choose Your Data Connection screen, create a connection to the Pubs database and choose Next. (See Figure 10-2)

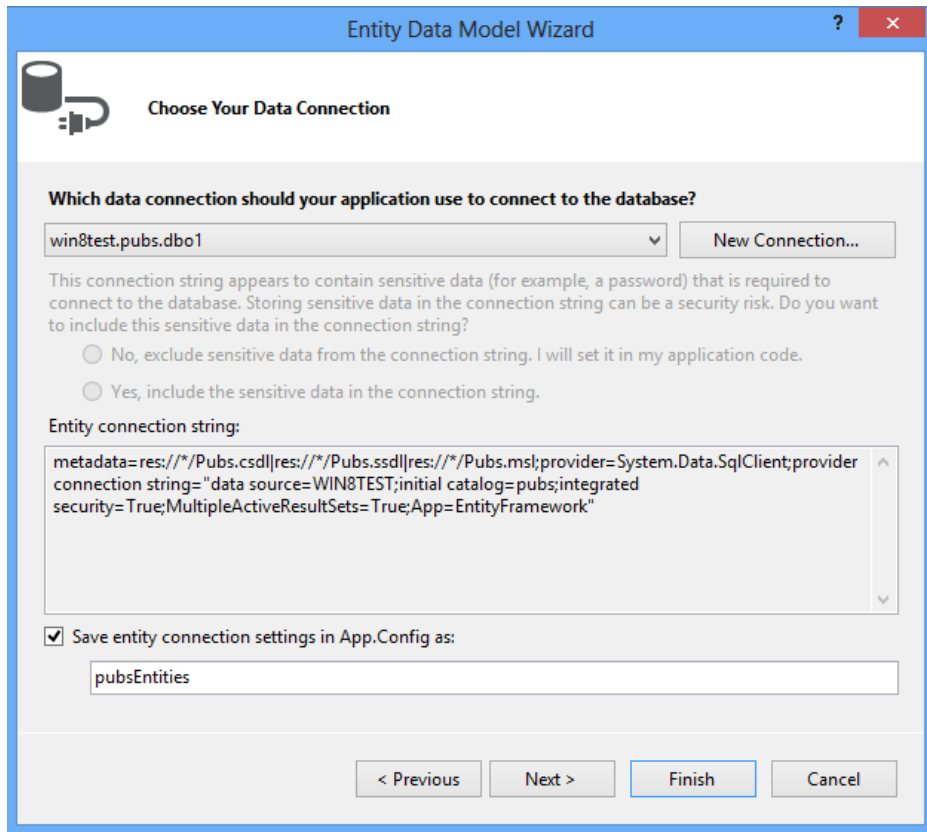


Figure 10-2. Creating a database connection with the Entity Data Model Wizard

7. In the Choose Your Database Objects screen, expand the Tables node and select the Sales, Stores, and Titles tables, as shown in Figure 10-3. Click Finish.

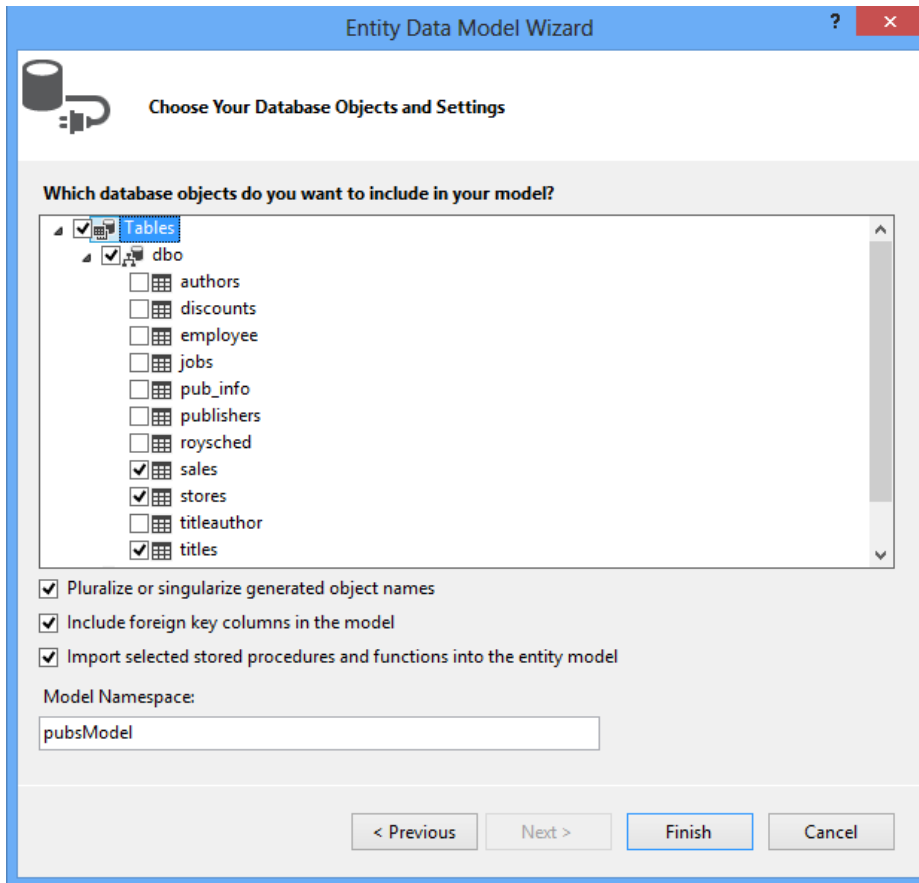


Figure 10-3. Selecting database objects for an Entity Data Model

8. You are presented with the Entity Model Designer containing the sales, store, and title entities, as shown in Figure 10-4.

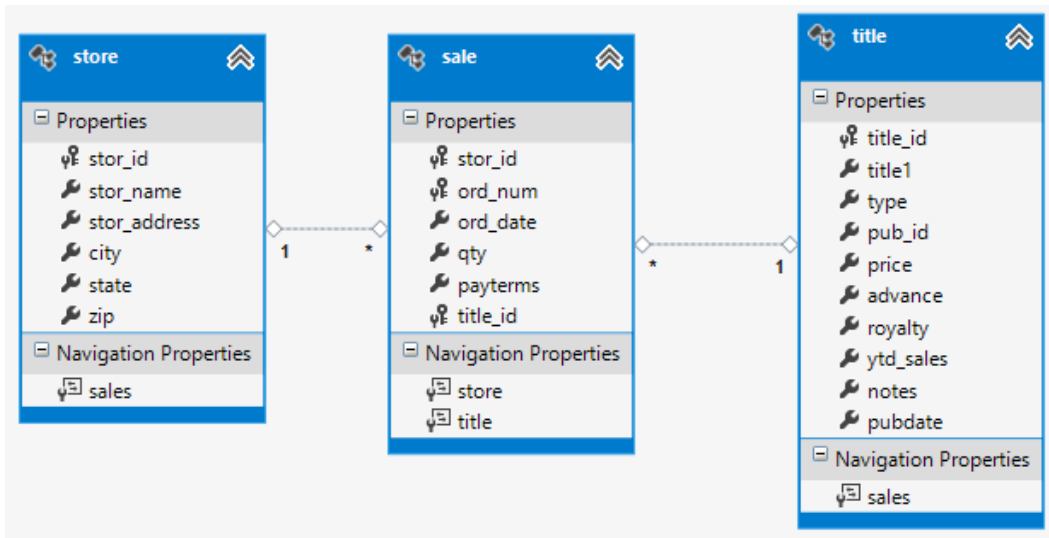


Figure 10-4. Entity Model Designer

- In the Entity Model Designer right-click on the `title` entity and select `rename`. Rename it to `book`. In the `book` entity, rename the `title1` property to `title`.

Querying an Entity Data Model

To query this entity data model using LINQ, follow these steps:

- Open the `Program.cs` file in the Code Editor Window.
- Add the following method to select the `book` entities and write their titles to the Console window:

```
private static void GetTitles()
{
    var context = new pubsEntities();
    var query = from b in context.books select b;
    var books = query.ToList();
    foreach (book b in books)
    {
        Console.WriteLine(b.title);
    }
    Console.ReadLine();
}
```

- Call the `GetTitles` method from the `Main` method.

```
static void Main(string[] args)
{
    GetTitles();
}
```


4. Run the program in debug mode. You should see the titles listed in the Console window. When you are done testing, stop the debugger.
5. Add the following method that gets books in the 10 to 20 dollar range and orders them by price:

```
private static void GetTitlesByPrice()
{
    var context = new pubsEntities();
    var query = from b in context.books
                where b.price >= (decimal)10.00
                   && b.price <= (decimal)20.00
                orderby b.price
                select b;
    var books = query.ToList();
    foreach (book b in books)
    {
        Console.WriteLine(b.price + " -- " + b.title);
    }
    Console.ReadLine();
}
```

6. Call the GetTitlesByPrice method from the Main method.

```
static void Main(string[] args)
{
    //GetTitles();
    GetTitlesByPrice();
}
```

7. Run the program in debug mode. You should see the titles and prices listed in the Console window. When you are done testing, stop the debugger.
8. Add the following method to list the book titles and the sum of their sales amount. Notice that this query gets the sales amount by adding up the book's related sales entities.

```
private static void GetBooksSold()
{
    var context = new pubsEntities();
    var query = from b in context.books
                select new
                {
                    BookID = b.title_id,
                    TotalSold = b.sales.Sum(s =>(int?) s.qty)
                };
    foreach (var item in query)
    {
        Console.WriteLine(item.BookID + " -- " + item.TotalSold);
    }
    Console.ReadLine();
}
```

9. Call the `GetBooksSold` method from the `Main` method.

```
static void Main(string[] args)
{
    //GetTitles();
    //GetTitlesByPrice();
    GetBooksSold();
}
```

10. Run the program in debug mode. You should see the book IDs and amount sold listed in the Console window. When you are done testing, stop the debugger and exit Visual Studio.
-

Summary

This chapter is the first in a series that will show you how to build the various tiers of an OOP application. To implement an application's data access layer, you learned about ADO.NET and the classes used to work with relational data sources. You looked at the various classes that make up the `System.Data.SqlClient` namespace; these classes retrieve and update data stored in a SQL Server database. You also examined the `System.Data` namespace classes that work with disconnected data. In addition, you were exposed to the Entity Framework and LINQ and saw how they allow you to query the data using OOP constructs. You wrote queries in terms of entities, and the framework translated the queries into the query syntax of the datasource, retrieved the data, and loaded the entities.

In Chapter 11, you will look at implementing the user interface (UI) tier of a Windows application. Along the way, you will take a closer look at the classes and namespaces of the .NET Framework used to create rich Windows-based user interfaces.



Developing WPF Applications

In Chapter 10, you learned how to build the data access layer of an application. To implement its logic, you used the classes of the `System.Data` namespace. These classes retrieve and work with relational data, which is a common requirement of many business applications. You are now ready to look at how users will interact with your application. Users interact with an application through the user interface layer. This layer, in turn, interacts with the business logic layer, which, in turn, interacts with the data access layer. In this chapter, you will learn how to build a user interface layer with the .NET Windows Presentation Foundation (WPF). WPF is commonly used to develop desktop business productivity applications running on Windows 7 and above. Business productivity applications target business users who need to query and update data stored in a backend database. It consists of a comprehensive set of application-development features that include Extensible Application Markup Language (XAML), controls, data binding, and layout. The concepts discussed in this chapter will also lay a foundation for Chapter 13, in which we look at creating the new Windows Store app user interface for Windows 8. These applications use a similar model for creating interfaces with XAML and data binding controls.

After reading this chapter, you will be comfortable performing the following tasks:

- using XAML markup to design a user interface
- working with layout controls
- working with display controls
- responding to control events
- using data binding controls
- creating and using control templates

Windows Fundamentals

Windows are objects with a visual interface that are painted on the screen to provide users a way to interact with programs. Like most objects you work with in object-oriented languages, .NET windows expose properties, methods, and events. A window's properties define its appearance. Its `Background` property, for example, determines its color. The methods of a window define its behaviors. For example, calling its `Hide` method hides it from the user. A window's events define interactions with the user (or other objects). You can use the `MouseDown` event, for example, to initiate an action when the user clicks the right mouse button on the window.

Controls are components with visual interfaces that give users a way to interact with the program. A window is a special type of control, called a container control, which hosts other controls. You can place many different types of controls on windows. Some common controls used on windows are `TextBoxes`, `Labels`, `OptionButtons`, `ListBoxes`, and `CheckBoxes`. In addition to the controls provided by the .NET Framework, you can also create your own custom controls or purchase controls from third-party vendors.

Introducing XAML

WPF user interfaces are built using a declarative markup language called XAML. XAML declares the controls that will make up the interface. An opening angle bracket (<) followed by the name of the control type and a closing bracket define the control. For example, the following markup defines a button control inside a Grid.

```
<Grid>
  <Button/>
</Grid>
```

Notice the Grid needs a formal closing tag because it contains the Button control. Since the Button control does not contain any other elements, you can use a forward slash (/) in front of the end bracket to close it.

The next step is to define the properties of the controls. For example, you may want to set the background color of the button to red and write some text on it. The properties of the control are set by using attribute syntax, which consists of the property name followed by an equal sign and the attribute value in quotation marks. The following markup shows the Button control with some attributes added:

```
<Grid>
  <Button Content="Click Me" Background="Red"/>
</Grid>
```

For some properties of an object element a syntax known as property element syntax is used. The syntax for the property element start tag is <typeName.propertyName>. For example, you can create rows and columns in the layout grid to control placement of controls in the grid, as shown:

```
<Grid.ColumnDefinitions>
  <ColumnDefinition Width="100" />
  <ColumnDefinition Width="*" />
</Grid.ColumnDefinitions>
<Grid.RowDefinitions>
  <RowDefinition Height="25" />
  <RowDefinition Height="25" />
  <RowDefinition Height="25" />
</Grid.RowDefinitions>
```

Controls are positioned in the grid by including a Grid.Row and Grid.Column attribute, as shown:

```
<Label Grid.Column="0" Grid.Row="0" Content="Name:" />
<Label Grid.Column="0" Grid.Row="1" Content="Password:" />
<TextBox Name="txtName" Grid.Column="1" Grid.Row="0"/>
<TextBox Name="txtPassword" Grid.Column="1" Grid.Row="1"/>
<Button Grid.Column="1" Grid.Row="3"
Content="Click Me" HorizontalAlignment="Right"
MinWidth="80" Background="Red"/>
```

Figure 11-1 shows the window with two textboxes created by the previous XAML code.




Figure 11-1. A window created with XAML

Using Layout Controls

Although you can use fixed positioning to place controls on a WPF window, it's not recommended. Using fixed positioning usually works well for a fixed resolution size but it doesn't scale well to different resolutions and devices. To overcome the limitations of fixed positioning, WPF offers several layout controls. A layout control allows you to position other controls within it using a relative positioning format. One of the main layout controls for positioning other controls is the Grid. As seen previously, a Grid control contains columns and rows to control the placement of its child controls. The height and width of the columns and rows can be set to a fixed value, auto, or *. The auto setting takes up as much space as needed by the contained control. The * setting takes up as much space as is available. The Grid control is often used to lay out data entry forms. The following code lays out a simple data entry form used to collect user information. The resulting form (in the Visual Studio designer) is shown in Figure 11-2.

```
<Grid>
  <Grid.RowDefinitions>
    <RowDefinition Height="Auto" />
    <RowDefinition Height="Auto" />
    <RowDefinition Height="Auto" />
    <RowDefinition Height="Auto" />
    <RowDefinition Height="28" />
    <RowDefinition Height="*" />
  </Grid.RowDefinitions>
  <Grid.ColumnDefinitions>
    <ColumnDefinition Width="Auto" />
    <ColumnDefinition Width="200" />
    <ColumnDefinition Width="*" />
  </Grid.ColumnDefinitions>
  <Label Grid.Row="0" Grid.Column="0" Content="Name:" />
  <Label Grid.Row="1" Grid.Column="0" Content="Old Password:" />
  <Label Grid.Row="2" Grid.Column="0" Content="New Password:" />
  <Label Grid.Row="3" Grid.Column="0" Content="Confirm Password:" />
  <TextBox Grid.Column="1" Grid.Row="0" Margin="3" />
  <TextBox Grid.Column="1" Grid.Row="1" Margin="3" />
  <TextBox Grid.Column="1" Grid.Row="2" Margin="3" />
  <TextBox Grid.Column="1" Grid.Row="3" Margin="3" />
  <Button Grid.Column="1" Grid.Row="4" HorizontalAlignment="Right"
    MinWidth="80" Margin="0,0,0,8" Content="Submit" />
</Grid>
```

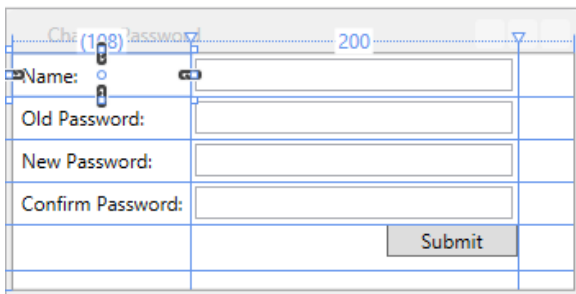


Figure 11-2. Input form window

Another useful layout control is the `StackPanel`. It lays out child controls either vertically or horizontally depending on the orientation setting. The following code shows two buttons in a `StackPanel` control:

```
<StackPanel Grid.Column="1" Grid.Row="4" Orientation="Horizontal" >
  <Button MinWidth="80" Margin="0,0,0,8" Content="Submit" />
  <Button MinWidth="80" Margin="0,0,0,8" Content="Cancel" />
</StackPanel>
```

Some other layout controls available are the `DockPanel`, `WrapPanel`, and `Canvas`. The `DockPanel` is used to provide docking of elements to the left, right, top, bottom, or center of the panel. The `WrapPanel` acts like a `StackPanel` but will wrap child controls to a new line if no room is left. The `Canvas` control is used to lay out its child elements with absolute positioning relative to one of its sides. It is typically used for graphics elements and not to lay out user interface controls.

Adding Display Controls

The goal of most business applications is to present data to users and allow them to update the data and save it back to a database. Some common controls used to facilitate this process are the `TextBox`, `ListBox`, `ComboBox`, `CheckBox`, `DatePicker`, and `DataGrid`. You have already seen the `TextBox` used on a window; the following code shows how to add a `ListBox` and `ComboBox` to a window. Figure 11-3 shows how the window is rendered.

```
<Grid>
  <Grid.ColumnDefinitions>
    <ColumnDefinition Width="*" />
    <ColumnDefinition Width="*" />
  </Grid.ColumnDefinitions>
  <ListBox Margin="20" Grid.Column="0">
    <ListBoxItem>Red</ListBoxItem>
    <ListBoxItem>Blue</ListBoxItem>
    <ListBoxItem>Green</ListBoxItem>
    <ListBoxItem>Yellow</ListBoxItem>
  </ListBox>
  <ComboBox Grid.Column="1" VerticalAlignment="Top">
    <ComboBoxItem>Small</ComboBoxItem>
    <ComboBoxItem>Medium</ComboBoxItem>
```

```

        <ComboBoxItem>Large</ComboBoxItem>
        <ComboBoxItem>X-Large</ComboBoxItem>
    </ComboBox>
</Grid>

```

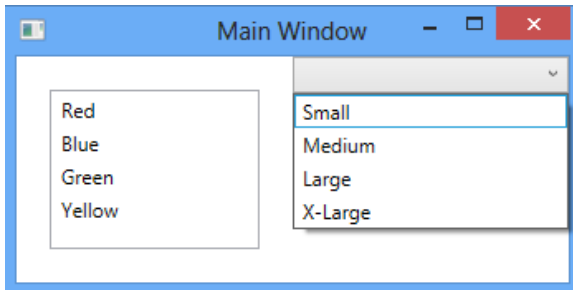


Figure 11-3. Window containing a Listbox and ComboBox

Although you can code the items displayed in these controls directly in the XAML markup, it is more likely that you will use data binding to display their values. You'll look at data binding shortly.

Using the Visual Studio Designer

Even though it's quite possible to create your window entirely through code using a text editor, you would probably find this process quite tedious and not a very productive use of your time. Thankfully, the Visual Studio IDE includes an excellent designer for creating your WPF windows. Using the designer, you can drag and drop controls from the Toolbox to the Visual Studio designer, set its properties using the Visual Studio Properties window, and get the benefits of auto completion and syntax checking as you enter code using the XAML editor. Figure 11-4 shows a window in the Visual Studio designer.

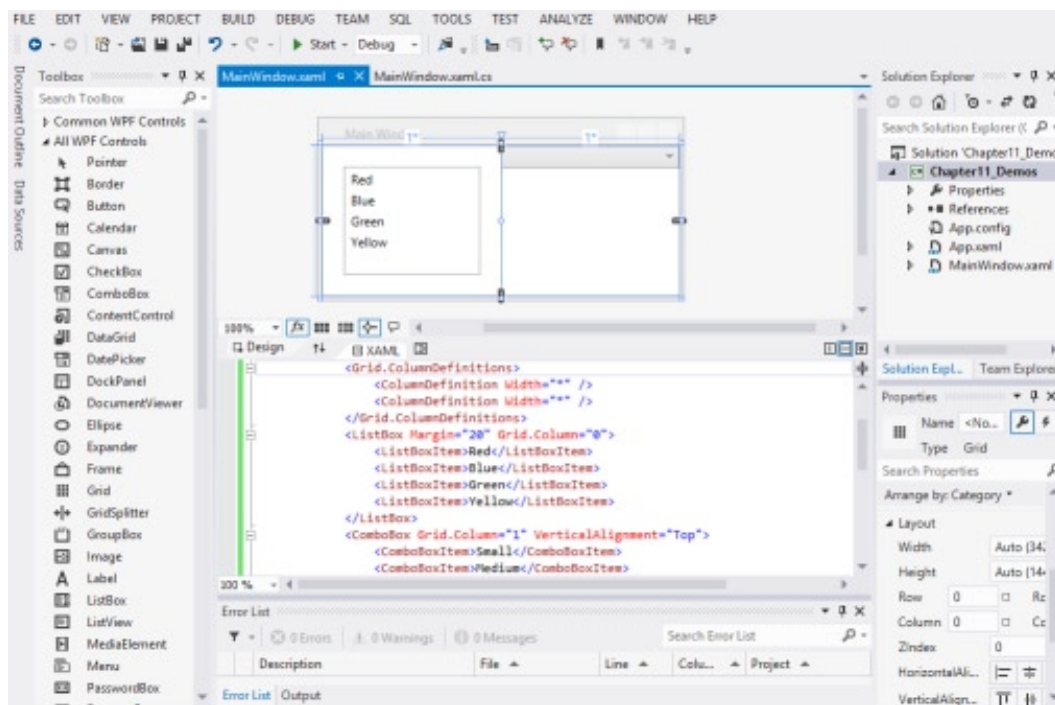


Figure 11-4. Designing a window in Visual Studio

Handling Control Events

Windows graphical user interface (GUI) programs are event-driven. Events are actions initiated by either a user or the system, whenever a user clicks a button, for example, or a `SqlConnection` object issues a `StateChange` event. Event-driven applications respond to the various events that occur by executing code that you specify. To respond to an event, you define the event handler to execute when a particular event occurs. As you saw in Chapter 8, the .NET Framework uses delegation to bind an event, with the event handler procedures written to respond to the event. A delegation object maintains an invocation list of methods that have subscribed to receive notification when the event occurs. When an event occurs—for example, a button is clicked—the control will raise the event by invoking the delegate for the event, which in turn will call the event handler methods that have subscribed to receive the event notification. Although this sounds complicated, the framework classes do most of the work for you.

In Visual Studio, you can add an event to a WPF control either by writing XAML code or by selecting it in the control's Properties window. Figure 11-5 shows wiring up an event handler in the XAML Editor window; Figure 11-6 shows wiring up an event handler using the Events tab of the Properties window. Note that when working with controls in code, you need to give them a unique name using the `Name` attribute.

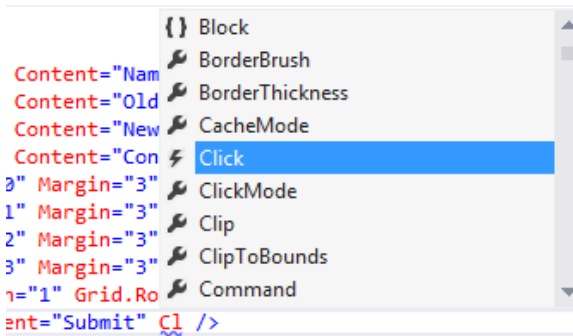


Figure 11-5. Wiring up an event handler in the XAML editor

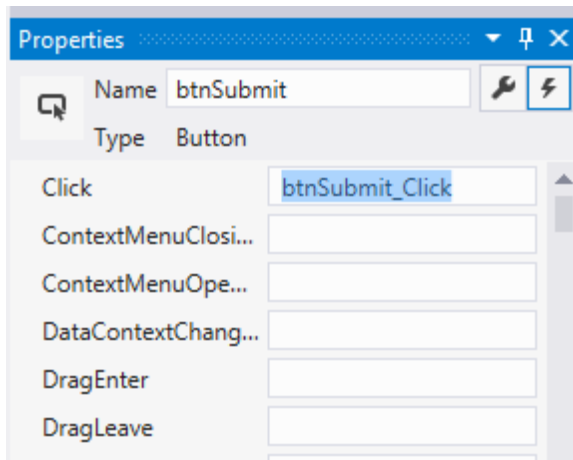


Figure 11-6. Wiring up an event handler in the Properties window

The following code shows the event handler method inserted in the code behind file for the button click event:

```
private void btnSubmit_Click(object sender, RoutedEventArgs e)
{
}

```

By convention, the name of the event handler method begins with the name of the object issuing the event followed by an underscore (`_`) and the name of the event. The actual name of the event handler, however, is unimportant. The `Click` attribute in the XAML code adds this method to the invocation list of the event's delegation object.

All event handlers must provide two parameters, which are passed to the method when the event is fired. The first parameter is the sender, which represents the object that initiated the event. The second parameter, of type `System.Windows.RoutedEventArgs`, is an object used to pass any information specific to the particular event.

Because the .NET Framework uses delegates for event notification, you can use the same method to handle more than one event, provided the events have the same signature. For example, you could handle a button click event and a menu click event with the same event handler, but not a text box key press event, because it has a different signature.

The following code demonstrates how to handle the button click event of two buttons that use the same handler method. The sender parameter is cast as a Button type and interrogated to determine which button fired the event.

```
private void Button_Click(object sender, RoutedEventArgs e)
{
    Button btn = (Button)sender;
    if (btn.Name == "btnCancel")
    {
        //Cancel code goes here
    }
    else if (btn.Name == "btnSubmit")
    {
        //Submit code goes here
    }
}
```

In the following activity, you will work with forms and controls to construct a simple memo viewer application that will allow users to load and view memo documents.

ACTIVITY 11-1. WORKING WITH WINDOWS AND CONTROLS

In this activity, you will become familiar with the following:

- creating a Windows Form-based GUI application
- working with Menu, StatusStrip, and Dialog controls
- working with Control events

Creating the Memo Viewer Interface

To create the memo viewer interface, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose a WPF Application under the C# Projects folder. Rename the project to Activity11_1 and click the OK button.
3. The project contains a MainWindow.xaml file. This file is where you design the user interface. The project also contains a MainWindow.xaml.cs file. This is the codebehind file and it is where you will add the code to respond to the events. If not already open, open the MainWindow.xaml file in the XAML Editor Window.
4. In the Window tag of the XAMLmarkup, add a Name attribute with a value of "MemoViewer". Change the Title attribute to "Memo Viewer".

```
<Window x:Name="Memoviewer" x:Class="Act11_1.MainWindow"
        xmlns="http://schemas.microsoft.com/winfx/2006/xaml/presentation"
        xmlns:x="http://schemas.microsoft.com/winfx/2006/xaml"
        Title="Memo Viewer" Height="350" Width="525">
```

5. Add a DockPanel control in the Grid control.

```
<Grid>
  <DockPanel LastChildFill="True">
    </DockPanel>
</Grid>
```

6. Add a Menu control inside the DockPanel and dock it to the top using the following XAML:

```
<DockPanel LastChildFill="True">
  <Menu DockPanel.Dock="Top">
    <MenuItem Header="_ File">
      <MenuItem Name="mnuNew" Header="_New..." />
      <Separator />
      <MenuItem Name="mnuOpen" Header="_Open..." />
      <Separator />
      <MenuItem Name="mnuSave" Header="_Save" />
      <MenuItem Name="mnuSaveAs" Header="_Save As..." />
      <Separator />
      <MenuItem Name="mnuExit" Header="_Exit" />
    </MenuItem>
    <MenuItem Header="_ Edit">
      <MenuItem Header="_ Cut..." />
      <MenuItem Header="_ Copy..." />
      <MenuItem Header="_ Paste" />
    </MenuItem>
  </Menu>
</DockPanel>
```

7. Add a StatusBar control by inserting the following code between the ending Menu tag and the ending DockPanel tag. Note that you are using a Grid control inside the StatusBar control to lay out the items in the StatusBar.

```
<StatusBar DockPanel.Dock="Bottom">
  <Grid>
    <Grid.RowDefinitions>
      <RowDefinition Height="*" />
    </Grid.RowDefinitions>
    <Grid.ColumnDefinitions>
      <ColumnDefinition Width="4*" />
      <ColumnDefinition Width="*" />
    </Grid.ColumnDefinitions>
  </Grid>
  <StatusBarItem Grid.Column="0" HorizontalAlignment="Left">
    <TextBlock Name="sbTextbox1">File Name</TextBlock>
  </StatusBarItem>
  <StatusBarItem Grid.Column="1" HorizontalAlignment="Right">
    <TextBlock Name="sbTextbox2">Date</TextBlock>
  </StatusBarItem>
</StatusBar>
```

8. Add a RichTextBox control after the StatusBar end tag and before the DockPanel end tag.

```

    </StatusBar>
    <RichTextBox Name="rtbMemo" />
</DockPanel>

```

9. Note that as you add the XAML, the Visual Designer updates the appearance of the window. The Memo Viewer window should look similar to the one shown Figure 11-7.




Figure 11-7. The completed MemoEditor window

10. Build the solution. If there are any errors, fix them and rebuild.

Coding the Control Events

To code the control events, follow these steps:

1. In the XAML Editor window, add the Loaded event attribute to the Window, as shown:

```

<Window x:Class="Act11_1.MainWindow"
  xmlns="http://schemas.microsoft.com/winfx/2006/xaml/presentation"
  xmlns:x="http://schemas.microsoft.com/winfx/2006/xaml"
  Name="MemoViewer" Title="Memo Viewer" Height="350" Width="525"
  Loaded="MemoViewer_Loaded">

```

- Open the codebehind file by right-clicking the XAML code editor and selecting View Code. Add the following code to the MemoViewer_Loaded event handler. When the window loads, it should show the message on the left side of the StatusPanel and the date on the right.

```
private void MemoViewer_Loaded(object sender, RoutedEventArgs e)
{
    sbTextbox1.Text = "Ready to load file";
    sbTextbox2.Text = DateTime.Today.ToShortDateString();
}
```

- In the XAML editor, add the Click event to the mnuOpen control.

```
<MenuItem Name="mnuOpen" Header="_Open..."
    Click="mnuOpen_Click"/>
```

- In the Code Editor window of the codebehind file, add the following code to the menu click event. This code configures and launches an Open File Dialog box, which returns the file path. The file path is then passed to a FileStream object, which loads the file into the RichTextBox. The file path is also loaded into the StatusBar TextBox.

```
private void mnuOpen_Click(object sender, RoutedEventArgs e)
{
    // Configure open file dialog box
    Microsoft.Win32.OpenFileDialog dlg = new Microsoft.Win32.OpenFileDialog();
    dlg.FileName = "Document"; // Default file name
    dlg.DefaultExt = ".txt"; // Default file extension
    dlg.Filter = "Text documents (.txt)|*.txt"; // Filter files by extension
    // Show open file dialog box
    Nullable<bool> result = dlg.ShowDialog();
    // Process open file dialog box results
    if (result == true)
    {
        // Open document and load RichTextBox
        string fileName = dlg.FileName;
        TextRange range;
        System.IO.FileStream fStream;
        if (System.IO.File.Exists(fileName))
        {
            range = new TextRange(rtbMemo.Document.ContentStart,
                rtbMemo.Document.ContentEnd);
            fStream = new System.IO.FileStream
                (fileName, System.IO.FileMode.OpenOrCreate);
            range.Load(fStream, System.Windows.DataFormats.Text );
            fStream.Close();
        }
        sbTextbox1.Text = fileName;
    }
}
```

5. Add a click event for the `mnuExit` control with the following code to close the window:

```
private void mnuExit_Click(object sender, RoutedEventArgs e)
{
    this.Close();
}
```

6. Build the solution and fix any errors.
7. Create a Memos folder on the C drive. Using Notepad, create a text file containing a test message. Save the file as `Test.txt`.
8. Select `Debug > Start`. Test the application by loading the `Test.txt` file. After viewing the file, close the window by clicking the Exit menu.
9. After testing the application, exit Visual Studio.

Creating and Using Dialog Boxes

Dialog boxes are special windows often used in Windows-based GUI applications to display or retrieve information from users. The difference between a normal window and a dialog box is that a dialog box is displayed modally. A modal window prevents the user from performing other tasks within the application until the dialog box has been dismissed. When you start a new project in Visual Studio, you are presented with a New Project dialog box, as shown in Figure 11-8. You can also use dialog boxes to present the user with critical information and query them for a response. For example, if you try to run an application in debug mode and a build error is encountered, the Visual Studio IDE presents you with a dialog box asking whether you want to continue (see Figure 11-9).

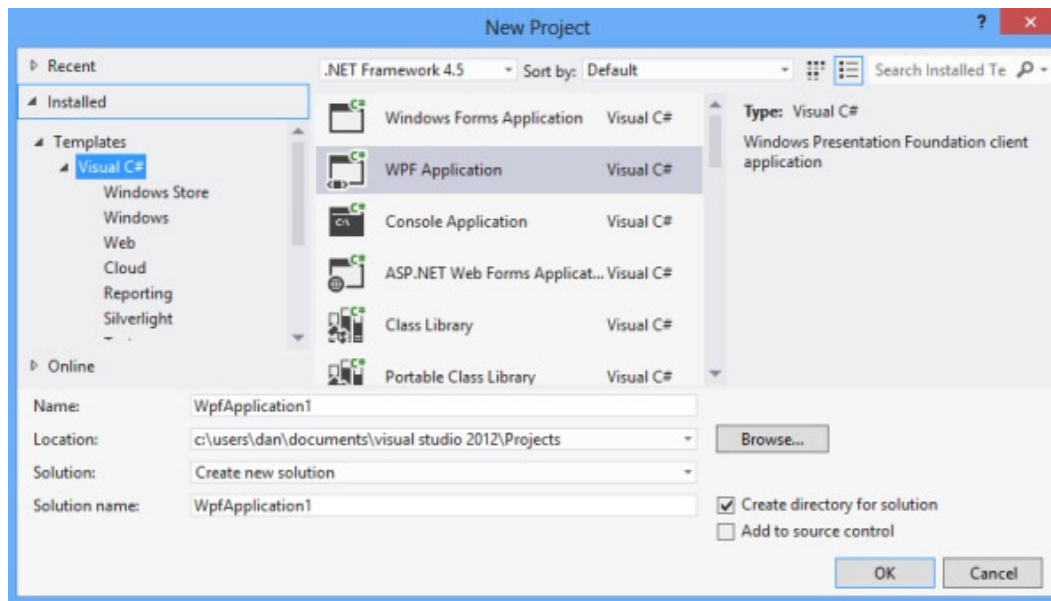


Figure 11-8. The New Project dialog box




Figure 11-9. Displaying critical information using a dialog box

Presenting a MessageBox to the User

The dialog box shown in Figure 11-9 is a special predefined type called a `MessageBox`. The `MessageBox` class is part of the `System.Windows` namespace. The `MessageBox` class can display a standard Windows message dialog box. To display a `MessageBox` to the user, you call the static `Show` method of the `MessageBox`, like so:

```
MessageBox.Show("File Saved");
```

The `Show` method is overloaded so that you can optionally show a `MessageBox` icon, show a title, change the buttons displayed, and set the default button. The only required setting is the text message to be displayed on the form. Figure 11-10 shows the `MessageBox` displayed by the previous code.



Figure 11-10. A basic `MessageBox`

The following code calls the `Show` method using some of the other parameters. Figure 11-11 shows the resulting `MessageBox` that gets displayed. For more information on the various parameters and settings available, look up the `MessageBox` class in the Visual Studio help file.

```
MessageBox.Show("Are you sure you want to quit?",
    "Closing Application", MessageBoxButton.OKCancel,
    MessageBoxImage.Question);
```

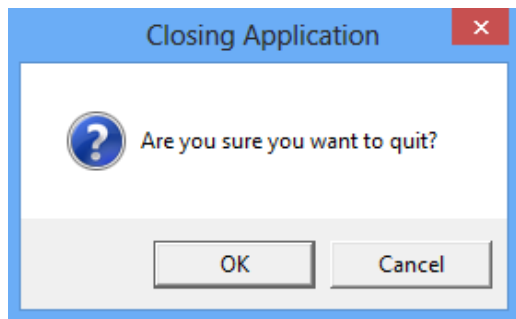


Figure 11-11. A more complex *MessageBox*

You will often use a `MessageBox` to query for a user response to a question. The user responds by clicking a button. The result is passed back as the return value of the `MessageBox.Show` method in the form of a `MessageBoxResult` enumeration. The following code captures the dialog box result entered by a user and closes the window (or not), depending on the result:

```
MessageBoxResult result = MessageBox.Show("Are you sure you want to quit?",
    "Closing Application", MessageBoxButton.OKCancel,
    MessageBoxImage.Question);
if (result == MessageBoxResult.OK)
{
    this.Close();
}
```

Creating a Custom Dialog Box

One of the most exciting features about the .NET Framework is its extensibility. Although there are many types of dialog boxes, you can use “right-out-of-the-box” ones for such tasks as printing, saving files, and loading files. You can also build your own custom dialog boxes. The first step in creating a custom dialog box is to add a new window to the application. Next, add any controls needed to interact with the user. Figure 11-12 shows a dialog box you might use to verify a user’s identity.

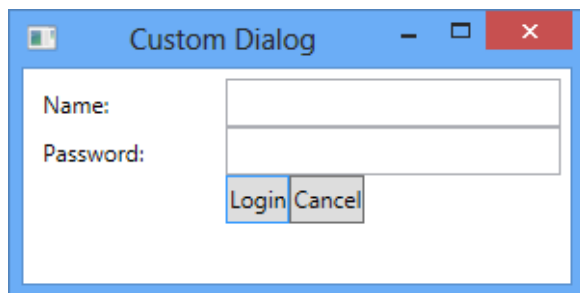


Figure 11-12. A custom dialog box

Setting the `IsCancel` property of the Cancel button to true associates it with the keyboard shortcut of the ESC key. Setting the `IsDefault` property of the Login button to true associates it with the keyboard Enter key. This is shown in the following XAML code:

```
<StackPanel Grid.Column="1" Grid.Row="3" Orientation="Horizontal">
  <Button Name="loginButton" IsDefault="True">Login</Button>
  <Button Name="cancelButton" IsCancel="True">Cancel</Button>
</StackPanel>
```

When the Login button is clicked, the click event of the button is responsible for validating the user input and setting the `DialogResult` property to either true or false. This value is returned to the window that called the `Show` method of the `DialogWindow` for further processing. The following code shows the `LoginDialog` window called and the `DialogResult` property being interrogated. Notice that the calling window has access to the objects defined on the `DialogWindow`. In this case, it is interrogating the `txtName` textbox's `Text` property.

```
LoginDialog dlg = new LoginDialog();
dlg.Owner = this;
dlg.ShowDialog();
if (dlg.DialogResult == false)
{
    string user = dlg.txtName.Text;
    MessageBox.Show("Invalid login for " + user, "Warning",
        MessageBoxButton.OK, MessageBoxImage.Exclamation);
    this.Close();
}
```

Data Binding in Windows-Based GUIs

Once you have retrieved the data from the business logic tier, you must present it to the user. The user may need to read through the data, edit the data, add records, or delete records. Many of the controls you'll want to add to a window can display data. The choice of which control to use often depends on the type of data you want to display, the ways you want to manipulate it, and the design you have in mind for your interface. Among the controls .NET developers commonly use to present data are the `TextBox`, `DataGrid`, `Label`, `ListBox`, `CheckBox`, and `Calendar`. When different fields of a data source are presented to the user in separate controls (for example, a first name `TextBox` and last name `TextBox`), it is important that the controls remain synchronized to show the same record.

The .NET Framework encapsulates much of the complexity of synchronizing controls to a data source through a process called data binding. When you create a binding between a control and some data, you are binding a binding target to a binding source. A binding object handles the interaction between the binding source and the binding target. `OneWay` binding causes changes to the source property to automatically update the target property, but changes to the target property are not propagated back to the source property. This is useful for read-only scenarios. `TwoWay` binding causes changes to either the source property or the target property to automatically update the other. This is useful for full data updating scenarios.

Binding Controls Using a DataContext

To bind a control to data, you need a data source object. The `DataContext` of a container control allows child controls to inherit information from their parent controls about the data source that is used for binding. The following code sets the `DataContext` property of the top level `Window` control. It uses a `DataSet` and a `TableAdapter` to fill a `Table` object and set it to the `DataContext` of the `Window`.

```
private void Window_Loaded(object sender, RoutedEventArgs e)
{
    pubsDataSet dsPubs = new pubsDataSet();
    pubsDataSetTableAdapters.storesTableAdapter taStores =
        new pubsDataSetTableAdapters.storesTableAdapter();
    taStores.Fill(dsPubs.stores);
    this.DataContext = dsPubs.stores.DefaultView;
}
```

The following XAML code binds the DataGrid columns to the Store table columns using the Path attribute. Using Binding for the source means “look up the container hierarchy until a DataContext is found.” In this case, it’s the DataContext of the Window container.

```
<DataGrid AutoGenerateColumns="False" ItemsSource="{Binding}">
  <DataGrid.Columns>
    <DataGridTextColumn x:Name="stor_idColumn"
      Binding="{Binding Path=stor_id}" Header="Id" />
    <DataGridTextColumn x:Name="stor_nameColumn"
      Binding="{Binding Path=stor_name}" Header="Name" />
    <DataGridTextColumn x:Name="stateColumn"
      Binding="{Binding Path=state}" Header="State" />
    <DataGridTextColumn x:Name="zipColumn"
      Binding="{Binding Path=zip}" Header="Zip" />
  </DataGrid.Columns>
</DataGrid>
```

The resulting DataGrid loaded with store data is shown in Figure 11-13.



The screenshot shows a window titled "Stores" with a DataGrid containing the following data:

Id	Name	State	Zip
6380	Eric the Read Books	WA	98056
7066	Barnum's	CA	92789
7067	News & Brews	CA	96745
7131	Doc-U-Mat: Quality Laundry and Books	WA	98014
7896	Fricative Bookshop	CA	90019
8042	Bookbeat	OR	89076

Figure 11-13. *Displaying store data with a DataGrid*

In the following activity, you will bind a DataGrid control to a DataTable containing data from Pubs database. You will also use a DataAdapter to update data changes made in the DataGrid control back to the Pubs database.

ACTIVITY 11-2. WORKING WITH DATA BOUND CONTROLS

In this activity, you will become familiar with the following:

- binding a DataGrid to a DataTable
- updating data using the DataAdapter

Binding a DataGrid to a DataTable

To bind a DataGrid to a DataTable object, follow these steps:

Create a DataSet

1. Start Visual Studio. Select File ► New ► Project.
2. Choose WPF Application. Rename the project to Activity11_2 and click the OK button.
3. After the project loads, locate the Data Sources window located on the left side of the screen. Click on the Add New Data Source link.
4. In the Data Source Configuration wizard, choose a data source type of Database and click Next to continue.
5. In the Choose a Database Model window, select the Dataset and click Next.
6. In the Choose your Data Connection window, select or create a connection to the Pubs database and click Next.
7. On the next screen, save the connection to the application configuration file.
8. In the Choose Your Database Objects window, expand the Tables' node and select the authors table. Click the Finish button.
9. Note in the Solutions Explorer window a pubsDataSet.xsd file has been added to the file. This file represents a strongly typed dataset object based on the pubs database. Double-click the file node in Solution Explorer to launch the dataset visual editor.
10. The visual editor contains an authors table. Select the authorsTableAdapter, as shown in Figure 11-14. In the Properties window, notice that the select, insert, update, and delete commands have been generated for you (see Figure 11-15).

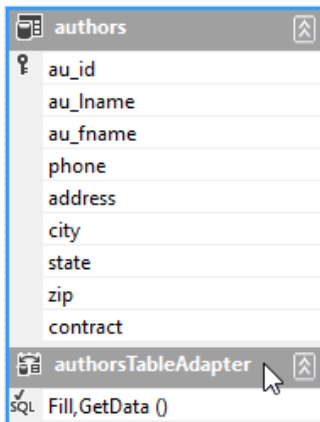


Figure 11-14. Selecting *authorsTableAdapter*

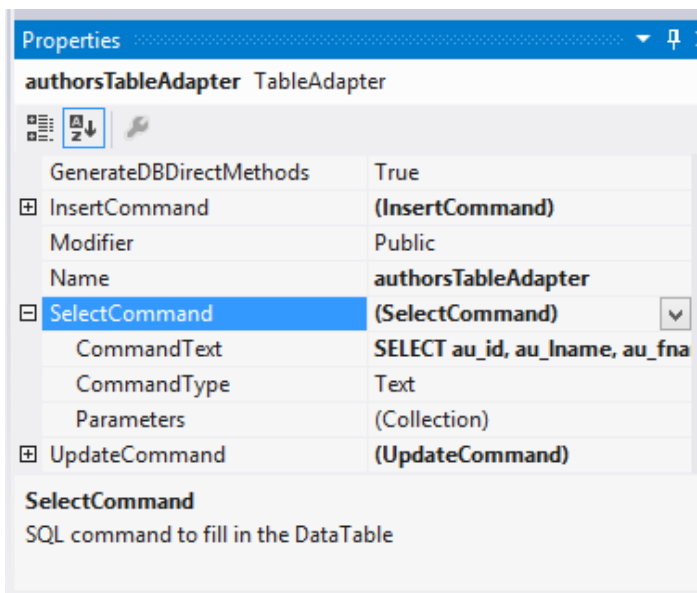


Figure 11-15. Viewing the generated command text

Create the Window Layout

1. Open the `MainWindow` in the XAML Editor window. Change the title of the Window to “Phone List”.
2. Inside the `Grid` tags, add a `DockPanel` control. Inside the `DockPanel`, add a `StackPanel`.

```

<Grid>
  <DockPanel>
    <StackPanel DockPanel.Dock="Top" Orientation="Horizontal">

      </StackPanel>
    </DockPanel>
  </Grid>

```

3. Inside the StackPanel, add two buttons—one for getting data and one for updating data. Add a Click event handler for each button.

```

<StackPanel DockPanel.Dock="Top" Orientation="Horizontal">
  <Button Name="btnGetData" Content="Get Data"
    Click="btnGetData_Click" />
  <Button Name="btnSaveData" Content="Save Data" />
</StackPanel>

```

4. Outside the StackPanel but inside the DockPanel, add a DataGrid.

```

<DataGrid Name="dgAuthors" AutoGenerateColumns="True"
  DockPanel.Dock="Bottom" />
</DockPanel>
</Grid>

```

Load the DataGrid

1. Open the MainWindow.xaml.cs file in the Code Editor window.
2. Add three class level variables of type pubsDataset, authorsTableAdapter, and authorsDataTable.

```

public partial class MainWindow : Window
{
  pubsDataSet _dsPubs;
  pubsDataSetTableAdapters.authorsTableAdapter _taAuthors;
  pubsDataSet.authorsDataTable _dtAuthors;

```

3. In the btnGetData_Click event, add code to fill the _taAuthors table and set it equal to the DataContext of the dgAuthors grid.

```

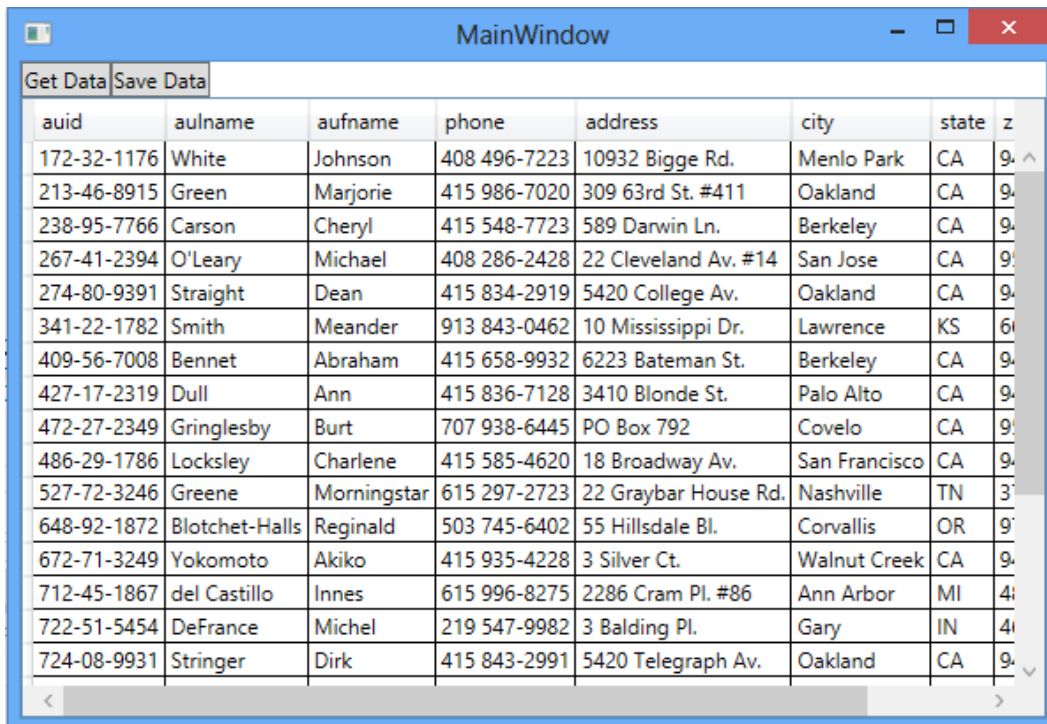
private void btnGetData_Click(object sender, RoutedEventArgs e)
{
  _dsPubs = new pubsDataSet();
  _taAuthors = new pubsDataSetTableAdapters.authorsTableAdapter();
  _dtAuthors = new pubsDataSet.authorsDataTable();
  _taAuthors.Fill(_dtAuthors);
  this.dgAuthors.DataContext = _dtAuthors;
}

```

- Switch back to the XAML Editor Window and add the ItemsSource binding to the DataGrid's XAML code. This will bind it to the DataContext.

```
<DataGrid Name="dgAuthors" AutoGenerateColumns="True"
          DockPanel.Dock="Bottom" ItemsSource="{Binding}" />
```

- Select Debug ► Start. Test the application by clicking the Get Data button. The DataGrid will load with the Authors' data (see Figure 11-16). Notice that since the AutoGenerateColumns property of the DataGrid is set to true, the grid loads with all the columns in the table. The headers of the grid columns are also the same name as the author's table columns.
- After viewing the window, stop the debugger.



aid	aulname	aufname	phone	address	city	state	z
172-32-1176	White	Johnson	408 496-7223	10932 Bigge Rd.	Menlo Park	CA	9
213-46-8915	Green	Marjorie	415 986-7020	309 63rd St. #411	Oakland	CA	9
238-95-7766	Carson	Cheryl	415 548-7723	589 Darwin Ln.	Berkeley	CA	9
267-41-2394	O'Leary	Michael	408 286-2428	22 Cleveland Av. #14	San Jose	CA	9
274-80-9391	Straight	Dean	415 834-2919	5420 College Av.	Oakland	CA	9
341-22-1782	Smith	Meander	913 843-0462	10 Mississippi Dr.	Lawrence	KS	6
409-56-7008	Bennet	Abraham	415 658-9932	6223 Bateman St.	Berkeley	CA	9
427-17-2319	Dull	Ann	415 836-7128	3410 Blonde St.	Palo Alto	CA	9
472-27-2349	Gringlesby	Burt	707 938-6445	PO Box 792	Covelo	CA	9
486-29-1786	Locksley	Charlene	415 585-4620	18 Broadway Av.	San Francisco	CA	9
527-72-3246	Greene	Morningstar	615 297-2723	22 Graybar House Rd.	Nashville	TN	3
648-92-1872	Blotch-Halls	Reginald	503 745-6402	55 Hillsdale Bl.	Corvallis	OR	9
672-71-3249	Yokomoto	Akiko	415 935-4228	3 Silver Ct.	Walnut Creek	CA	9
712-45-1867	del Castillo	Innes	615 996-8275	2286 Cram Pl. #86	Ann Arbor	MI	4
722-51-5454	DeFrance	Michel	219 547-9982	3 Balding Pl.	Gary	IN	4
724-08-9931	Stringer	Dirk	415 843-2991	5420 Telegraph Av.	Oakland	CA	9

Figure 11-16. The authors DataGrid

Updating Data

1. Open the MainWindow.xaml.cs file in the Code Editor window. Add the following code to update the data in the btnSaveData_Click event handler. This code uses the table adapter's update command to send the changes back to the database. You must also alter the btnSaveData's XAML to call this event when it is clicked.

```
private void btnSaveData_Click(object sender, RoutedEventArgs e)
{
    try
    {
        _taAuthors.Update(_dtAuthors);
        MessageBox.Show("Data Saved.",
            "Information", MessageBoxButton.OK,
            MessageBoxImage.Information);
    }
    catch (Exception ex)
    {
        MessageBox.Show("Could not save data!",
            "Warning", MessageBoxButton.OK,
            MessageBoxImage.Warning);
    }
}
```

2. Update the Grid's XAML code to only show the first name, last name, and phone columns.

```
<DataGrid Name="dgAuthors" AutoGenerateColumns="False"
    DockPanel.Dock="Bottom" ItemsSource="{Binding}">
    <DataGrid.Columns>
        <DataGridTextColumn Header="Last Name"
            Binding="{Binding Path='au_lname'}" />
        <DataGridTextColumn Header="First Name"
            Binding="{Binding Path='au_fname'}" />
        <DataGridTextColumn Header="Phone"
            Binding="{Binding Path='phone'}" />
    </DataGrid.Columns>
</DataGrid>
```

3. Select Debug ► Start. Test the application by clicking the Get Data button. Update some of the Names. Click the Save Data button and then click the Get Data button to verify the names were saved to the database.
 4. After testing, stop the debugger and exit Visual Studio.
-

Creating and Using Control and Data Templates

In WPF, every control has a template that manages its visual appearance. If you don't explicitly set its `Style` property, then it uses a default template. Creating a custom template and assigning it to the `Style` property is an excellent way to alter the look and feel of your applications. Figure 11-17 shows a rounded button created by using a control template.

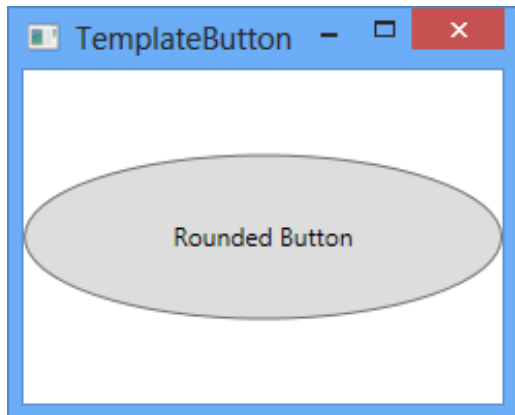


Figure 11-17. Creating a rounded button with a custom template

The following XAML is the markup that defines the custom template used to create the rounded button in Figure 11-17.

```
<Window.Resources>
  <Style x:Key="RoundedButtonStyle" TargetType="Button">
    <Setter Property="Template">
      <Setter.Value>
        <ControlTemplate TargetType="{x:Type Button}">
          <Grid>
            <Ellipse Fill="{TemplateBinding Background}"
              Stroke="{TemplateBinding BorderBrush}"/>
            <ContentPresenter HorizontalAlignment="Center"
              VerticalAlignment="Center"/>
          </Grid>
        </ControlTemplate>
      </Setter.Value>
    </Setter>
  </Style>
</Window.Resources>
```

The following XAML code is used to bind the custom style to a button using the button's `Style` property:

```
<Button Content="Rounded Button" Style="{StaticResource RoundedButtonStyle}"
```


Along with control style templates, you can also create data templates. Data templates let you customize how your business objects will look when you bind them in your UI. A good example of when you need to use a custom data template is the list box. By default, it renders data as a single line of text. When you try to bind it to a list of employee objects, it calls the `Tostring()` method and writes it out to the display. As you can see in Figure 11-18, this is clearly not what you want.

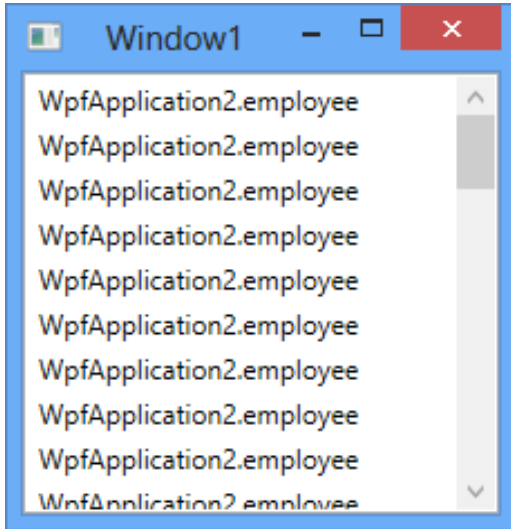


Figure 11-18. *ListBox using the default DataTemplate*

By adding a `DataTemplate` to the `ListBox` control, you can not only get the employee data to display, but you can also control how it gets displayed. The following XAML adds a `DataTemplate` to the `ListBox`, and Figure 11-19 shows the result:

```
<ListBox ItemsSource="{Binding}" >
  <ListBox.ItemTemplate>
    <DataTemplate>
      <StackPanel Orientation="Horizontal">
        <TextBlock FontWeight="Bold" Text="{Binding Path='lname'}" />
        <TextBlock Text=", " />
        <TextBlock Text="{Binding Path='fname'}" />
        <TextBlock Text=" " />
        <TextBlock Text="{Binding Path = 'minit'}" />
      </StackPanel>
    </DataTemplate>
  </ListBox.ItemTemplate>
</ListBox>
```

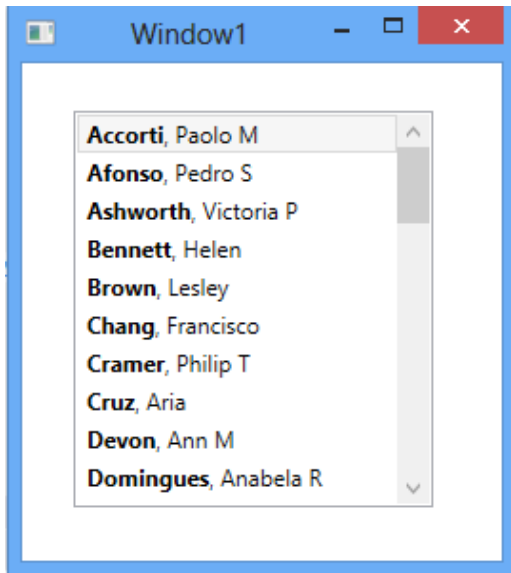


Figure 11-19. *ListBox using a custom DataTemplate*

In the following activity, you will bind a `ListBox` control to an entity created from the Pubs database using an entity data model. You will also create a master detail view by synchronizing a `ListBox` control and a `DataGrid` control.

ACTIVITY 11-3. WORKING WITH DATA TEMPLATES

In this activity, you will become familiar with the following:

- binding a `ListBox` to an Entity
- creating a `DataTemplate`
- creating a Master Detail View

Binding a `ListBox` to an Entity

To bind a `ListBox` to an entity object, follow these steps:

1. Start Visual Studio. Select **File** ► **New** ► **Project**.
2. Choose **WPF Application** under the **Windows** templates. Rename the project to `Activity11_3` and click the **OK** button.
3. After the project loads locate the **Data Sources** window (It should be to the left of the designer window). Click on the **Add New Data Source** link.
4. In the **Data Source Configuration** wizard, choose a data source type of **Database**.
5. In the **Choose a Database Model** window, select the **Entity Data Model**.

6. In the Choose Model Contents window, select the Generate from database option.
7. In the Choose your Data Connection window, select or create a connection to the Pubs database. Save the connection to the application configuration file.
8. In the Choose Your Database Objects window, expand the Tables node and select the stores and sales tables. Click the Finish button.
9. Notice in the Solutions Explorer window a Model1.edmx file has been added to the project. This file contains the relational mapping between the entities and the tables in the pubs database.

Creating the Data Template

1. Add a DockPanel and a ListBox control in the XAML Editor window between the default Grid tags.

```
<Grid>
  <DockPanel>
    <ListBox Name="StoresList" DockPanel.Dock="Left" ItemsSource="{Binding}">
      </ListBox>
    </DockPanel>
  </Grid>
```

2. Add a window Loaded event attribute to the Window's opening tag.

```
<Window x:Class="Activity11_3.MainWindow"
  xmlns="http://schemas.microsoft.com/winfx/2006/xaml/presentation"
  xmlns:x="http://schemas.microsoft.com/winfx/2006/xaml"
  Title="MainWindow" Height="350" Width="525"
  Loaded="Window_Loaded_1">
```

3. Add a Window_Loaded event handler in the code file that sets the DataContext of the ListBox to the stores entities.

```
private void Window_Loaded_1(object sender, RoutedEventArgs e)
{
  pubsEntities pEntities = new pubsEntities();
  this.StoresGrid.DataContext = pEntities.stores.ToList<store>();
}
```

4. Add a DataTemplate to display the store name in a TextBlock control inside the ListBox control.

```
<ListBox Name="StoresList" DockPanel.Dock="Left"
  ItemsSource="{Binding}">
  <ListBox.ItemTemplate>
    <DataTemplate>
      <TextBlock FontWeight="Bold" Text="{Binding Path='stor_name'}" />
    </DataTemplate>
  </ListBox.ItemTemplate>
</ListBox>
```

5. Select Debug ► Start. Make sure the ListBox shows the store names. When you're done viewing the ListBox, stop the debugger.
6. To implement a master/detail data view, add a DataGrid control to the DockPanel control after the ListBox control. The Binding of the grid is set to the same as the list box, which is the store entity, but the binding path is set to the sales entity. This will cause the data grid to show the sales items of the store selected in the list box.

```
<DataGrid Name="SalesGrid" DockPanel.Dock="Right"
          ItemsSource="{Binding Path='sales'}" AutoGenerateColumns="False">
  <DataGrid.Columns>
    <DataGridTextColumn Header="Order Number"
      Binding="{Binding Path='ord_num'}"/>
    <DataGridTextColumn Header="Order Date"
      Binding="{Binding Path='ord_date'}"/>
  </DataGrid.Columns>
</DataGrid>
```

7. Add the following property to the ListBox control in the XAML code. This will ensure that the ListBox control and DataGrid control will remain in sync.

```
IsSynchronizedWithCurrentItem="True"
```

8. Launch the application in the debugger. Your window should look similar to Figure 11-20. Click on different stores in the list box. You should see the data grid update with the store's sales data. After testing, stop the debugger and close Visual Studio.

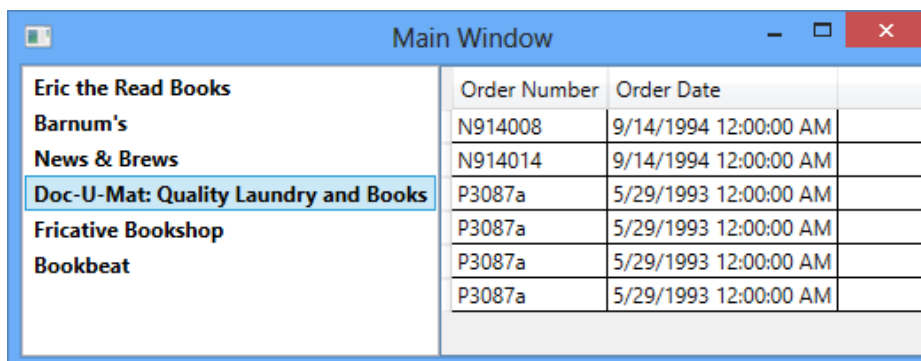


Figure 11-20. Viewing master/detail data

Summary

In this chapter, you looked at implementing the interface tier of an application. You implemented the user interface through a WPF-based application front end. Along the way, you took a closer look at the classes and namespaces of the .NET Framework used to implement rich Windows-based user interfaces. You saw how to use XAML syntax to define the controls and layout of the interface. You also saw how easy it is to bind the controls to the data and present it to the users. In Chapter 13, you will reexamine these concepts and apply them to developing the new Windows 8 and Windows Store applications.

In the next chapter, you will revisit the UI tier of a .NET application, but instead of implementing the GUI using WPF, you will implement the GUI as a web-based application. Along the way, you will take a closer look at the namespaces available for creating web-based GUI applications and the techniques involved in implementing the classes contained in these namespaces.



Developing Web Applications

In Chapter 11, you learned how to build a simple Windows-based graphical user interface (GUI) using C# and WPF. Although WPF gives programmers the ability to build extremely rich user interfaces easily, it is not always practical to assume users will access your programs through a traditional Windows-based PC. With the proliferation of intranets, web applications, and mobile devices, applications now need to allow users the ability to access the interface through a variety of browsers and devices. This chapter shows you how to build a web-based user interface using ASP.NET. ASP.NET is a web framework and supports two approaches to building web applications: Web Forms and ASP.NET MVC applications (which are based on the Model-View-Controller pattern). For the purposes of just starting out with web programming I suggest using Web Forms. Web Forms provide the familiar control and event-based programming model that you used in Chapter 11 when building windows applications. (For a discussion of the differences between developing in Web Forms and MVC see the official ASP.NET web site at <http://www.asp.net/>)

In this chapter, you will be performing the following tasks:

- creating Web Pages
- working with controls
- responding to page and control events
- maintaining state
- implementing data-bound controls

Web Pages and Web Forms

ASP.NET Web Pages provide the user interface for your web applications. While you can build a web site with traditional HTML pages, ASP.NET pages offer many advantages. They support the implementation of both client-side and server-side application logic. The programming model is based on a familiar object oriented event-driven paradigm. ASP.NET is built on top of the .NET Framework and receives all the benefits it supplies, such as type safety, managed code, and a robust class library. ASP.NET Web Pages are designed to work in many of the popular desktop browsers like Internet Explorer and Chrome. When an ASP.NET Web Page is rendered, the type of browser requesting the page is detected. The page is then dynamically rendered using the correct HTML to take advantage of browser-specific features such as layout, formatting, and styles.

A Web Page consists of two parts: the visual interface and the programming logic. The visual interface consists of a text file containing HTML markup tags, references to include files, processing directives, and client side scripting blocks. This text file is given the default extension of .aspx and is referred to as a page. Contained within the page is a Web Form. The Web Form acts as a container for the text and controls that will be displayed in the browser. Figure 12-1 shows the .aspx file for an ASP.NET page containing a Web Form, which contains a TextBox control, a Button control, and a Label control. When the Web Page is compiled, the code is combined into a new class file. This file is then compiled into a new class that inherits from System.Web.UI.Page class. It is this class code that executes when the .aspx page is requested. Figure 12-2 shows the page rendered in a browser (Internet Explorer).

```

<%@ Page Language="C#" AutoEventWireup="true" CodeBehind="Page1.aspx.cs"
    Inherits="Chapter12Demos.Page1" %>

<!DOCTYPE html>

<html xmlns="http://www.w3.org/1999/xhtml">
<head runat="server">
    <title></title>
</head>
<body>
    <form id="form1" runat="server">
        <div>
            <asp:Label ID="Label1" runat="server" Text="Enter Your Name:"></asp:Label>
            <asp:TextBox ID="TextBox1" runat="server"></asp:TextBox>
            <br />
            <asp:Button ID="Button1" runat="server" Text="OK" />
        </div>
    </form>
</body>
</html>

```

Figure 12-1. Code contained in the .aspx file




Figure 12-2. Web page rendered in IE

Looking at the page directive at the top of the page reveals the second piece of the programming model, which is referred to as the code-behind file and is indicated by the .aspx.cs extension. The code-behind file contains a partial class file. Code is placed in the partial class file to handle any necessary events exposed by the page or controls contained within the page. The following class code is contained in the code-behind file linked to the .aspx file of Figure 12-1:

```

public partial class Page1 : System.Web.UI.Page
{
    protected void Page_Load(object sender, EventArgs e)
    {
    }
}

```

When a browser requests the Web page, the page file and the code-behind file combine, along with a partial class generated for the .aspx file, and compile into a single executable Page class file. It is this Page class that intercepts and processes incoming requests. After processing the incoming request, the Page class dynamically creates and sends an HTML response stream back to the browser. Because the Page class is compiled code, execution is much faster than technologies that rely on interpreting script. The Internet Information Services (IIS) Web server is also able to cache the execution code in memory to further increase performance.

Web Server Control Fundamentals

The .NET Framework provides a set of Web server controls specifically for hosting within a Web Form. The types of Web server controls available include common form controls such as a TextBox, Label, and Button, as well as more complex controls such as a GridView and a Calendar. The Web server controls abstract out the HTML coding from the developer. When the Page class sends the response stream to the browser, the Web server control is rendered on the page using the appropriate HTML. The HTML sent to the browser depends on such factors as the browser type and the control settings that have been made.

The following code is used to place a TextBox Web server control in the Web Form:

```
<asp:TextBox ID="txtName" runat="server" BorderStyle="Dashed"
  ForeColor="#0000C0"></asp:TextBox>
```

The control is then rendered in Internet Explorer (IE) 10.0 as the following HTML code:

```
<input name="txtName" type="text" id="txtName"
  style="color:#0000C0;border-style:Dashed;" />
```

If the TextMode property of the TextBox control is set to MultiLine, the code in the Web Page is altered to reflect the property setting:

```
<asp:TextBox ID="txtName" runat="server" BorderStyle="Dashed" ForeColor="#0000C0"
  TextMode="MultiLine"></asp:TextBox>
```

Although the change in the code for the Web server control was minimal, the HTML code rendered to the browser changes to a completely different HTML control:

```
<textarea name="txtName" rows="2" cols="20" id="txtName"
  style="color:#0000C0;border-style:Dashed;"></textarea>
```

Web server controls offer .NET programmers many advantages, including a familiar event-driven object model, automatic browser detection with dynamic rendering, data binding capabilities, and automatic control state maintenance, just to name a few.

Understanding Web Page and Web Server Control Inheritance Hierarchy

At the top of the Web Page interface code contained in the .aspx file is the following page directive:

```
<%@ Page Language="C#" AutoEventWireup="true" CodeBehind="Page2.aspx.cs"
  Inherits="Chapter12Demo1.Page2" %>
```


This code reveals that the Web Page interface code inherits from the Page2 code-behind class located in the Page2.aspx.cs file. The code-behind class in turn inherits from the Page class located in the System.Web.UI namespace:

```
public partial class Page2 : System.Web.UI.Page
```

The Page class exposes important functionality needed to program a Web application and interact with the Web server. For example, it enables access to the Application, Session, Response, Request, and Server objects. The Application object enables sharing of values among all users of the Web application. The Request object enables the reading of values sent by the Web Page during a Web request. The Page class also exposes functionality such as working with Postback events initiated by script embedded in the Web Page, initiating page-level validation, and registration of hidden fields required by server controls.

If you trace the inheritance chain of the Web Page further, you discover that the Page class inherits functionality from the TemplateControl class. This class adds support for loading user controls, which are custom controls commonly created to partition and reuse user interface (UI) functionality in several Web Pages. User controls are created in .ascx files and are hosted by the Web Pages. Along with managing the user controls inserted into a Web Page, the TemplateControl class adds transactional support and error handling functionality to the Page class.

The TemplateControl class inherits from the Control class. The Control class exposes much of the functionality needed by all server controls. This class includes important properties such as the ID and Page properties. The ID property gets and sets the programmatic identifier of a control, and the Page property gets a reference to the Page object that contains the control. The Control class exposes methods for rendering the HTML code for the control and any child controls. There are methods for handling and raising events such as the Load, Unload, and PreRender events. The Control class also exposes the functionality needed to add data binding support to Web server controls. Figure 12-3 shows the hierarchy chain of the Page class in the Class View window.

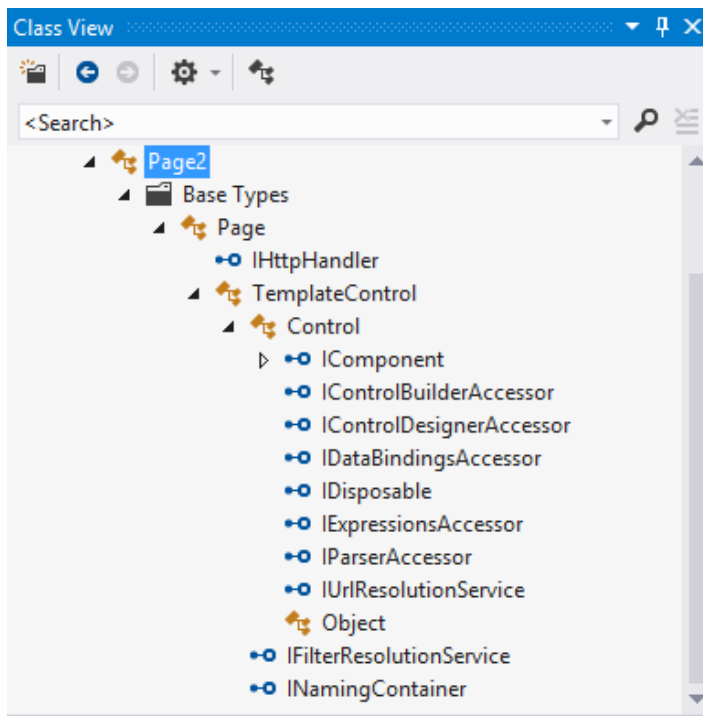


Figure 12-3. Hierarchy chain of the Page class

Tracing the hierarchy chain of a Web server control—for example, a TextBox control—reveals that they also gain much of their functionality from the Control class. When you place a TextBox server control on a Web Form, at runtime it instantiates an object instance of the TextBox class, which is part of the `System.Web.UI.WebControls` namespace. This class exposes properties, methods, and events needed by a TextBox. For example, it defines such properties as the `Text`, `ReadOnly`, and `Wrap` properties. The TextBox class also adds functionality for raising and handling the `TextChanged` event.

All Web server control classes inherit common functionality from the `WebControl` class. The `WebControl` class provides properties that control the look of the control when it gets rendered in the browser. For example, the `ForeColor`, `BackColor`, `Height`, and `Width` properties are defined in the `WebControl` class. It also defines properties that control the behavior of the control such as the `Enabled` and `TabIndex` properties. The `WebControl` class inherits from the `Control` class discussed previously. Figure 12-4 shows the hierarchy chain of a TextBox Web server control in a class diagram.

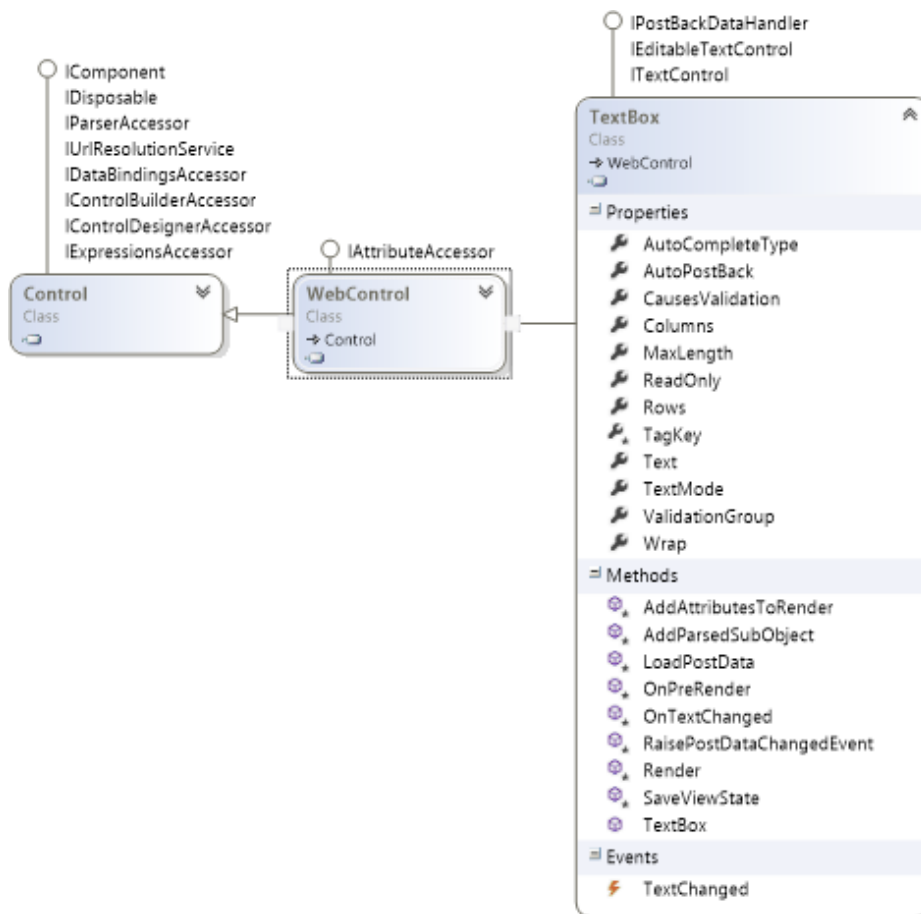


Figure 12-4. Hierarchy chain of a TextBox control

Using the Visual Studio Web Page Designer

The Visual Studio Integrated Development Environment (IDE) includes an excellent Web Page Designer. Using the designer, you can drag and drop controls onto the Web Page from the Toolbox and set control properties using the Properties window. Figure 12-5 shows some of the controls available in the Visual Studio designer toolbox.

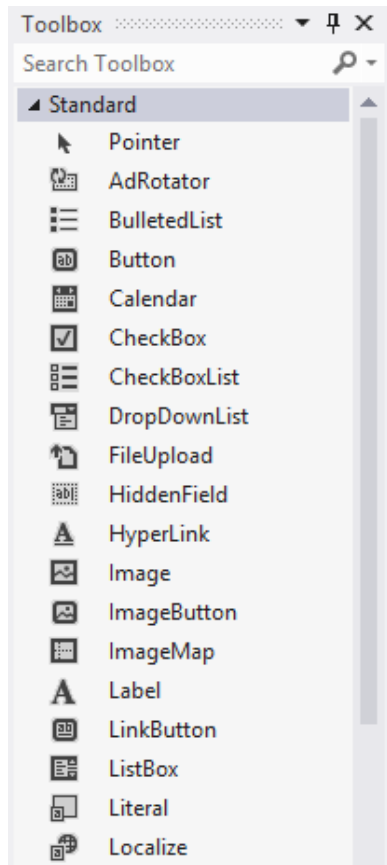


Figure 12-5. Visual Studio Web Page Designer Toolbox

The Web Page Designer contains three tabs at the bottom of the window—Design, Split, and Source. These tabs switch the view in the window from a visual representation of the Web Page (Design) to the ASP.NET and HTML markup (Source). The split view shows both the designer and the source code as a split window. The Source view shows the tag markup code used to render the page and allows you to insert client-side script into the page. Figure 12-6 shows a Web Page displayed in the Split view of the designer.

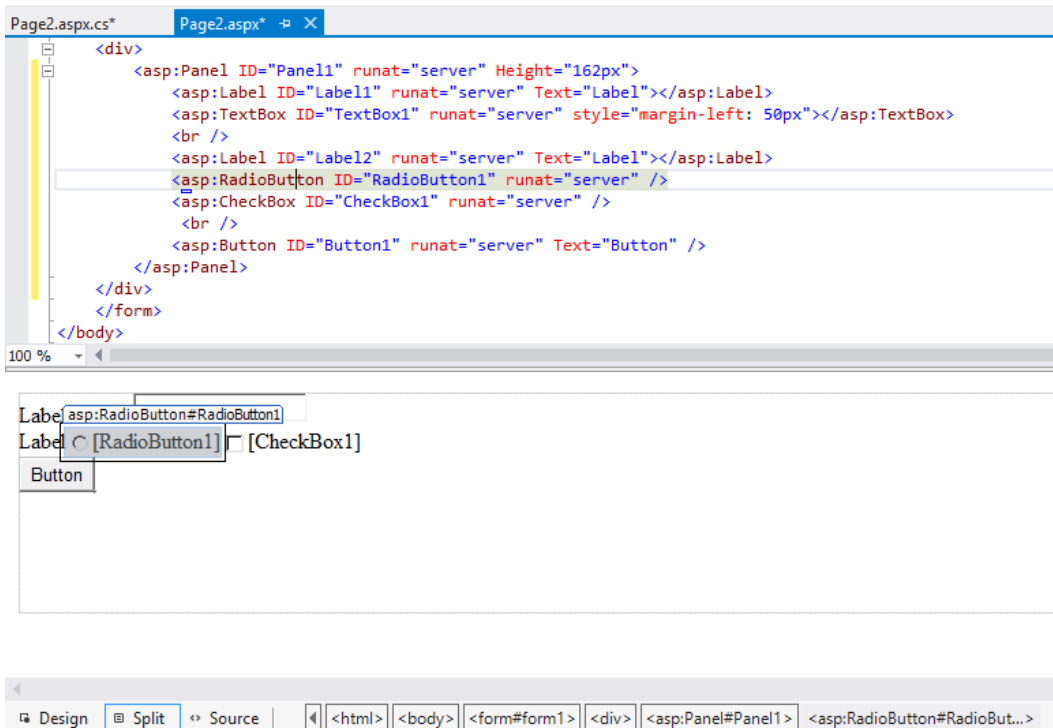


Figure 12-6. Visual Studio Web Page Designer in the Split view

The Web Page Life Cycle

When a Web page is requested, it runs through a series of stages, each with a defined purpose. For example, during the Initialization stage, any themes are applied to the page and any controls on the page have their UniqueID property set. In the Load stage, if the page is being posted back (has already rendered once and is sending information back), control properties are loaded from information stored in hidden controls called the view state and control state (view state is described in more detail later in the chapter). The Postback event handling stage is where control event handlers are called. In order to properly debug web pages, it is important that you understand the various stages of the page life cycle and when they occur. Table 12-1 summarizes the various stages that occur when a page is requested.

Table 12-1. Stages of a Page Request

Stage	Description
Page request	ASP.NET determines whether the page needs to be parsed and compiled or whether a cached version of the page can be sent.
Start	Page properties such as Request and Response are set. Page sets the IsPostBack property.
Initialization	Controls on the page are available and each control's UniqueID property is set. A master page and themes are also applied to the page, if applicable.

(continued)

Table 12-1. (continued)

Stage	Description
Load	If the request is a postback, control properties are loaded with information recovered from view state and control state.
Postback event handling	If the request is a postback, control event handlers are called. After that, the <code>Validate</code> method of all validator controls is called, which sets the <code>IsValid</code> property of individual validator controls and of the page.
Rendering	Before rendering, view state is saved for the page and all controls. During the rendering stage, the page calls the <code>Render</code> method for each control.
Unload	Page objects such as <code>Response</code> and <code>Request</code> are unloaded and cleanup is performed.

As the page runs through its lifecycle, it raises various events that you can respond to by creating an event handler that runs your code. The web pages support auto event wiring so it looks for methods with standard naming to run when an event occurs. These methods are in the form of `Page_event`, so to wire up an event handler for the page load event you would name the method `Page_Load`. The `Page_Load` event occurs at the beginning of a page request after the controls in the Page have been initialized. This event is often used to read data from and load data into the server controls. The `IsPostBack` Page property is interrogated to determine if the page is being viewed for the first time or in response to a Postback event. The following code uses the page load event to write a message on the page if it is a post back.

```
protected void Page_Load(object sender, EventArgs e)
{
    if (IsPostBack)
    {
        Response.Write("This is a post back.");
    }
}
```

Some other common page events you may create event handlers for are the `Page_Init`, `Page_PreRender`, and the `Page_Unload` events.

Control Events

Just like their Windows Graphical User Interface (GUI) counterparts, Web controls interact with users based on an event-driven model. When an event occurs on the client—for example, a button click—the event information is captured on the client and the information is transmitted to the server via a Hypertext Transfer Protocol (HTTP) post. On the server, the event information is intercepted by an event delegate object, which in turn informs any event handler methods that have subscribed to the invocation list. Although this sounds complicated, the .NET Framework classes abstract and encapsulate most of the process from you.

Because of the overhead of handling events through postbacks requiring a roundtrip from the browser to the server, only a limited set of events are exposed by Web Forms and server controls. Events that can occur frequently such as mouse movement and keypress events are not supported through server-side event handlers. These events are supported through client-side event handlers written in script. One common event you may respond to using a server-side event handler is the button click event.

```
protected void Button1_Click(object sender, EventArgs e)
{
    Response.Write ("Hello " + this.TextBox1.Text);
}
```

In order to wire up the event to the button control, the `onClick` attribute is added to the button's markup code. Controls that are wired up to respond to events must also have an ID unique to the form container set.

```
<asp:Button ID="Button1" runat="server" Text="OK" OnClick="Button1_Click"/>
```

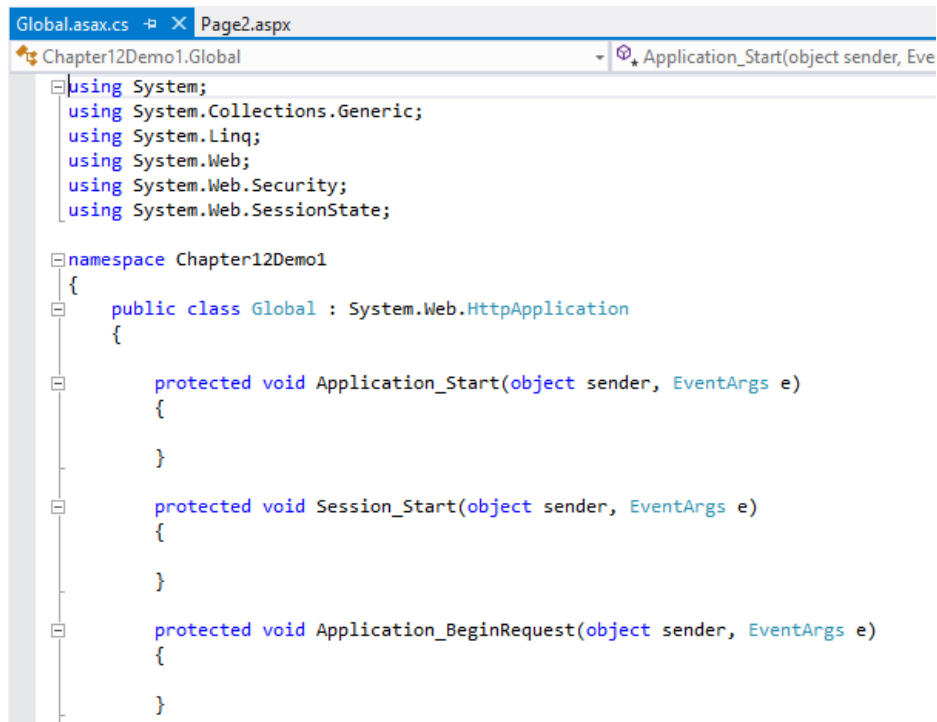
By convention, the name of the event handler method for the control is the name of the object issuing the event followed by an underscore character and the name of the event (which is similar to Page level event handling). The actual name of the event handler, however, is unimportant.

All event handlers must provide two parameters that will be passed to the method when the event is fired. The first parameter is the sender, which represents the object that initiated the event. The second parameter, `e` of type `System.EventArgs`, is an object used to pass any information specific to a particular event. For example, the `EditCommand` event is raised when the Edit button for an item in a `DataList` control is clicked. The event argument `e` is used to pass information about the item being edited in the `DataList` control. The following code checks to see if the item selected in the `DataList` is enabled:

```
protected void DataList1_EditCommand(object source, DataListCommandEventArgs e)
{
    if (e.Item.Enabled)
    {
        //Add code here.
    }
}
```

Understanding Application and Session Events

Along with the Page and control events, the .NET Framework provides the ability to intercept and respond to events raised when a Web session starts or ends. A Web session starts when a user requests a page from the application and ends when the session is abandoned or it times out. In addition to the session events, the .NET Framework exposes several application-level events. The `Application_Start` event occurs the first time anyone requests a page in the Web application. `Application_BeginRequest` occurs when any page or service is requested. There are corresponding events that fire when ending requests, authenticating users, raising errors, and stopping the application. The session-level and application-level event handlers are in the `Global.asax.cs` code-behind file. Figure 12-7 shows the `Global.asax.cs` file in the code editor.



```
Global.asax.cs Page2.aspx
Chapter12Demo1.Global Application_Start(object sender, Eve
using System;
using System.Collections.Generic;
using System.Linq;
using System.Web;
using System.Web.Security;
using System.Web.SessionState;

namespace Chapter12Demo1
{
    public class Global : System.Web.HttpApplication
    {
        protected void Application_Start(object sender, EventArgs e)
        {
        }

        protected void Session_Start(object sender, EventArgs e)
        {
        }

        protected void Application_BeginRequest(object sender, EventArgs e)
        {
        }
    }
}
```

Figure 12-7. The *Global.asax.cs* file open in the code editor

When the application begins, the *Global.asax* compiles into a dynamically generated class that derives from the *HttpApplication* base class. This class exposes and defines the methods, properties, and events common to all application objects within an ASP.NET application. A detailed explanation of the *HttpApplication* class is beyond the scope of this book. For more information, see the *HttpApplication* object in the help files.

■ **Note** The *Global.asax* file is optional and must be added to the web application.

ACTIVITY 12-1. WORKING WITH WEB PAGES AND SERVER CONTROLS

In this activity you will become familiar with the following:

- creating a Web Form–based GUI application
- working with page and server control events

Creating the Web Application

To create the Web application, follow these steps:

1. Start Visual Studio. On the File menu, choose New Web Site.
2. Choose an ASP.NET Empty Web Site. Change the name of the web site at the end of location path Activity12_1. By default, web sites are stored in a folder under the path Documents\Visual Studio 2012\WebSites. Figure 12-8 shows the New Web Site dialog box for creating a Web Site or Service.




Figure 12-8. The New Web Site dialog box

3. Click the OK button in the New Web Site dialog box.
4. After the Web Site is created, you need to add a web form. Under the Website menu, select “Add new item” and add a Web Form named Page1.aspx.
5. You are presented with the Web Page Designer in Source view. Switch to Design view by clicking the tab in the lower left corner of the window.

- Select the `<div>` tag at the bottom of the Design window (see Figure 12-9).




Figure 12-9. Selecting the div tag

- In the Properties window, select the style property and click on the ellipses to launch the Modify Style dialog window. Under the Position category, set the height to 200 pixels (see Figure 12-10) and click the OK button

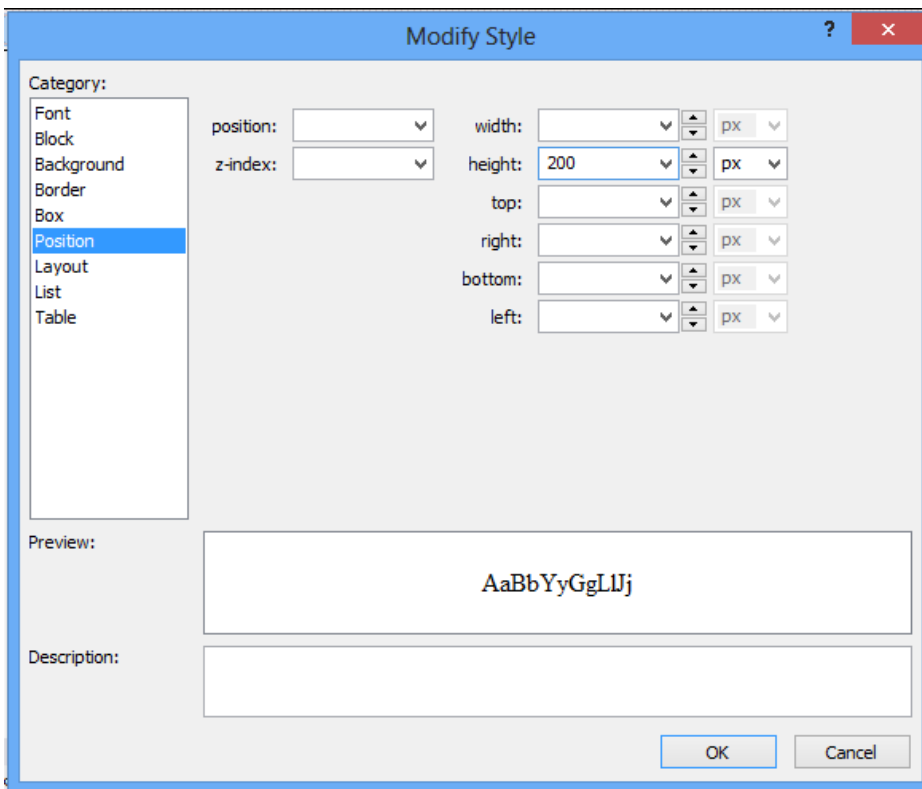


Figure 12-10. Setting the div height

- Switch the Page Design window to split mode. Drag a Label control from the toolbox and drop it between the div tags in the source pane. Change the text to "Hello Stranger..." and the ID to lblMessage.

9. You may get a message asking you to synchronize the design view and the source view. After synchronizing the views, your designer should look similar to Figure 12-11.

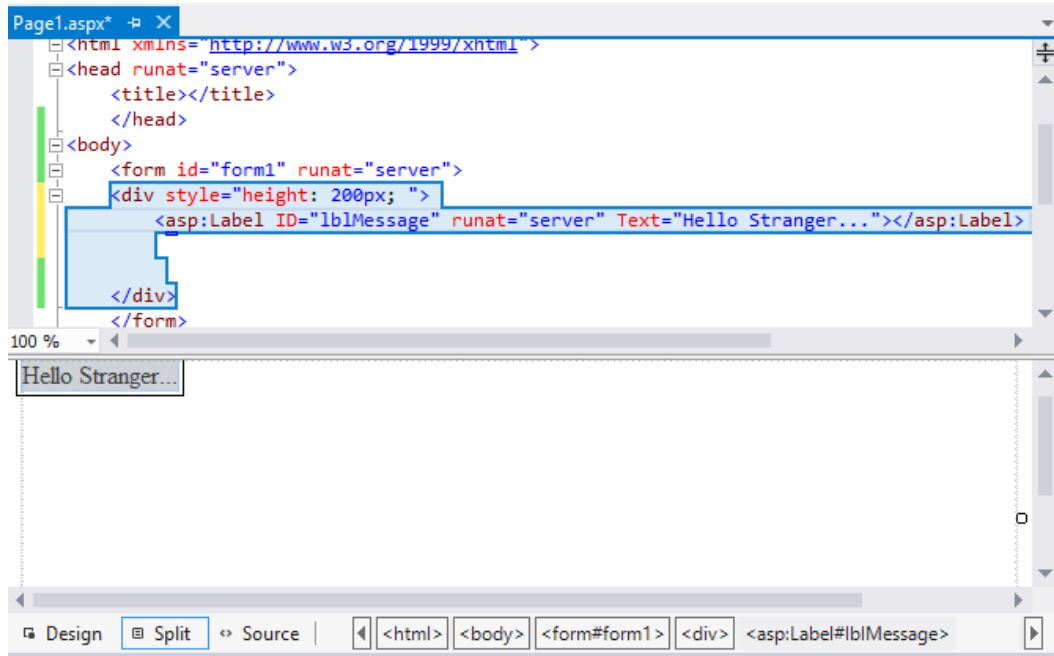


Figure 12-11. Adding a Label control

10. In a similar process to step 9, add the controls in Table 12-2 between the div tags. Don't worry about the positioning on the form yet.

Table 12-2. Web Control Properties

Control	Property	Value
Label	ID	lblName
	Text	Enter your name:
TextBox	ID	txtName
	Text	[Blank]
Button	ID	btnSubmit
	Text	Submit

11. To position a control, select the control in the designer and in the format menu, select Set Position and check Absolute (see Figure 12-12).

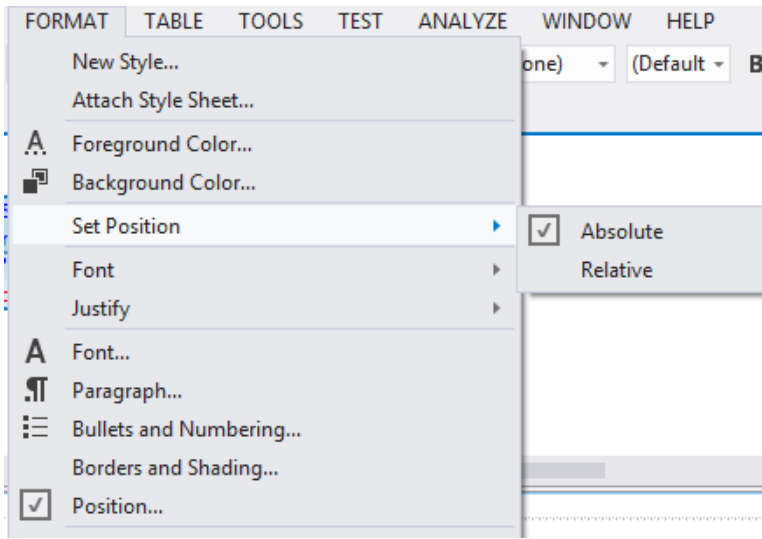


Figure 12-12. Turning on absolute positioning

12. Turn on absolute positioning for the rest of the controls. You should now be able to drag them around in the visual designer to position them. Arrange the controls on the form so they look similar to Figure 12-13.

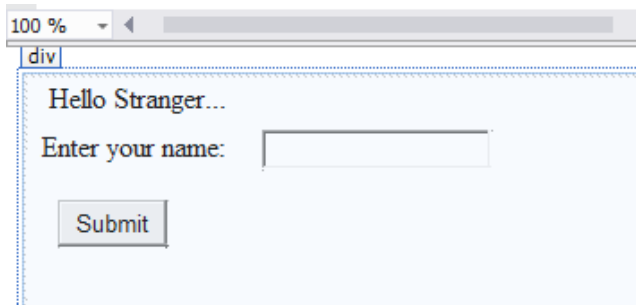


Figure 12-13. Positioning the controls

Creating Server-Side Control Event Handlers

To create server-side control event handlers, follow these steps:

1. In the Design pane, double click the Submit button. This action will open up the Page1.aspx.cs code behind page and create an event handler method for the button's click event.
2. In the button's click event handler, check to see if the page request is the result of a postback. If it is, change the lblMessage text to say hello using the text contained in the txtName control:

```
protected void btnSubmit_Click(object sender, EventArgs e)
{
    if (IsPostBack)
    {
        lblMessage.Text = "Hello " + txtName.Text;
    }
}
```

3. Using the Debug menu, launch the page in debug mode. The first time you launch the debugger for a web project, you will see a dialog box asking if you want to enable debugging in the Web.config file. Click the OK button to enable debugging. When the page displays in the browser, enter your name and click the Submit button. Verify that the page redisplay with your name included in the Hello... message. After testing, close the browser.

Creating Non-PostBack Server Control Events

To create non-postback server control events, follow these steps:

4. Using the Page Designer, add a RadioButtonList control to Page1. Change its positioning to absolute and place it below the textbox.
5. Change the ID property to rblFontColor. Click the Items collection in the Properties window to display the ListItem Collection Editor dialog box. Add a Black item and a Red item to the list (see Figure 12-14).




Figure 12-14. Adding list items

- In the Properties window for the `rb1FontCoLor`, click the events button to display the event list (see Figure 12-15).




Figure 12-15. Selecting control events

7. Double click on the `SelectedIndexChanged` event. This will take you to the code behind page and set up the method to handle the event. Add the following code to change the font color of the `lblMessage` control in the event handler:

```
protected void rblFontColor_SelectedIndexChanged(object sender, EventArgs e)
{
    if (rblFontColor.SelectedItem.Value == "Red")
    {
        lblMessage.ForeColor = System.Drawing.Color.Red;
    }
    else
    {
        lblMessage.ForeColor = System.Drawing.Color.Black;
    }
}
```

8. Using the Debug menu, launch the page in debug mode. When the page displays in the browser, click the Red radio button. Notice that the event does not cause a postback to occur. This event will queue until an `AutoPostBack` event triggers a post back.
9. Enter your name in the `txtName` control and click the Submit button. When the postback is processed, both the `SelectedIndexChanged` event of `rblFontColor` and the `Click` event of the Submit button are processed.
10. After testing, close the browser and exit Visual Studio.

■ **Note** This section has explored only server-side event handling. If the browser supports client-side scripting and Dynamic HTML, you can add JavaScript code to `.aspx` pages to take advantage of the client-side processing capabilities.

Storing and Sharing State in a Web Application

Web applications use a stateless protocol (HTTP) to communicate between the browser and the server. In many Web applications, however, some sort of state maintenance needs to exist during a user's session. Traditionally, maintaining state efficiently has been challenging for Web programmers. To help alleviate the challenge of state management, the .NET Framework provides several options to manage state on both the client and the server.

Maintaining View State

One type of state management that needs to occur in a Web application is the preservation of control property values when a form posts back to the client. Because a Web page is re-created and destroyed each time the page is requested, any changes to control made by the client are lost when the page is posted back. For example, any information entered into a `TextBox` control would be lost when posted back. To overcome this type of state loss, the server controls have a `ViewState` property, which provides a dictionary object used to retain the value of the control between post backs. When a page is processed, the current state of the page and the controls are hashed into a string and saved as a

hidden field on the page. During a post back, the values are retrieved and restored when the controls are initialized for rendering of the page back to the browser. The following code shows the ViewState hidden control that gets rendered to the browser:

```
<input type="hidden " name="__VIEWSTATE"
value="dDw1Njg2NjE2ODY7O2w8XNOMTowOz4+upC3lZ6nNLX/ShtpGHJAmi8mpH4=" />
```

Using Query Strings

While view state is used to store information between post backs to the same page, you often need to pass information from one page to another. One popular way to accomplish this is through the use of query strings. A query string is appended to the page request URL; for example, the following URL passes an ID field to the page where it can be used to look up values in a database.

<http://LocalHost/LookUpAuthorInfo.aspx?AuthorId=30>

You can pass more than one query string value by separating the key value pairs with an ampersand.

<http://LocalHost/LookUpAuthorInfo.aspx?AuthorId=30&BookTitleId=355>

To retrieve the values from the query string, you use the `QueryString` property of the `Request` object.

```
int ID = int.Parse(Request.QueryString["AuthorId"]);
```

■ **Caution** One downside of using query strings is that they are readily visible in the browser and users can easily change the values. Never pass sensitive data using query strings.

Using Cookies

You can use cookies to store small amounts of data in a text file located on the client device. The `HttpResponse` class's `Cookie` property provides access to the `Cookies` collection. The `Cookies` collection contains cookies transmitted by the client to the server in the `Cookies` header. This collection contains cookies originally generated on the server and transmitted to the client in the `Set-Cookie` header. Because the browser can only send cookie data to the server that originally created the cookie, and cookie information can be encrypted before being sent to the browser, it is a fairly secure way of maintaining user data. A common use for cookies is to send a user identity token to the client that can be retrieved the next time the user visits the site. This token is then used to retrieve client specific information from a database. The use of cookies is a good way to maintain client state between visits to the site. In the following code, a `Cookie` object is created to hold the date and time of the user's visit. This `Cookie` object is then added to the `Cookies` collection and sent to the browser:

```
HttpCookie visitDate = new HttpCookie("visitDate");
DateTime now = DateTime.Now;
visitDate.Value = now.ToString();
visitDate.Expires = now.AddHours(1);
Response.Cookies.Add(visitDate);
```

To read a cookie, you use the Request object. The following code reads the cookies value and loads it into a Textbox control.

```
if (Request.Cookies["visitDate"] != null)
{
    txtLastVisit.Text = Request.Cookies["visitDate"].Value;
}
```

Maintaining Session and Application State

Another type of state management often needed in Web applications is session state. Session state is the ability to maintain information pertinent to a user as they request the various pages within a Web application. Session state is maintained on the server and is provided by an object instance of the `HttpSessionState` class. This class provides access to session state values and session-level settings for the Web application. Session state values are stored in a key-value dictionary structure only accessible by the current browser session. The following code checks the `Session` object to see if a phone number for the current user is stored in a session key-value pair. The `Session` object exposes an object instance of the `HttpSessionState` class:

```
if (Session["phone"]==null)
{
    txtphone.Visible = true;
}
```

By default, session state is stored in the memory of the web server. There are times when you want session state to be more durable and maintain values in cases of server restarts. In these cases, you can store session values in a SQL Server database or an ASP.NET session state server.

Although session state is scoped on a per-session basis, there are times a Web application needs to share a state among all sessions in the application. You can achieve this globally scoped state using an object instance of the `HttpApplicationState` class. The application state is stored in a key-value dictionary structure similar to the session state, except that it is available to all sessions and from all forms in the application. The first time a client requests any URL resource from the Web application, a single instance of an `HttpApplicationState` class is created. This instance is exposed through the intrinsic `Application` object. The following code uses the `Application` object to store a connection string attribute for a `SqlConnection` object:

```
string pubsConnectionString;
pubsConnectionString = "Integrated Security=SSPI;Initial Catalog=pubs;" +
    "Data Source=localhost";
Application["PubsConnectionString"] = pubsConnectionString;
```

■ **Note** Application state can also be stored in cache using the `Cache` class in the `System.Web.Caching` namespace.

ACTIVITY 12-2. MAINTAINING STATE IN A WEB APPLICATION

In this activity you will become familiar with the following:

- how view state is maintained in a web form
- how to use cookies to store and retrieve data
- the difference between Session and Application state

Maintaining View State

To investigate view state, follow these steps:

1. Start Visual Studio. On the File menu, choose New Web Site.
2. Choose an ASP.NET Empty Web Site. Change the name of the web site at the end of location path Activity12_2.
3. After the Web Site is created, add a Web Form named Page1.aspx. On the Web form add a Textbox control and a Button control using the following mark up.

```
<form id="form1" runat="server">
<div>
    <asp:TextBox ID="TextBox1" runat="server"></asp:TextBox>
    <asp:Button ID="Button1" runat="server" Text="Button" />
</div>
</form>
```

4. Launch the page using the debugger. After the page loads, right-click on it and chose View source. You should see the mark up for a hidden control used to store the view state.

```
<div class="aspNetHidden">
<input type="hidden" name="__VIEWSTATE" id="__VIEWSTATE"
value="s8jATaLWa00KhQ+roB18XjQi6d2d0KSSbBx8tZAvxtIM8ih2HwPQotse+N13fzpd3g+
w7fbor7Hi2uGUuJ5+H+IomzQJHTf5011dSwwLD3s=" />
</div>
```

5. Stop the debugger and add a click event handler to the button.

```
<asp:Button ID="Button1" runat="server" Text="Button" OnClick="Button1_Click"/>
```

6. In the code behind class (Page1.aspx.cs) add the method to handle the button click event. In this event you will save and retrieve a value from the view state that records how many times the button is clicked.

```
protected void Button1_Click(object sender, EventArgs e)
{
    int count;
    if (IsPostBack)
    {
```

```

        if (ViewState["count"] != null)
        {
            count = (int)ViewState["count"];
            count += 1;
            ViewState["count"] = count;
        }
        else
        {
            count = 1;
            ViewState.Add("count", count);
        }
    }
    Response.Write(ViewState["count"]);
}

```

7. Launch the page using the debugger. Click on the button; you should see a counter at the top of the page increment with each click. When done testing stop the debugger.

Reading and Writing Cookies

To investigate using cookies, follow these steps:

1. Add a second Web Form to the project named Page2.
2. On Page1 add a second button and add the following code to the button's click event handler. This code saves the textbox text to a cookie. It then transfers execution to Page2.

```

protected void Button2_Click(object sender, EventArgs e)
{
    HttpCookie visitor = new HttpCookie("visitor");
    visitor.Expires = DateTime.Now.AddDays(1);
    visitor.Value = TextBox1.Text;
    Response.Cookies.Add(visitor);
    Server.Transfer("Page2.aspx");
}

```

3. In Page2's Page_Load method, retrieve the value of the cookie using the Request object and write it out to the page using the Response object.

```

protected void Page_Load(object sender, EventArgs e)
{
    if (Request.Cookies["visitor"] != null)
    {
        Response.Write("Hello " + Request.Cookies["visitor"].Value);
    }
}

```

4. Launch the Page1 using the debugger. Type your name in the textbox and click on the second button; you should see Page2 displayed with the hello message. When done testing stop the debugger.

Session and Application State

To investigate session and application state, complete the following steps:

1. Using the Add New Item Dialog, add a Global Application Class named Global.asax.
2. In the Global.asax code page add the following code to the `Session_Start` event handler.

```
void Session_Start(object sender, EventArgs e)
{
    Session["SessionStartTime"] = DateTime.Now;
}
```

3. Add a third Web Form to the project named Page3.
4. In the `Form_Load` event handler for Page3 add the following code to display the session start time.

```
protected void Page_Load(object sender, EventArgs e)
{
    Response.Write(Session["SessionStartTime"]);
}
```

5. In the Solution Explorer window right-click Page3.aspx and select View in Browser. You should see the session start time displayed.
6. After 30 seconds or so refresh the page. The time should not change because you are in the same session.
7. Launch a new instance of the browser and copy the web address for Page3 from the first browser window to the second browser. You should see Page3 displayed with a new time because this is a new session. When done testing close the browsers.

8. In the Global.asax code page add the following code to the `Application_Start` event handler.

```
void Application_Start(object sender, EventArgs e)
{
    Application["ApplicationStartTime"] = DateTime.Now;
}
```

9. In the `Form_Load` event handler for Page3 add the following code to display the application start time.

```
protected void Page_Load(object sender, EventArgs e)
{
    Response.Write(Session["SessionStartTime"]);
    Response.Write("<br/>");
    Response.Write(Application["ApplicationStartTime"]);
}
```

10. In the Solution Explorer window right-click Page3.aspx and select View in Browser. You should see both the session and application start times displayed.

11. Launch a new instance of the browser and copy the web address for Page3 from the first browser window to the second browser. You should see Page3 displayed with a different session start time but the same application start time because this is a new session but the same application start time. When done testing close the browsers and exit Visual Studio.

Data-Bound Web Controls

Just as with traditional Windows based GUIs, programs developed using a Web based GUI often need to interact with data. Users of business applications need to view, query, and update data held in a backend data store. Fortunately, the .NET Framework makes this easy by exposing the functionality needed to implement data binding controls in Web Forms.

Data-bound controls automate the process of presenting data in a web form. ASP.NET provides various controls to display the data depending on the type of data, how it should be displayed, and whether it can be updated. Some of these controls are used to present read-only data, for example, the Repeater control displays a list of read-only data. If you need to update the data in a list, you could use the ListView control. If you need to display many columns and rows of data in a tabular format, you would use the GridView control. If you need to display a single record at a time, you could use the DetailsView or the FormView controls.

Data binding automates the process of transferring the data from the data store and a data-bound control. In order to bind the controls, you add a DataSource control to the web page. DataSource controls do not render UI markup. Instead, they are responsible for providing the link between the data source and the data-bound control. There are several types of data source controls available depending on the data source. Table 12-3 lists the types of DataSources provided by the .NET Framework.

In the following markup, a SqlDataSource control is used to provide data to a Repeater control that repeats a label control for each name. Notice the use of the data binding expressions Eval and Bind. Eval is used for one-way binding (read-only) and Bind is used for two way binding (read and update)

Table 12-3. DataSources provided by the .NET Framework

DataSource	Use
EntityDataSource	Enables you to bind to data that is based on the Entity Data Model.
SqlDataSource	Enables you to work with Microsoft SQL Server, OleDb, ODBC, or Oracle databases
ObjectDataSource	Enables you to work with a business object or other class
LinqDataSource	Enables you to use Language-Integrated Query
XmlDataSource	Enables you to work with an XML file.
SiteMapDataSource	Used with ASP.NET site navigation.

```
<asp:SqlDataSource ID="SqlDataSource1" runat="server"
    ConnectionString="<%"$ ConnectionStrings:pubsConnectionString %">
    SelectCommand="SELECT [au_id], [au_lname] FROM [authors]">
</asp:SqlDataSource>
```

```
<asp:Repeater ID="Repeater1" runat="server" DataSourceID="SqlDataSource1">
  <ItemTemplate>
    <asp:Label ID="lblName" runat="server" Text='<%= Eval("au_lname") %>'></asp:Label>
    <br />
  </ItemTemplate>
</asp:Repeater>
```

DataSource controls are convenient for rapid application development. They allow you to create web pages that support CRUD (Create Read Update Delete) operations while writing a minimum amount of code.

Model Binding

While DataSource controls work well for simple CRUD operations and simple applications, they are inflexible and cumbersome to work with as your application logic and data access logic become more complex. They mix the data retrieval logic with the presentation logic, which violates the n tier development principle of separating data access logic, business logic, and presentation logic into distinct tiers. In order to overcome these limitations, ASP.NET includes Web Forms model binding.

Model binding is a more code focused data access logic used to data bind controls. With model binding you declare the type of data you are binding a control to using the ItemType property. When the ItemType property of a control is set, the Item and the BindItem expressions are used to bind the control to a data source item. The Item expression implements one-way binding (read-only) while the BindItem expression is used for two-way binding (read and update). The following code shows using a repeater control to host a label control and bind its Text property to the Employee.Name property. Notice the ItemType property of the Repeater control is set to the Employee class.

```
<asp:Repeater ID="Repeater1" runat="server"
  ItemType="Chapter12Demo1.Employee" SelectMethod="GetEmployee">
  <ItemTemplate>
    <asp:Label ID="Label1" runat="server" Text='<%= Item.Name %>'></asp:Label>
    <br />
  </ItemTemplate>
</asp:Repeater>
```

To configure a data control to use model binding to select data, you set the control's SelectMethod property to the name of a method in the page's code. In the case above, the SelectMethod is set to the GetEmployee method which returns a list of Employee objects. The data control calls the method at the appropriate time in the page life cycle and automatically binds the returned data.

To implement two-way binding, you include a method to update the data and set it equal to the UpdateMethod property of the control. The Model Binding will send the updated values to the update method where the values get updated back to the data source. The following code defines a GridView that uses Model Binding to update employees. Notice the use of BindItem and the UpdateMethod to implement two-way binding.

```
<asp:GridView ID="employeesGrid" runat="server" DataKeyNames="ID"
  ItemType="Chapter12Demo1.Employee" AutoGenerateColumns="false" AutoGenerateEditButton="true"
  SelectMethod="GetEmployees" UpdateMethod="UpdateEmployee">
  <Columns>
    <asp:TemplateField>
      <ItemTemplate>
        <asp:Label ID="lblID" runat="server" Text='<%= Item.ID %>'></asp:Label>
      </ItemTemplate>
    </asp:TemplateField>
```

```

<asp:TemplateField>
  <ItemTemplate>
    <asp:Label ID="lblName" runat="server" Text='<%= BindItem.Name %>'></asp:Label>
  </ItemTemplate>
  <EditItemTemplate>
    <asp:TextBox ID="txtName" runat="server" Text='<%= BindItem.Name %>'></asp:TextBox>
  </EditItemTemplate>
</asp:TemplateField>
</Columns>
</asp:GridView>

```

■ **Note** For new development that targets ASP.NET 4.5, Microsoft recommends using Model Binding. For applications targeting previous versions, DataSource controls are the recommended method.

ACTIVITY 12-3. BINDING A GRIDVIEW CONTROL

In this activity, you will become familiar with the following:

- how to implement one-way binding to a Repeater control using a DataSource control
- how to implement two-way binding to a GridView control using Model Binding

Displaying Data in an ASP.NET Repeater Control

To display data in an ASP.NET Repeater control, follow these steps:

1. Start Visual Studio. On the File menu, choose New ► Web Site.
2. Choose an ASP.NET Empty Web Site. Change the name of the web site at the end of location path Activity12_3.
3. Add a new Web Form to the site and Name it AuthorListPage.
4. Open the AuthorListPage in Split view. Drag a SqlDataSource control located under the Data node to the form. You should see a SqlDataSource Tasks pane as in figure 12-16.




Figure 12-16. The SqlDataSource Task Pane

5. Click on the Configure Data Source link. This launches a wizard to help configure the DataSource control.
6. Create a new connection to your SQL Server using the .NET Framework Data Provider for SQL Server (see Figure 12-17).

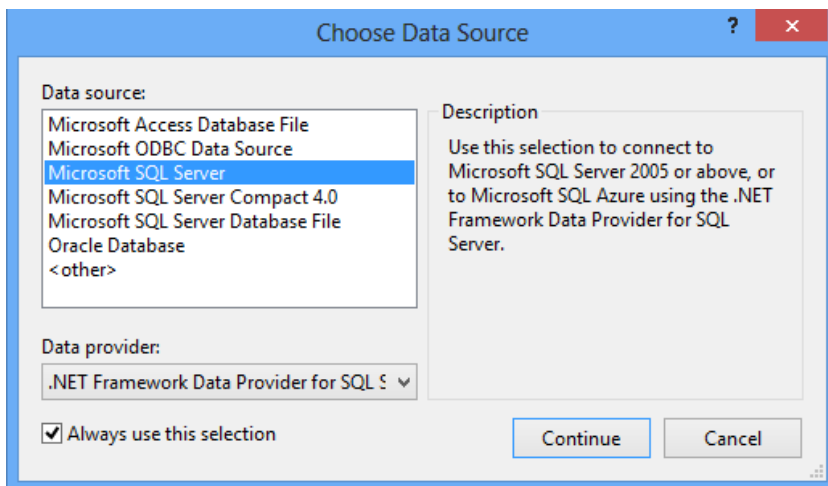


Figure 12-17. Selecting the SQL Server Data Source

7. In the Add Connection dialog, enter your server name and select the Pubs database. Test the connection and click the OK button.
8. In the next screen, check yes to save the connection string in the configuration file.
9. In the Configure the Select Statement screen select the au_id and au_lname columns from the authors table as in Figure 12-18.

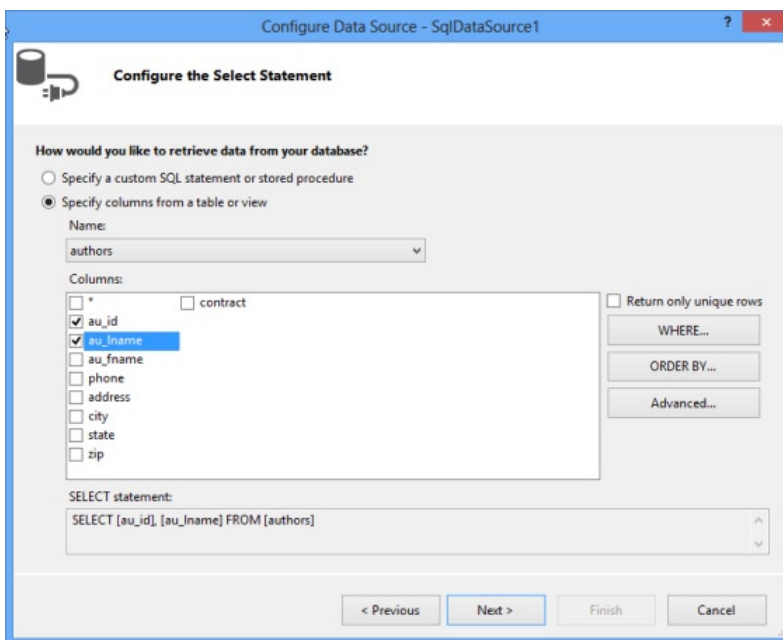


Figure 12-18. Configuring the Select Statement

10. In the next screen you can test the query. After testing the query click the Finish button.
11. In the page source you should see the mark up added to the page to set up the SqlDataSource control.

```
<asp:SqlDataSource ID="SqlDataSource1" runat="server"
    ConnectionString="<%= $ ConnectionStrings:pubsConnectionString %>"
    SelectCommand="SELECT [au_id], [au_lname] FROM [authors]">
</asp:SqlDataSource>
```

12. In the Design view for the AuthorListPage drag a Repeater control from the Data node in the Toolbox to the form. Place it after the SqlDataSource control.
13. In the Repeater Tasks pop out window select the SQLDataSource1 control.
14. Switch to the Source window and add the ItemTemplate mark up and the asp:Label mark up inside the Repeater control mark up. Notice the Text property of the label is bound to the au_lname column using the Eval method.

```
<asp:Repeater ID="Repeater1" runat="server" DataSourceID="SqlDataSource1">
    <ItemTemplate>
        <asp:Label ID="lblName" runat="server" Text='<%=# Eval("au_lname") %>'>
        </asp:Label>
        <br />
    </ItemTemplate>
</asp:Repeater>
```

15. Right click on the AutherListPage.aspx node in the Solution Explorer window and select “Set as Start Page”. Start the debugger. You should see the page displayed with a list of author last names.
16. When done testing, stop the debugger.

Displaying Data using an ASP.NET GridView Control

To display and update data using an ASP.NET GridView control, follow these steps:

1. In the Solution Explore window, right click the project node and select Add ► Add New Item. Add an ADO.NET Entity Data Model and Name it Pubs.edmx. You will get a message about adding it to the App_Code folder. Select Yes.
2. In the Entity Data Model Wizard select Generate from database.
3. In the Chose Your Data Connection, you can either reuse the connection created previously or create a new connection to the Pubs database. Be sure the save entity connection settings in the Web.Config is checked and it is named pubsEntities.
4. In the Choose Your Database Objects and Settings window, select the employee table. Leave the default check boxes checked and name the model pubsModel (see Figure 12-19).




Figure 12-19. Selecting the Data Objects

5. After completing the wizard you should see the employee entity in the Entity Designer. Close the Entity designer and add a new Web Form to the designer. Name the page EmployeeGridViewPage.aspx.
6. Add a GridView control to the form and add the mark up below. Notice you are setting the ItemType to the employee entity and the SelectMethod to GetEmployees.

```
<asp:GridView ID="grdEmployee" runat="server"
    ItemType="employee" SelectMethod="GetEmployees"
    AutoGenerateColumns="false">
    <Columns>
        <asp:BoundField DataField="emp_id" HeaderText="ID" ReadOnly="true" />
        <asp:BoundField DataField="fname" HeaderText="First Name" />
        <asp:BoundField DataField="lname" HeaderText="Last Name" />
    </Columns>
</asp:GridView>
```

7. In the code behind file for the page (EmployeeGridViewPage.aspx.cs) add the following code to retrieve the employee data using the employee entity.

```
public partial class EmployeeGridViewPage : System.Web.UI.Page
{
    pubsEntities dbContext = new pubsEntities();
    protected void Page_Load(object sender, EventArgs e)
    {

    }
    public IQueryable<employee> GetEmployees()
    {
        var Employees = dbContext.employees;
        return Employees;
    }
}
```

8. Right click on the EmployeesGridViewPage.aspx node in the solution Explorer window and select "Set as Start Page". Start the debugger. You should see the page displayed with a table of employee information.
9. When done testing, stop the debugger.

Updating Data using an ASP.NET GridView Control

1. Update the GridView control's mark up to allow updating and set the UpdateMethod attribute to the UpdateEmployee method. The DataKeyNames attribute is set to the unique identifier of the employee entity.

```
<asp:GridView ID="grdEmployee" runat="server"
    ModelType="employee" SelectMethod="GetEmployees" DataKeyNames="emp_id"
    AutoGenerateEditButton="true" UpdateMethod="UpdateEmployee"
    AutoGenerateColumns="false">
    <Columns>
        <asp:BoundField DataField="emp_id" HeaderText="ID" ReadOnly="true" />
        <asp:BoundField DataField="fname" HeaderText="First Name" />
        <asp:BoundField DataField="lname" HeaderText="Last Name" />
    </Columns>
</asp:GridView>
```

2. In the code behind file for the page (EmployeeGridViewPage.aspx.cs) add the following code to update the employee data using the employee entity. This code looks up the old values held in the entity and tries to update the model. If it is successful it saves the changes back to the database.

```
public void UpdateEmployee(employee Employee)
{
    var Employee2 = dbContext.employees.Find(Employee.emp_id);
    TryUpdateModel(Employee2);
}
```

```

        if (ModelState.IsValid)
        {
            dbContext.SaveChanges();
        }
    }
}

```

3. By default Web Forms have `UnobtrusiveValidationMode` turned on. To implement this you need to add a `ScriptManager` and some `jQuery` include files. For now you can just turn it off by adding the following to the `Page_Load` event handler method.

```
this.UnobtrusiveValidationMode = System.Web.UI.UnobtrusiveValidationMode.None;
```

4. Start the debugger. You should see the page displayed with a table of employee information and an Edit button. Click on the Edit button for a row. This sets the grid in edit mode.
 5. Update the first name and click the Update button. The grid will refresh and show the new values.
-

Summary

In this chapter you looked at implementing the interface tier of an application using a Web Form-based front end. Along the way, you took a closer look at the classes and namespaces of the .NET Framework that are used to implement Web Forms-based user interfaces. You were also exposed to data binding Web server controls, in particular, the `GridView` control. My intent in this chapter was to expose you to some of the techniques used to develop data centric web applications. To become a skilled web developer, you need to dig deeper into these topics and other topics such as using Master pages, client side scripting, using style sheets, and web security.

In Chapter 13, you will look at a new type of application introduced with Windows 8 - the Windows Store app. Windows Store apps are a sort of hybrid between traditional client apps and web apps. The interface is developed using HTML or XAML and the application logic can be developed in C#, Visual Basic, or JavaScript. They support the ever-popular touch screen devices and represent a shift in the way users interact with your applications.



Developing Windows Store Applications

In the previous chapters, you learned how to build a data-centric application based on a traditional Windows client (WPF) user interface and a web-based (ASP.NET) user interface. In this chapter, you will learn how to build a user interface using the new Windows Store app. Windows Store apps are designed to run on Windows 8 devices. Windows Store apps have a new look and feel designed to dynamically support different display sizes and devices. As tablets and phones increasingly become the most popular devices for interacting with information, business users are demanding the ability to use these devices for interacting with their information stores. Fortunately, Microsoft offers a variety of application programming interfaces (APIs) that can be used to create applications to run on these devices. They have exposed a managed .NET framework that allows you to use C# for the code behind and a XAML framework for developing the user interface. Developing these apps is a cross between developing WPF apps and web apps.

After reading this chapter, you will be comfortable performing the following tasks:

- using XAML markup to design a user interface
- working with layout controls
- working with display controls
- responding to control events
- using data-bound controls

Building the User Interface

Building the UI of a Windows Store app is very similar to developing WPF applications. If you recall from Chapter 11, the visual interface of an application contains objects. These objects, like most objects you work with in object-oriented languages, expose properties, methods, and events. In a Windows Store app, the main object is a page. A page has properties (for example the Background property), methods (for example the focus method), and events (for example the DoubleTapped event).

Controls are components with visual interfaces that give users a way to interact with the program. A page is a container control, which hosts other controls. You can place many different types of controls on pages. Some common controls used on pages are TextBoxes, TextBlocks, Buttons, ListViews, and GridViews.

Just like WPF, Windows Store user interfaces are built using the declarative markup language XAML. For example, the following markup defines a button control inside a Grid.

```
<Grid Background="#E5951D1D">
  <Button Content="Button" HorizontalAlignment="Left" Margin="56,124,0,0"
    VerticalAlignment="Top" Background="#FF2457A0"/>
</Grid>
```

Notice the Grid needs a formal closing tag because it contains the Button control. Since the Button control does not contain any other controls, you can use a forward slash (/) in front of the end bracket to close it. Notice the properties of the controls are set by using attribute syntax.

Controls are positioned either absolutely or relatively using grid rows and columns. Notice the code above uses absolute positioning of the button in the grid which is achieved using the Margin attribute. The code below shows the button positioned relatively by using grid rows and columns.

```
<Grid Background="#E5951D1D">
  <Grid.RowDefinitions>
    <RowDefinition Height="140"/>
    <RowDefinition Height="*" />
  </Grid.RowDefinitions>

  <Grid.ColumnDefinitions>
    <ColumnDefinition Width ="140"/>
    <ColumnDefinition Width="*" />
  </Grid.ColumnDefinitions>

  <Button Content="Button" HorizontalAlignment="Left"
    Grid.Column="1" Grid.Row="1"
    VerticalAlignment="Top" Background="#FF2457A0"
    Height="78" Width="162" />
</Grid>
```

Figure 13-1 shows the page with the button created by the previous XAML code.




Figure 13-1. A window created with XAML

As seen previously, a Grid control contains columns and rows to control the placement of its child controls. The height and width of the columns and rows can be set to a fixed value, auto, or *. The auto setting takes up as much space as needed by the contained control. The * setting takes up as much space as is available. The following code lays out a simple data entry form used to collect user information. The resulting form is shown in Figure 13-2.

```
<Grid Background="#E5B1BDB9">
    <Grid.RowDefinitions>
        <RowDefinition Height="Auto" />
        <RowDefinition Height="Auto" />
        <RowDefinition Height="Auto" />
        <RowDefinition Height="Auto" />
    <RowDefinition Height="Auto" />
        <RowDefinition Height="*" />
    </Grid.RowDefinitions>
    <Grid.ColumnDefinitions>
        <ColumnDefinition Width="Auto" />
        <ColumnDefinition Width="200" />
        <ColumnDefinition Width="*" />
    </Grid.ColumnDefinitions>
    <TextBlock Grid.Row="0" Grid.Column="0" Text="Name:"
        Foreground="Black" VerticalAlignment="Center"
        Margin="10" FontSize="16"/>
    <TextBlock Grid.Row="1" Grid.Column="0" Text="City:"
        Foreground="Black" VerticalAlignment="Center"
        Margin="10" FontSize="16"/>
    <TextBlock Grid.Row="2" Grid.Column="0" Text="State:"
        Foreground="Black" VerticalAlignment="Center"
        Margin="10" FontSize="16" />

    <TextBox Grid.Column="1" Grid.Row="0" Margin="3" />
    <TextBox Grid.Column="1" Grid.Row="1" Margin="3" />
    <TextBox Grid.Column="1" Grid.Row="2" Margin="3" />
    <Button Grid.Column="1" Grid.Row="4" HorizontalAlignment="Right"
        MinWidth="80" MinHeight="50" Margin="0,0,0,8" Content="Submit"
        Foreground="Black"/>
</Grid>
```

The screenshot shows a data entry form on a light gray background. It consists of three rows of text labels followed by input fields. The first row is 'Name:' followed by a white text box with a vertical cursor. The second row is 'City:' followed by a white text box. The third row is 'State:' followed by a white text box. At the bottom right of the form is a white button with a black border and the text 'Submit' in black.

Figure 13-2. Input form page

Some other layout controls available are the `StackPanel`, `WrapGrid`, and `Canvas`. The following code shows two buttons in a `StackPanel` control:

```
<StackPanel Grid.Column="1" Grid.Row="4" Orientation="Horizontal" HorizontalAlignment="right">
    <Button MinWidth="80" MinHeight="50" Margin="0,0,0,8" Content="Submit"
        Foreground="Black"/>
    <Button MinWidth="80" MinHeight="50" Margin="0,0,0,8" Content="Cancel"
        Foreground="Black"/>
</StackPanel>
```

Using Style Sheets

An import aspect of developing Windows Store apps is maintaining a consistent look and feel across pages and applications. The interface needs to scale across large and small displays and handle changes in orientation smoothly. The creation and styling of these apps can be a time-consuming, tedious process. Fortunately, there are pre-built styles that you can use across all pages and controls in your application. These styles are contained in a XAML file and are referenced by the controls using a Static Resource Markup Extension. The following code uses a static resource to style two button controls.

```
<Button Content="Submit" Style="{StaticResource TextPrimaryButtonStyle}" />
<Button Content="Cancel" Style="{StaticResource TextSecondaryButtonStyle}"/>
```

The following markup defining the styles is in the `StandardStyles.xaml` file located in the `Common` folder. Notice you can reference a style within a style definition.

```
<Style x:Key="TextPrimaryButtonStyle" TargetType="ButtonBase"
    BasedOn="{StaticResource TextButtonStyle}">
    <Setter Property="Foreground"
        Value="{StaticResource ApplicationHeaderForegroundThemeBrush}"/>
</Style>
```

Handling Control Events

Just like WPF and ASP.NET apps, Windows Store apps interact with users through events. The difference is how the events are initiated. While WPF and ASP.NET applications are largely driven by mouse and keyboard events, Windows Store apps are largely driven by touch interactions (keyboard and mouse events are still supported). Some common interacting actions include tapping, press and hold, sliding, and swiping.

In Visual Studio, you can add an event to a control either by writing XAML code or by selecting it in the control's Properties window. Figure 13-3 shows wiring up an event handler in the XAML Editor window; Figure 13-4 shows wiring up an event handler using the Events tab of the Properties window. Remember when working with controls in code, you need to give them a unique name using the `Name` attribute.

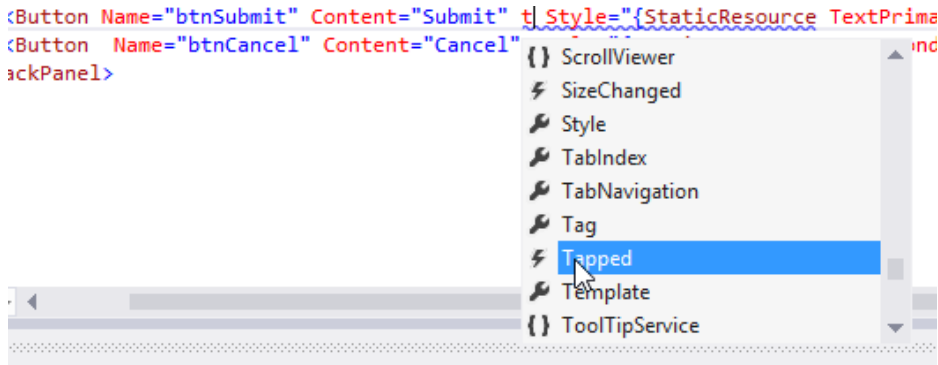


Figure 13-3. Wiring up an event handler in the XAML editor

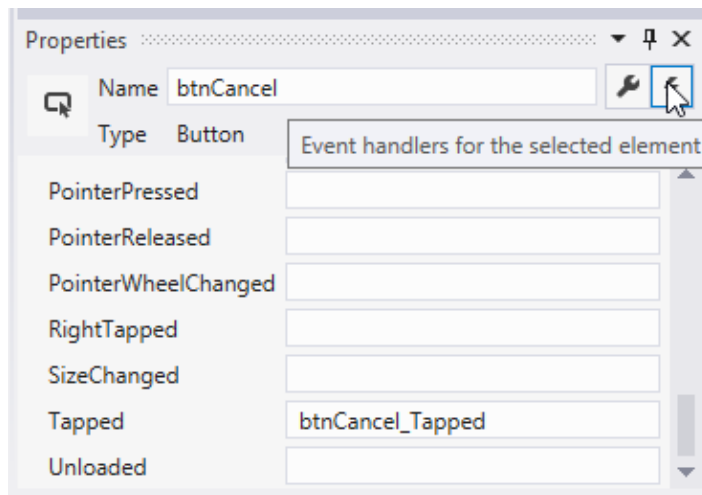


Figure 13-4. Wiring up an event handler in the Properties window

Regardless of how you wire up an event handler, the Visual Studio code editor inserts an empty event handler method in the code behind file. The following code shows the event handler method inserted for the button tap event:

```
private void btnSubmit_Tapped(object sender, TappedRoutedEventArgs e)
{
}
```

Just as in WPF, two parameters are passed to the method when the event is fired. The first parameter is the sender, which represents the object that initiated the event. The second parameter, in this case of type `Windows.UI.Xaml.Input.RoutedEventArgs`, is an object used to pass any information specific to the particular event. The following code shows checking if they tapped with a pen.


```
private void btnSubmit_Tapped(object sender, TappedRoutedEventArgs e)
{
    if (e.PointerDeviceType == PointerDeviceType.Pen)
    {
        //additional code...
    }
}
```

In the following activity, you will investigate working with controls in a simple Windows Store app.

ACTIVITY 13-1. WORKING WITH CONTROLS

In this activity, you will become familiar with the following:

- creating a Windows Store App
- laying out a Page
- working with Control events

Creating the Windows Store App Interface

To create the application interface, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Click on the Windows Store node under the Visual C# template node. Select the Blank App template. Change the name to Activity13_1 and click the OK button.
3. Open the MainPage.xaml in the design window. Notice it is using the default dark theme. To change this to the light theme, open the App.xaml page. Add the RequestedTheme attribute to the Application mark up.

```
<Application
  x:Class="Activity13_1.App"
  xmlns="http://schemas.microsoft.com/winfx/2006/xaml/presentation"
  xmlns:x="http://schemas.microsoft.com/winfx/2006/xaml"
  xmlns:local="using:Activity13_1"
  RequestedTheme="Light">
```

4. Close and reopen the MainPage.xaml file. You should now see a white background.
5. Update the Grid XAML to define two columns and 5 rows.

```
<Grid Background="{StaticResource ApplicationPageBackgroundThemeBrush}">
  <Grid.ColumnDefinitions>
    <ColumnDefinition Width="300"/>
    <ColumnDefinition Width="Auto"/>
  </Grid.ColumnDefinitions>
```

```

        <Grid.RowDefinitions>
            <RowDefinition Height="Auto"/>
            <RowDefinition Height="Auto"/>
            <RowDefinition Height="Auto"/>
            <RowDefinition Height="Auto"/>
            <RowDefinition Height="Auto"/>
        </Grid.RowDefinitions>

    </Grid>
</Grid>

```

6. Add the following XAML after the closing `Grid.RowDefinitions` tag to put two text boxes on the page. One accepts a single line of text and one accepts multiple lines of text.

```

<TextBlock Grid.Row="0" Grid.Column="0" Margin="10"
            Style="{StaticResource BasicTextStyle}">
    Single line text input
</TextBlock>
<TextBox Grid.Row="0" Grid.Column="1" Margin="10"
          Width="200" HorizontalAlignment="Left"/>
<TextBlock Grid.Row="1" Grid.Column="0" Margin="10"
            Style="{StaticResource BasicTextStyle}">
    Multi-line input
</TextBlock>
<TextBox Grid.Row="1" Grid.Column="1" Margin="10"
          Width="200" Height="100" TextWrapping="Wrap"
          AcceptsReturn="True" HorizontalAlignment="Left"/>

```

7. Add a password entry box by inserting the following code.

```

<TextBlock Grid.Row="2" Grid.Column="0" Margin="10"
            Style="{StaticResource BasicTextStyle}">
    Password entry
</TextBlock>
<PasswordBox Grid.Row="2" Grid.Column="1" Margin="10"
              Width="200" HorizontalAlignment="Left"/>

```

8. Add two more `TextBox` controls, one for numeric input and one for email address entry.

```

<TextBlock Grid.Row="3" Grid.Column="0" Margin="10"
            Style="{StaticResource BasicTextStyle}">
    Number input
</TextBlock>
<TextBox Grid.Row="3" Grid.Column="1" Margin="10"
          Width="200" InputScope="Number" HorizontalAlignment="Left"/>
<TextBlock Grid.Row="4" Grid.Column="0" Margin="10"
            Style="{StaticResource BasicTextStyle}" TextWrapping="Wrap">
    Email address
</TextBlock>
<TextBox Grid.Row="4" Grid.Column="1" Margin="10"
          Width="200" InputScope="EmailSntpAddress" />

```

- Note that as you add the XAML, the Visual Designer updates the appearance of the window. The page should look similar to the one shown in Figure 13-5.

Single line text input	<input type="text"/>
Multi-line input	<input type="text"/>
Password entry	<input type="password"/>
Number input	<input type="text"/>
Email address	<input type="text"/>

Figure 13-5. The data entry page

- Build the solution. If there are any errors, fix them and rebuild.
- To test the page, change the debugger target to the Simulator in the toolbar (See Figure 13-6) and start the debugger.

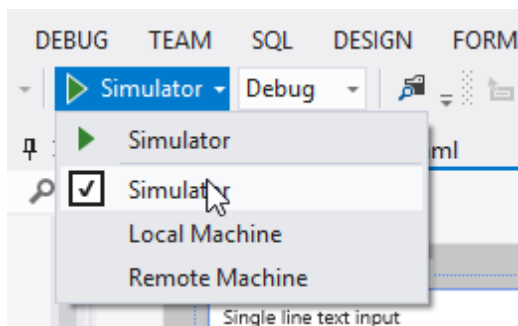


Figure 13-6. Launching the simulator

- When the Simulator launches, change the mode to basic touch (See Figure 13-7). In this mode, the mouse is used to simulate a finger press.




Figure 13-7. Changing the interaction mode

13. When you press in the single line input TextBox, the soft key board is displayed. Enter some text and try to hit enter. You should not be able to add additional lines. Test the multi-line input and the password input.
14. Press in the number input TextBox. Notice the keyboard displays the number pad. It knew to do this because you added the `InputScope` attribute to the control.
15. Pressing in the email address input TextBox adds the “@” and “.com” keys to the keyboard.
16. When finished testing, you can close the Simulator by pressing the `Ctrl + Alt + F4` keys.

Coding Control Events

To code the control events, follow these steps:

1. In the XAML Editor window for the `MainPage.xaml` file in the `Grid.RowDefinitions` section add a new row to the grid. In the new row, place two button controls in a stack panel using the code below.

```
<StackPanel Grid.Row="5" Grid.Column="1" Orientation="Horizontal">
    <Button Name ="btnSave" Content ="Save" Margin="10"
        Tapped="btnSave_Tapped"/>
```

```

        <Button Name="btnCancel" Content ="Cancel" Margin="10"
            Tapped="btnCancel_Tapped"/>
    </StackPanel>

```

2. After the closing `StackPanel` tag, add the following markup to display a progress ring on the page.

```

<ProgressRing x:Name="SaveProgress" Grid.Row="5" Grid.Column="0"
    IsActive="False" Width="113" Height="80"/>

```

3. The XAML defining the buttons above includes an event handler method for the `Tapped` event. Right-click on the `btnSave`'s `Tapped` event and select "Navigate to Event Handler" in the pop-up menu.
4. In the Code Editor window of the codebehind file, add the following code to the `btnSave_Tapped` event handler method. This code makes the progress ring show on the page.

```

this.SaveProgress.IsActive = true;

```

5. In the `btnCancel_Tapped` method, deactivate the progress ring.

```

this.SaveProgress.IsActive = false;

```

6. Build the solution and fix any errors.
 7. Test the buttons' tapped events using the Simulator.
 8. After testing the application, press `Ctrl + Alt + F4` keys to close the simulator and stop debugging.
-

Data Binding Controls

Binding a Windows Store app control to data is done in a way that is very similar to the way it's handled in WPF (covered in Chapter 10). When you do the binding with XAML, you use the `Binding` attribute available with each control. When you bind a control in code, you set its source with the `DataContext` property. When you set the `DataContext` for a parent element, such as a `Grid` control, the child elements will use the same `DataContext` unless their `DataContext` is explicitly set.

The .NET Framework encapsulates much of the complexity of synchronizing controls to a data source through the data binding process. The `Mode` property determines how the data binding flows and reacts to data changes. `OneTime` binding updates the target with data from the source when the binding is created. `OneWay` binding causes changes to the source property to automatically update the target property, but changes to the target property are not propagated back to the source property. This is useful for read-only scenarios and is the default binding. `TwoWay` binding causes changes to either the source property or the target property to automatically update the other. This is useful for full-data updating scenarios.

The following code shows the `DataContext` of a `Grid` control set to a `CollectionViewSource` that contains a list of authors. The `CollectionViewSource` allows you to move through the list of authors.

```

cvs = new CollectionViewSource();
cvs.Source = authorList;
this.AuthorList.DataContext = cvs;
cvs.View.MoveCurrentToFirst();

```

The following XAML code binds a `TextBox` control and a `CheckBox` control contained in the `Grid` control to the properties of the `Authors` class using the `Path` attribute. Using `Binding` to designate the source data means “look up the container hierarchy until a `DataContext` is found.” In this case, the `DataContext` will be the one specified for the `Grid` container.

```

<TextBox Grid.Column="1" Grid.Row="2" Height="23" HorizontalAlignment="Left"
        Margin="3" Name="txtLastName" Text="{Binding Path=LastName}"
        VerticalAlignment="Center" Width="120" />
<CheckBox Name="chkContract" Content="Under Contract"
        IsChecked="{Binding Path=UnderContract}"
        Grid.Row="4" Grid.ColumnSpan="2" FlowDirection="RightToLeft" />

```

Figure 13-8 shows a page containing controls bound to the author data.

Author Info

First Name:

Last Name:

Royalty:

Under Contract

Figure 13-8. Page displaying author data

To move through the records you just use the `CollectionViewSource` `View` methods. The following code shows how to move to the next record.

```

cvs.View.MoveCurrentToNext();

```

While some controls can only bind to one record at a time, other controls, such as the `GridView` control, bind to and display the entire collection. The following code sets the `ItemsSource` of a `GridView` to the list of authors. In this case, it’s not necessary to use a `CollectionViewSource`.

```

this.AuthorDataGrid.ItemsSource = authorList;

```

The following XAML creates the `GridView` and binds the controls contained in the view. The resulting page is shown in Figure 13-9.

```

<GridView Grid.Row="0" Grid.Column="2" x:Name="AuthorGridView"
        Header="Authors" FontSize="14" FontWeight="Bold"
        HorizontalAlignment="Center" >
    <GridView.ItemTemplate>

```

```

<DataTemplate>
  <StackPanel Orientation="Horizontal">
    <TextBlock Text="{Binding FirstName}"
      Style="{StaticResource ItemTextStyle}" Margin="10"/>
    <TextBlock Text="{Binding LastName}" TextWrapping="NoWrap"
      Style="{StaticResource BodyTextStyle}" Margin="10"/>
    <CheckBox Content="UnderContract" IsChecked="{Binding UnderContract}"
      IsEnabled="False" Margin="10" FontSize="10"/>
  </StackPanel>
</DataTemplate>
</GridView.ItemTemplate>
</GridView>

```

Authors

Clive Cussler UnderContract

Steve Berry UnderContract

Kate Morton UnderContract

Karma Wilson UnderContract

Figure 13-9. Page displaying author data in a GridView

In the following activity, you will build a page with controls bound to a collection of Author objects.

ACTIVITY 13-2. WORKING WITH DATA-BOUND CONTROLS

In this activity, you will become familiar with the following:

- binding controls to a collection
- investigating binding modes

Binding Controls to a Collection

To bind controls to a collection, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose a Windows Store Application and select the Blank App template. Rename the project to Activity13_2 and click the OK button.
3. Right-click on the project node in Solution Explorer and choose Add ► Class. Name the class Author.
4. At the top of the class file, add the following using statement:

```
using System.ComponentModel;
```

5. In the Author class, implement the INotifyPropertyChanged interface. This is needed to facilitate binding.

```
class Author : INotifyPropertyChanged
{
    public event PropertyChangedEventHandler PropertyChanged;

    void RaisePropertyChanged(string propertyName)
    {
        var handler = PropertyChanged;
        if (handler != null)
        {
            handler(this, new PropertyChangedEventArgs(propertyName));
        }
    }
}
```

6. Add the following properties. Note that when the values are changed, the PropertyChanged event is raised.

```
string _firstName;
public string FirstName
{
    get { return _firstName; }
    set
```



```

    {
        if (_firstName != value)
        {
            _firstName = value;
            RaisePropertyChanged("FirstName");
        }
    }
}
string _lastName;
public string LastName
{
    get { return _lastName; }
    set
    {
        if (_lastName != value)
        {
            _lastName = value;
            RaisePropertyChanged("LastName");
        }
    }
}
Boolean _underContract;
public Boolean UnderContract
{
    get { return _underContract; }
    set
    {
        if (_underContract != value)
        {
            _underContract = value;
            RaisePropertyChanged("UnderContract");
        }
    }
}
double _royalty;
public double Royalty
{
    get { return _royalty; }
    set
    {
        if (_royalty != value)
        {
            _royalty = value;
            RaisePropertyChanged("Royalty");
        }
    }
}
}

```

7. Add the following constructor to the Author class:

```
public Author(string firstName, string lastName,
             Boolean underContract, double royalty)
{
    this.FirstName = firstName;
    this.LastName = lastName;
    this.UnderContract = underContract;
    this.Royalty = royalty;
}
```

8. Build the project and make sure there are no errors. If there are, fix them and rebuild.
9. Open the App.xaml page and add the RequestedTheme attribute to the Application mark up.

```
<Application
    x:Class="Activity13_2.App"
    xmlns="http://schemas.microsoft.com/winfx/2006/xaml/presentation"
    xmlns:x="http://schemas.microsoft.com/winfx/2006/xaml"
    xmlns:local="using:Activity13_2"
    RequestedTheme="Light">
```

10. Add the following XAML markup to the MainPage.xaml file to create the page header:

```
<Grid Background="{StaticResource ApplicationPageBackgroundThemeBrush}">

    <Grid.RowDefinitions>
        <RowDefinition Height="140"/>
        <RowDefinition Height="*"/>
    </Grid.RowDefinitions>
    <Grid>
        <Grid.ColumnDefinitions>
            <ColumnDefinition Width="Auto"/>
            <ColumnDefinition Width="*"/>
        </Grid.ColumnDefinitions>
        <TextBlock x:Name="PageTitle"
            Text="Authors" Grid.Column="1" VerticalAlignment="Center"
            Margin="120,20,0,0" IsHitTestVisible="false"
            Style="{StaticResource PageHeaderTextStyle}"/>
    </Grid>
    <!-- Add the GridView here -->
</Grid>
```

11. You can define the DataTemplate under the ResourceDictionary of the Page.Resources section of the page. Add the following XAML markup to define the DataTemplate for the author information contained in the GridView (Add this above the first Grid tag).

```
<Page.Resources>
    <ResourceDictionary>
        <DataTemplate x:Key="AuthorTemplate">
```

```

<Grid Background="LightGray"
      Width="300"
      Height="200">
  <Grid.RowDefinitions>
    <RowDefinition Height="75"/>
    <RowDefinition Height="Auto"/>
    <RowDefinition Height="Auto"/>
  </Grid.RowDefinitions>
  <StackPanel Orientation="Horizontal" Grid.Row="0">
    <TextBlock Text="{Binding LastName}"
              Margin="20,10,0,0"
              FontSize="24"
              FontWeight="SemiBold"/>
    <TextBlock Text="{Binding FirstName}"
              Margin="20,10,0,0"
              FontSize="24"/>
  </StackPanel>
  <StackPanel Orientation="Horizontal" Grid.Row="1">
    <TextBlock Text="Royalty: "
              Margin="20,0,0,0"
              FontSize="18"/>
    <TextBlock Text="{Binding Royalty}"
              Margin="20,0,0,0"
              FontSize="18"/>
  </StackPanel>
  <StackPanel Orientation="Horizontal" Grid.Row="2">
    <TextBlock Text="Under Contract:"
              Margin="20,0,0,0"
              FontSize="18"/>
    <TextBlock Text="{Binding UnderContract}"
              Margin="20,0,0,0"
              FontSize="18"/>
  </StackPanel>
</Grid>
</DataTemplate>
</ResourceDictionary>
</Page.Resources>

```

12. Add the `GridView` to the page to display the author info. Notice the `ItemTemplate` is set to the `AuthorTemplate` you defined in step 11.

```

<GridView x:Name="AuthorsGridView"
          Grid.Row="1"
          Margin="110,50,0,0"
          SelectionMode="Single"
          IsSwipeEnabled="True"
          IsItemClickEnabled="True"
          ItemsSource="{Binding}"
          ItemTemplate="{StaticResource AuthorTemplate}">
  <GridView.ItemsPanel>

```

```

        <ItemsPanelTemplate>
            <WrapGrid Orientation="Horizontal" />
        </ItemsPanelTemplate>
    </GridView.ItemsPanel>
</GridView>

```

13. In the MainPage.xaml.cs code file, add a class level variable of Type List of Author.

```

public sealed partial class MainPage : Page
{
    List<Author> authors;

```

14. Add the following method to load the list.

```

private void PopulateAuthors()
{
    authors = new List272103_1_En();
    authors.Add(new Author("Clive", "Cussler", true, .15));
    authors.Add(new Author("Steve", "Berry", false, .20));
    authors.Add(new Author("Kate", "Morton", false, .20));
    authors.Add(new Author("Karma", "Wilson", true, .18));
}

```

15. In the onNavigatedTo event handler method, call the PopulateAuthors method and set the authors list to the GridView's ItemSource.

```

protected override void OnNavigatedTo(NavigationEventArgs e)
{
    PopulateAuthors();
    this.AuthorsGridView.ItemsSource = authors;
}

```

16. To test the page, start the debugger with the Simulator as the target. You should see a page similar to Figure 13-10.

Authors

<p>Cussler Clive</p> <p>Royalty: 0.15 Under Contract: True</p>	<p>Berry Steve</p> <p>Royalty: 0.2 Under Contract: False</p>	<p>Morton Kate</p> <p>Royalty: 0.2 Under Contract: False</p>	<p>Wilson Karma</p> <p>Royalty: 0.18 Under Contract: True</p>
---	---	---	--

Figure 13-10. Author information in a GridView

17. After testing the page, press Ctrl +Alt + F4 keys to close the simulator and stop debugging.

Investigating Binding Modes

18. Add another row to the main Grid on the MainPage.

```
<Grid Background="{StaticResource ApplicationPageBackgroundThemeBrush}">
    <Grid.RowDefinitions>
        <RowDefinition Height="140"/>
        <RowDefinition Height="Auto"/>
        <RowDefinition Height="Auto"/>
    </Grid.RowDefinitions>
```

19. Below the GridView closing tag, add XAML to create a StackPanel with two buttons in it.

```
<StackPanel Orientation="Horizontal" Grid.Row="2" VerticalAlignment="Top">
    <Button Name="btnUpdate" Click="btnUpdate_Click" Content="Update" />
    <Button Name="btnRefresh" Click="btnRefresh_Click" Content="Refresh" />
</StackPanel>
```

20. The XAML defining the btnUpdate button above includes an event handler method for the Click event. Right-click on the Click event and select “Navigate to Event Handler” in the pop up menu.

21. Add the following code to the Click event handler to update the first author’s last name.

```
private void btnUpdate_Click(object sender, RoutedEventArgs e)
{
    authors[0].LastName = "Webb";
}
```

22. Add the following code to the btnRefresh_Click method. This code clears and rebinds the GridView.

```
private void btnRefresh_Click(object sender, RoutedEventArgs e)
{
    this.AuthorsGridView.ItemsSource = null;
    this.AuthorsGridView.ItemsSource = authors;
}
```

23. Test the page in the Simulator. When you click the Update button the first author’s name should update on the page. This is because the binding mode is set to OneWay by default.

24. Change the binding mode of the txtFirstName TextBlock to OneTime.

```
<TextBlock Text="{Binding FirstName, Mode=OneTime}"
           Margin="20,10,0,0" FontSize="24"/>
```

25. Test the page in the Simulator. When you click the Update button, the first author’s name should not update on the page. This is because the binding mode is set to OneTime, which means it only changes when it binds. Click the refresh button to rebind the DataGrid. You should now see the new value in the GridView.

26. When you are done testing, close the Simulator and exit Visual Studio.

Page Navigation

Page navigation in a Windows Store app is similar in concept to page navigation in a web app. The XAML UI framework provides a built-in `Frame` control that keeps track of the navigation history. You can use the methods of the `Frame` object, such as the `GoBack`, `GoForward`, and `Navigate`, to easily navigate through the navigation history or jump to a particular page. The `Navigate` method also provides a parameter that allows you to pass data between pages.

The following code in the `App.xaml.cs` file is used to navigate to the main page of the app when it starts up. If the `Frame` object cannot navigate to the `MainPage`, it throws an exception.

```
if (!rootFrame.Navigate(typeof(MainPage), args.Arguments))
{
    throw new Exception("Failed to create initial page");
}
```

The following code calls the `Navigate` method to navigate to a details page passing an object representing the author shown in a `GridView` item (`e.ClickedItem`).

```
private void AuthorsGridView_ItemClick(object sender, ItemClickEventArgs e)
{
    Frame.Navigate(typeof(AuthorDetailPage), e.ClickedItem);
}
```

The `AuthorDetailPage` uses the event handler `OnNavigatedTo` to retrieve the author object passed and bind it to `DataContext` of the `AuthorDetails` grid.

```
protected override void OnNavigatedTo(NavigationEventArgs e)
{
    author = (Author) e.Parameter;
    this.AuthorDetails.DataContext = author;
}
```

ACTIVITY 13-3. WORKING WITH DATA TEMPLATES

In this activity, you will become familiar with the following:

- implementing Page Navigation
- passing Data between Pages

Passing Data and Page Navigation

To implement page navigation, follow these steps:

1. Start Visual Studio. Select **File** ► **Open** ► **Project**.
2. Choose the `Activity13_2` project you completed in the last activity.
3. After the project loads, add a new blank page called `AuthorDetail.xaml`.

4. Add the following markup code between the `Grid` tags in the `AuthorDetailPage`'s XAML file. This creates the page's title and adds a back button to the page.

```
<Grid.RowDefinitions>
    <RowDefinition Height="140"/>
    <RowDefinition Height="Auto"/>
</Grid.RowDefinitions>
<Grid>
    <Grid.ColumnDefinitions>
        <ColumnDefinition Width="Auto"/>
        <ColumnDefinition Width="*/>
    </Grid.ColumnDefinitions>
    <Button Name="btnBack" Style="{StaticResource BackButtonStyle}"
        Click="btnBack_Click"></Button>
    <TextBlock x:Name="PageTitle"
        Text="Author Info" Grid.Column="1" VerticalAlignment="Center"
        Margin="120,20,0,0" IsHitTestVisible="false"
        Style="{StaticResource PageHeaderTextStyle}"/>
</Grid>
```

5. Add this mark up after the `TextBlock` tag and before the ending `Grid` tag shown in step 4. This creates the display controls to display the author's details.

```
<Grid Grid.Row="1" Name="AuthorDetails" DataContext="{Binding Mode=TwoWay}"
    HorizontalAlignment="Left" Margin="20,0,0,0">
    <Grid.ColumnDefinitions>
        <ColumnDefinition Width="Auto" />
        <ColumnDefinition Width="Auto" />
    </Grid.ColumnDefinitions>
    <Grid.RowDefinitions>
        <RowDefinition Height="Auto" />
        <RowDefinition Height="Auto" />
        <RowDefinition Height="Auto" />
        <RowDefinition Height="Auto" />
        <RowDefinition Height="Auto" />
    </Grid.RowDefinitions>

    <TextBlock Text="First Name:" Grid.Column="0"
        Grid.Row="0" HorizontalAlignment="Left"
        Margin="3" VerticalAlignment="Center" FontSize="18" />
    <TextBox Grid.Column="1" Grid.Row="0" Height="23" HorizontalAlignment="Left"
        Margin="3" Name="txtFirstName" Text="{Binding Path=FirstName}"
        VerticalAlignment="Center" Width="120" FontSize="18" />
    <TextBlock Text="Last Name:" Grid.Column="0" Grid.Row="1"
        HorizontalAlignment="Left" Margin="3" VerticalAlignment="Center"
        FontSize="18" />
    <TextBox Grid.Column="1" Grid.Row="1" Height="23" HorizontalAlignment="Left"
        Margin="3" Name="txtLastName" Text="{Binding Path=LastName}"
        IsEnabled="False"
        VerticalAlignment="Center" Width="120" FontSize="18" />
```

```

<TextBlock Text="Royalty:" Grid.Column="0" Grid.Row="2"
    HorizontalAlignment="Left" Margin="3" VerticalAlignment="Center"
    FontSize="18" />
<TextBox Grid.Column="1" Grid.Row="2" Height="23" HorizontalAlignment="Left"
    Margin="3" Name="txtRoyalty" Text="{Binding Path=Royalty}"
    VerticalAlignment="Center" Width="120" FontSize="18" />
<CheckBox Name="chkContract" Content="Under Contract"
    IsChecked="{Binding Path=UnderContract}"
    Grid.Row="3" Grid.Column="0" FlowDirection="RightToLeft" FontSize="18" />
</Grid>

```

6. In the `AuthorDetailPage.xaml.cs` file, add the `btnBack_Click` event handler. This sends the navigation back to the calling page.

```

private void btnBack_Click(object sender, RoutedEventArgs e)
{
    Frame.GoBack();
}

```

7. In the `MainPage.xaml` file, find the `AuthorsGridView` opening tag and add an `ItemClick` attribute.

```

<GridView x:Name="AuthorsGridView"
    Grid.Row="1"
    Margin="110,50,0,0"
    SelectionMode="Single"
    IsSwipeEnabled="True"
    IsItemClickEnabled="True"
    ItemsSource="{Binding Mode=OneTime}"
    ItemTemplate="{StaticResource AuthorTemplate}"
    ItemClick="AuthorsGridView_ItemClick">

```

8. In the `MainPage.xaml.cs` file, add the `AuthorsGridView_ItemClick` event handler. This causes the application to navigate to the `AuthorDetailPage`.

```

private void AuthorsGridView_ItemClick(object sender, ItemClickEventArgs e)
{
    Frame.Navigate(typeof(AuthorDetailPage));
}

```

9. Test the navigation in the Simulator. When the `MainPage` displays, click on one of the gray boxes with the authors' information displayed. The `AuthorDetailPage` should display. Click on the back button, which will take you back to the `MainPage`. After testing, close the Simulator.

Passing Data between Pages

1. In the `MainPage.xaml.cs` file, update the `AuthorsGridView_ItemClick` event handler to pass the `e.ClickedItem` which represents an `Author` object to the `AuthorDetailPage`.

```
private void AuthorsGridView_ItemClick(object sender, ItemClickEventArgs e)
{
    Frame.Navigate(typeof(AuthorDetailPage), e.ClickedItem);
}
```

2. In the `AuthorDetailPage.cs` file, add the following code to the `OnNavigatedTo` event handler. This code takes the parameter passed to the page and assigns it to the `DataContext` of the `AuthorDetails` grid.

```
protected override void OnNavigatedTo(NavigationEventArgs e)
{
    Author author = (Author) e.Parameter;
    this.AuthorDetails.DataContext = author;
}
```

3. Test the application in the Simulator. When the `MainPage` displays, click on one of the gray boxes with the authors information displayed. The `AuthorDetailsPage` should display with the author information loaded into the controls on the page. After testing, close the Simulator.
-

Summary

In this chapter, you took a look at implementing the interface tier of an application using the new Windows Store App. You saw how to use XAML to create pages and layout controls. You also saw how easy it is to bind the controls to data and present them to the users. What's still missing from the story is information on how to retrieve data from a relational database on a server. In order to provide server-side data to a Windows Store application, you need to utilize a web service. In Chapter 14, you will look at creating web services using the Windows Communication Framework (WCF) and the new ASP.NET Web API. You will also look at consuming web services in a client application.



Developing and Consuming Web Services

In the previous three chapters, you examined the steps required to create the graphical user interface of an application. Graphical user interfaces created with WPF, ASP.NET, and Windows Store apps provide the user a way to interact with your applications and employ the services the application provides. This chapter shows you how to build another type of interface, one that is implemented using web service protocols and is meant to be consumed by an application. Such a service provides an application with a programmatic interface with which to access its functions, without the need for human interaction. You will first look at creating and consuming Windows Communication Foundation (WCF) web services. WCF services are robust services that support secure messaging over multiple protocols. WCF services are an excellent choice when you are responsible for both the service provider and the service consumer, for example providing services for consumption on a corporate intranet. You will also look at creating and consuming ASP.NET Web API services. The ASP.NET Web API provides a lighter weight protocol used to pass data between the client and server over HTTP. It is a better choice for passing data across the internet between client and servers created using different technologies and/or different groups.

After reading this chapter, you will have a clearer understanding of the following:

- what web services are and how they came about
- how WCF processes service requests
- how to create and consume a WCF service
- how ASP.NET Web API services process service requests
- how to create and consume an ASP.NET Web API service

What Are Services?

Microsoft first introduced the concept of services with its inclusion of web services support in .NET Framework 1.0. A web service provides a way for an application to request a service and receive a reply. This is essentially the same as a client object requesting a service (method) from a server object within the boundaries of your application. The difference is the location of the client objects and server objects. If they reside in the same application, then they can issue and receive binary messages and inherently understand each other because they are speaking the same “language.” As the applications you build grow more complex, it is common to split the application up into distinct components. When you segment an application into components, each designed to perform a distinct specialized service, you greatly enhance code maintenance, reusability, and reliability. Additionally, separate servers can host the various components for increased performance, better maintenance, and security.

Prior to the introduction of web services, the clients and servers of an application relied on distributed technologies such as Distributed Component Object Model (DCOM) from Microsoft and the Common Object Request Broker Architecture (CORBA) from the Object Management Group (OMG) to communicate. These early attempts at service-oriented architecture worked fairly well if the client and server applications utilized the same technology, but when the client and server utilized disparate technologies, it became very problematic. The power of web services lies in the fact that they use a set of open messaging and transport protocols. This means that client and server components utilizing different technologies can communicate in a standard way. For example, a Java-based application running on an Apache web server can request a service from a .NET-based application running on an IIS server. In addition, they can be located virtually anywhere in the world that has an Internet connection.

WCF Web Services

With the release of the .NET Framework 3.0, Microsoft introduced a new way to create web services in the form of Windows Communication Foundation services (WCF). Before WCF, Microsoft had a robust but confusing set of messaging technologies including ASP.NET Web Services, MSMQ, Enterprise services, and .NET Remoting. Microsoft decided to roll all these technologies into a single framework for developing service-oriented applications. This made developing service-oriented applications more consistent and less confusing for developers.

A WCF service is made up of three parts: the service, an end point, and a hosting environment. The service is a class that contains methods you want to expose to clients of the service. An end point is a definition of how clients can communicate with the service. It's worth noting that a service can have more than one endpoint defined. An endpoint consists of the base *address* of the service, its *binding* information, and its *contract* information (the three are often referred to as the ABCs of WCF). The hosting environment refers to the application hosting the service. For your purposes, this will be a web server, but there are other options that exist depending on the type of WCF service you implement.

Creating a WCF Web Service

Creating and consuming basic WCF services using Visual Studio 2012 is a fairly easy process. If you use the templates Visual Studio provides, much of the plumbing work is done for you. Figure 14-1 shows the available templates. To create a WCF web service, you use the WCF Service Application template.

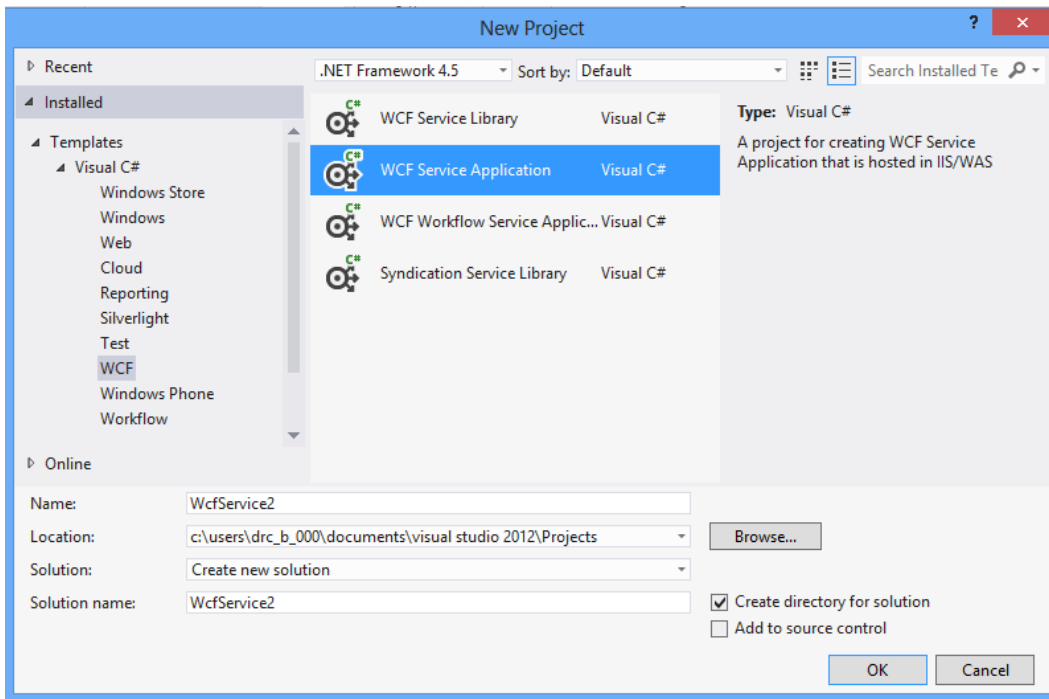


Figure 14-1. WCF templates provided by Visual Studio

Selecting a template adds two important files to the project: one defines the service contract using an interface and one is a class file that contains the service implementation code. In Figure 14-2, the `IService1.cs` file defines the interface and the `Service1.svc.cs` contains the class implementation for the service.




Figure 14-2. WCF interface and class files

When you create a service, you need to define the service contract. The contract is defined by an interface definition. The interface defines the methods exposed by the service, any input parameters expected by the methods, and any output parameters passed back by the methods. The following code shows the interface code for a tax service. The interface is marked with the `[ServiceContract]` attribute and any exposed methods are marked with the `[OperationContract]`.

```
[ServiceContract]
public interface ITax
{
    [OperationContract]
    double GetSalesTax(string statecode);
}
```

Once the interface is defined, the next step is to define the class that implements the interface. The following code implements the `ITax` interface and provides the code to implement its exposed methods.

```
public class Tax : ITax
{
    public double GetSalesTax(string stateCode)
    {
        if (stateCode == "PA")
        {
            return .06;
        }
        else
        {
            return .05;
        }
    }
}
```

Once the interface and class are defined, compiling and running the application produces the test client shown in Figure 14-3. You can test the web service by entering a state code and clicking the `Invoke` button. Clicking on the `XML` tab shows the request and response messages sent between the client and server as shown in Figure 14-4.

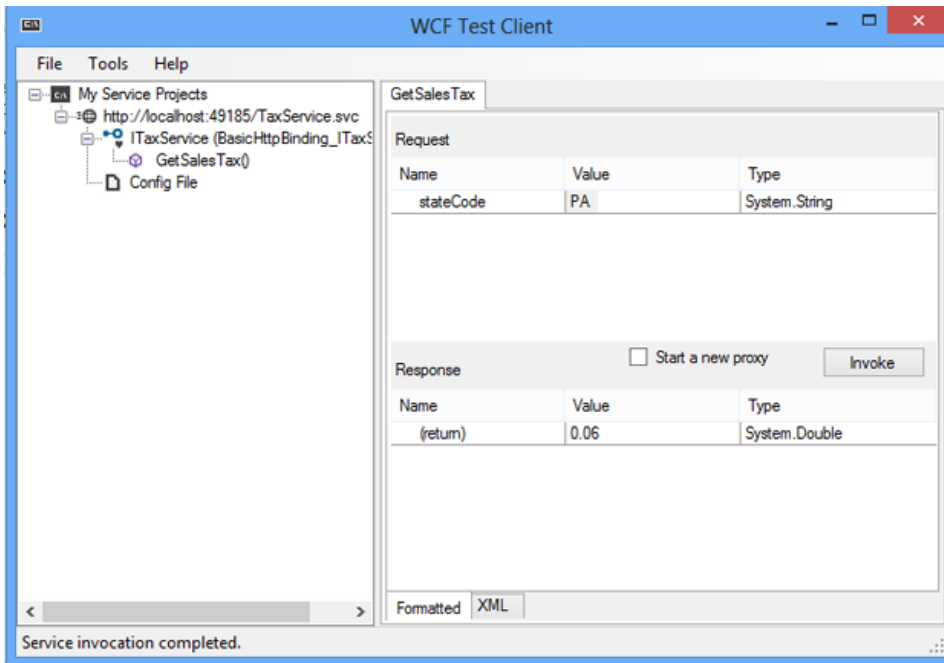


Figure 14-3. Testing the web service

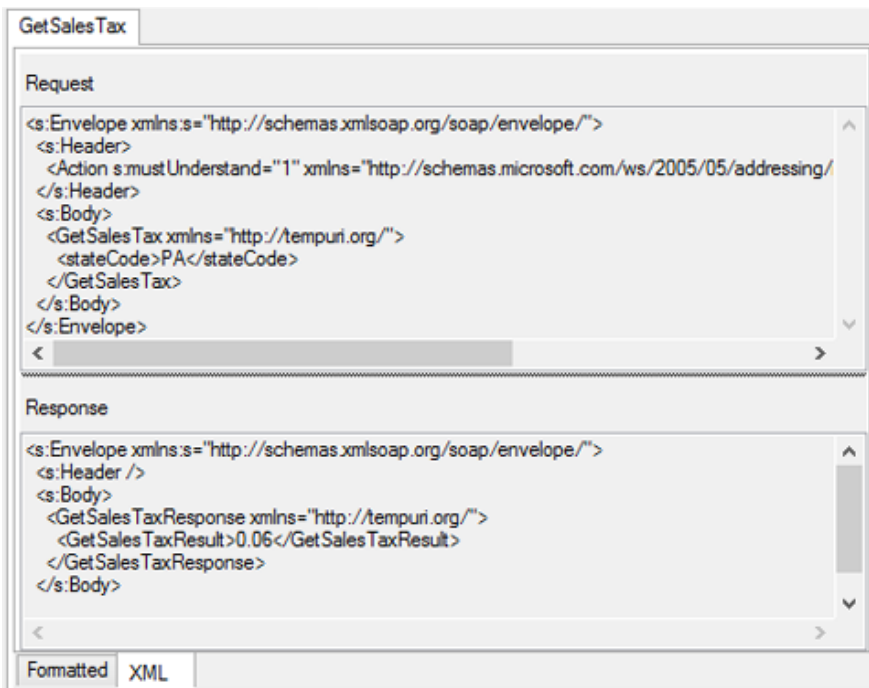


Figure 14-4. The Request and Response messages

Consuming a WCF Web Service

To consume a WCF service in a .NET client, you must add a service reference to the project. When you add a service reference in Visual Studio 2012, you are presented with an Add Reference window (see Figure 14-5). This window allows you to discover the services available and the operations they expose. You can also change the namespace that you use to program against the service.

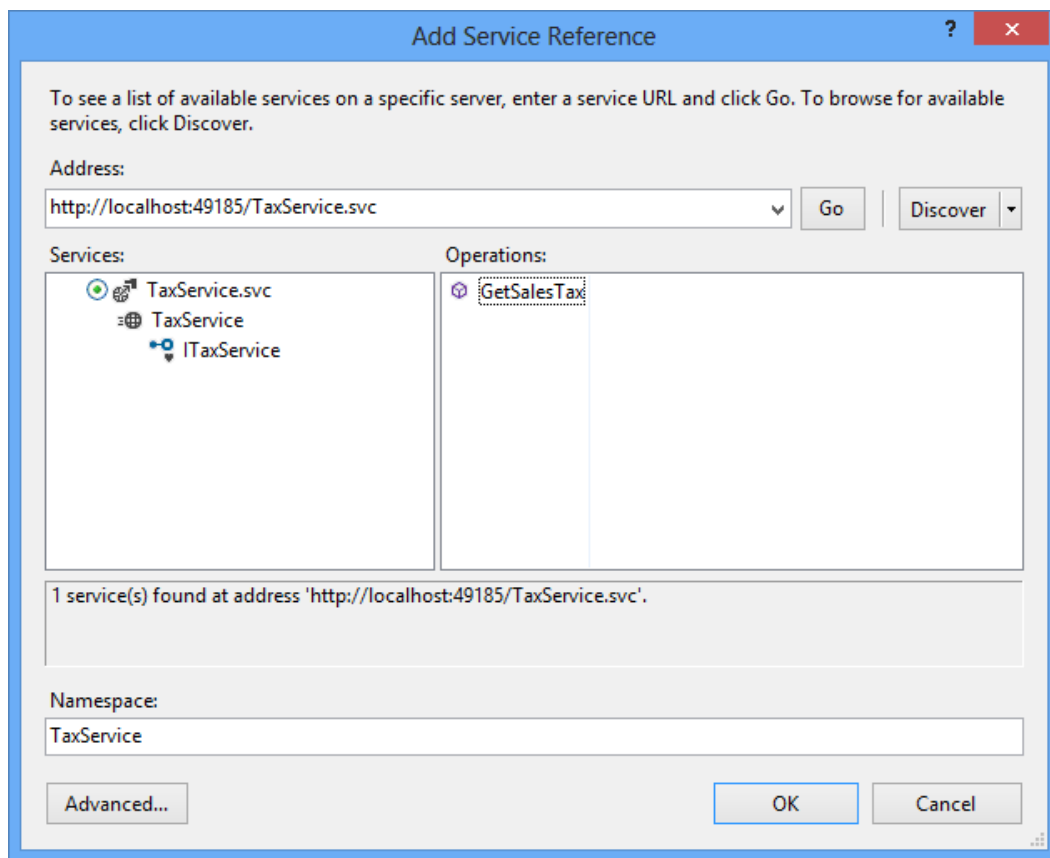


Figure 14-5. Adding a service reference

Once the service reference is added to the project, Visual Studio updates the application configuration file with the information needed to call the service. This includes the endpoint configuration with the address, binding, and contract information.

```
<endpoint address="http://localhost:49185/TaxService.svc" binding="basicHttpBinding"
  bindingConfiguration="BasicHttpBinding_ITaxService" contract="TaxService.
  ITaxService"
  name="BasicHttpBinding_ITaxService" />
```

A client proxy is also added to the client application. The client application uses this proxy to interact with the service. The following code shows a client console application calling the service using the `TaxServiceClient` proxy and writing the results out to the console window. Figure 14-6 shows the output in console window.

```
TaxService.TaxServiceClient webService = new TaxService.TaxServiceClient();
string state1 = "PA";
double salesTax1 = webService.GetSalesTax(state1);
Console.WriteLine("The sales tax for {0} is {1}", state1, salesTax1);
string state2 = "NJ";
double salesTax2 = webService.GetSalesTax(state2);
Console.WriteLine("The sales tax for {0} is {1}", state2, salesTax2);
webService.Close();
Console.ReadLine();
```

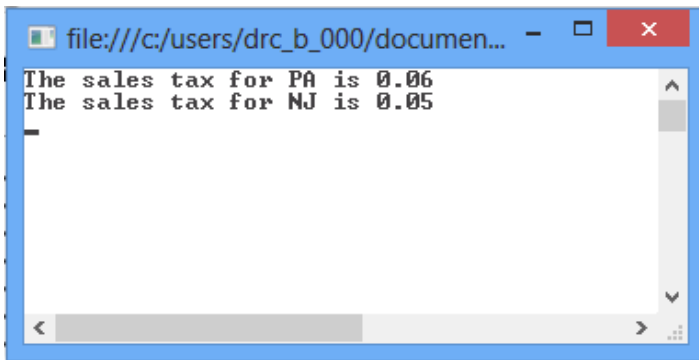


Figure 14-6. Output from calling the `TaxService`

Using Data Contracts

In the previous example, the WCF web service used only simple types to pass data back and forth between the service and the client. Simple types such as integer, double, and string do not require any special encoding to pass them between the client and server. There are times when you want to pass complex types between the client and server. Complex types are comprised of simple types. For example, you may have a service that takes an address type made up of street, city, state, and zip code and returns a location type made up of longitude and latitude. To facilitate the exchange of complex types, the WCF service uses data contracts. You create your data class as usual, then mark it with the `[DataContract]` attribute. The properties of the class that you want exposed are marked with the `[DataMember]` attribute. The following code exposes the `Location` class to clients of the service:


```
[DataContract]
public class Location
{
    double _longitude;
    double _latitude;
    [DataMember]
    public double Latitude
    {
        get { return _latitude; }
        set { _latitude = value; }
    }
    [DataMember]
    public double Longitude
    {
        get { return _longitude; }
        set { _longitude = value; }
    }
}
```

By marking the classes with the `[DataContract]` and `[DataMember]` attributes, an XML Schema Definition (XSD) file is created describing the complex types. Clients use this file to determine what to supply the service and what to expect as a return type. The following markup shows the portion of the XSD file created for the `Location` type returned by the service.

```
<xs:complexType name="Location">
  <xs:sequence>
    <xs:element minOccurs="0" name="Latitude" type="xs:double" />
    <xs:element minOccurs="0" name="Longitude" type="xs:double" />
  </xs:sequence>
</xs:complexType>
<xs:element name="Location" nillable="true" type="tns:Location" />
```

In the following activity, you will create a WCF service and consume the service in a .NET client application.

ACTIVITY 14-1. CREATING AND CONSUMING A WCF SERVICE

In this activity, you will become familiar with the following:

- creating a WCF Service
- consuming a WCF Service in a client application

Creating a WCF Service

To create the WCF Service, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose WCF under the Visual C# Templates folder and select the WCF Service Application project type. Rename the project to `Activity14_1` and click the OK button. (See Figure 14-7)




Figure 14-7. Creating a WCF Service Application

3. Delete the `IService1.cs` and the `Service1.svc` files. Right-click on the project node in Solution Explorer and select **Add** ► **New Item**. In the Add New Item dialog window select the WCF Service. Name the service `DiscountService`.
4. Open the `IDiscountService.cs` file in the code editor. Replace the `DoWork` operation contract with a `GetDiscount` operation contact.

```
[ServiceContract]
public interface IDiscountService
{
    [OperationContract]
    double GetDiscount(string discountCode);
}
```

5. Open the `DiscountService.svc.cs` file in the code editor. Replace the `DoWork` method with the following `GetDiscount` method.

```
public class DiscountService : IDiscountService
{
    public double GetDiscount(string discountCode)
    {
        if (discountCode == "XXXX")
        {
            return 20;
        }
    }
}
```

```

        else
        {
            return 10;
        }
    }
}

```

6. Build the project. If there are any errors, fix them, and then rebuild.
7. Make sure the DiscountService.svc node is selected in solution explorer and launch the debugger. You should be presented with the WCF Test Client. Test the GetDiscount operation. When done testing, close the WCF Test Client.

Creating the .NET Client

8. In the Solution Explorer, right-click the Activity14_1 solution node and select Add ► New Project. Under the Windows templates add a WPF Application named OfficeSupply.
9. Add the following XAML to the MainWindow between the Grid tags.

```

<TextBlock Text="Discount Code:" FontSize="20" Margin="10,25,350,250"/>
<TextBox Name="txtDiscountCode" FontSize="20" Margin="160,25,100,250"/>
<Button Name="btnGetDiscount" Content="Get Discount" Width="150"
        FontSize="20" Margin="150,75,213,200" />

```

10. In the Solution Explorer, right-click the OfficeSupply project node and select Add Service Reference. In the Add Service Reference dialog, click the Discover button. You should see the DiscountService.svc as shown in Figure 14-8. Click the OK button to add the reference.




Figure 14-8. Adding the service reference

■ **Note** The port number of your service address may change when you develop it locally.

11. Add a Click event handler method to the btnGetDiscount.
12. Add the following code to the click event handler. It is usually good practice to call web services asynchronously so the client application remains responsive while waiting for the response. Remember from Chapter 8 you need to use the `await` key word when calling an async method and also add the `async` key word to the method you are calling it from.

```
private async void btnGetDiscount_Click(object sender, RoutedEventArgs e)
{
    ServiceReference1.DiscountServiceClient service =
        new ServiceReference1.DiscountServiceClient();
    double discount = await service.GetDiscountAsync(txtDiscountCode.Text);
    MessageBox.Show(discount.ToString());
}
```

13. In the Solution Explorer, right-click the Activity14_1 solution node and select Properties. Click on the Startup Project node and select Multiple startup projects. Make sure the Action of each project is set to Start and click OK to continue. (See Figure 14-9)




Figure 14-9. Setting the startup projects

14. Start the debugger. When the form shows enter a discount code of “XXXX” and click the Get Discount button. You should see a message box showing 20 (see Figure 14-10). Try a different code; you should see a discount of 10. When finished testing, stop the debugger and exit Visual Studio.

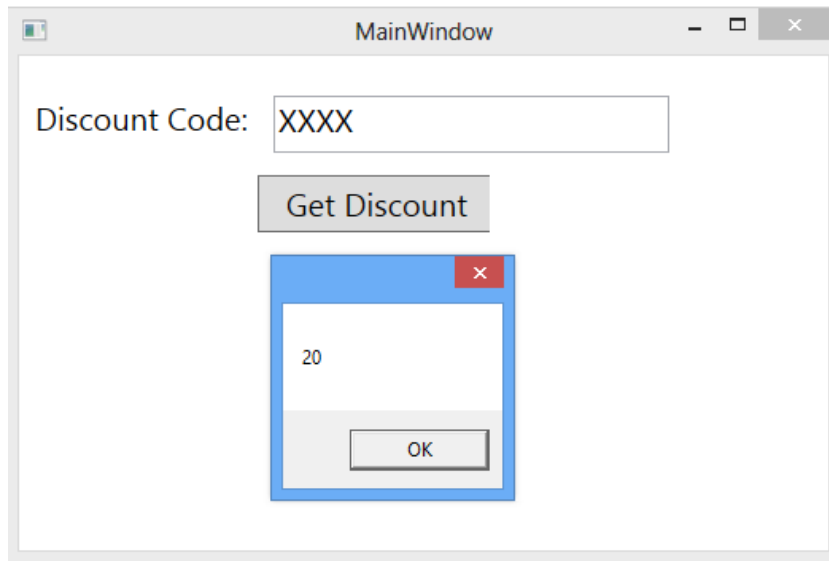


Figure 14-10. Testing the discount web service

RESTful Data Services

Even though WCF is a great technology for building web services, it contains a lot of overhead when passing messages between clients and servers. This is especially apparent when the client and server are passing data back and forth. Not only is the data passed, but also all the meta-data describing the data. Often the actual data being transferred comprises very little of the message being sent. Often all this meta-data is not needed and in fact strongly typing the data can cause problems when the client and server are built using different programming frameworks with different type systems. This is where Representational State Transfer- (REST) based services come in handy.

RESTful web services use simple communication over HTTP and deliver the data using plain old XML (POX) or JavaScript Object Notation (JSON). Data is accessed and changed by using the standard HTTP verbs GET, PUT, POST, and DELETE. Microsoft developed the ASP.NET WEB API as a RESTful alternative to WCF. It doesn't replace WCF, but provides a nice lighter weight alternative to WCF when you do not need to support transport protocols like TCP and UDP, or support binary encoding. It is very useful for passing data from client to server in a loosely coupled way. So while WCF is useful for exposing application logic (methods) as a service, ASP.NET Web API services are great for exposing data.

Creating an ASP.NET Web API Service

When creating an ASP.NET Web API service, you need to create a web application to host the service. Once the web application is created, you add a Web API Controller Class. The Web API Controller class provides the functionality necessary to process request messages, interact with the data model, and generate response messages. For example, when you want to request a list of employees, you would request the following URI.

<http://localhost/DemoChapter14/employees>

Where localhost is the name of the web server, DemoChapter14 is the name of the web application and employees is the name of the controller. This URI gets mapped to the get method of the controller that returns an array of strings containing the employee data. The output can be in the form of POX or JSON depending on the request.

If you want to get a single employee's data you would use the same URI with the employee key tacked on the end of it.

<http://localhost/DemoChapter14/employees/10>

This is mapped to a Get method on the controller that has an input parameter of type integer.

```
public string Get(int id)
{
    //Retrieve employee data
}
```

To map the URI requests to the appropriate method of the Web API Controller class you need to add the following code to the Application_Start event handler in the Global.asax.cs file. This code adds the mapping information to the RouteTable.

```
RouteTable.Routes.MapHttpRoute(
    name: "DefaultApi",
    routeTemplate: "DemoChapter14/{controller}/{id}",
    defaults: new
    {
        id = RouteParameter.Optional
    }
);
```

In the following activity, you'll create a Web API Service that supplies book data from the Pubs database.

ACTIVITY 14-2. CREATING A WEB API SERVICE

In this activity, you will become familiar with the following:

- creating a Web API Service

To create a Web API Service, follow these steps:

1. Start Visual Studio. Select File ► New ► Project.
2. Choose Web under the Visual C# Templates folder and select the ASP.NET Empty Web Application project type. Rename the project to Activity14_2 and click the OK button. (See Figure 14-11)

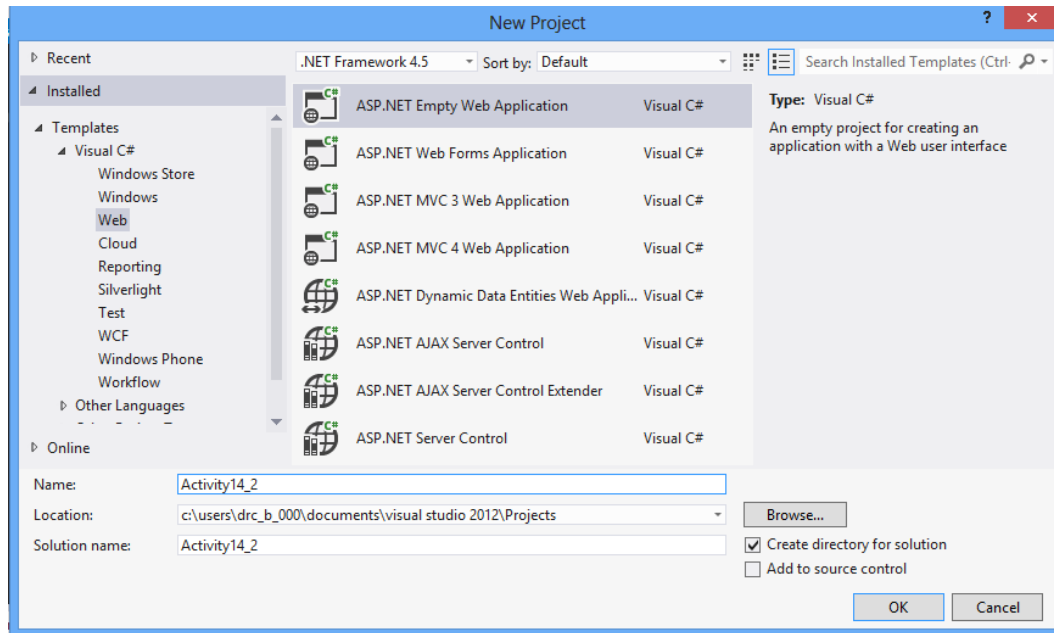


Figure 14-11. Adding an empty Web Application project

3. Right-click on the Activity14_2.Web project node in the solution explorer window and select Add ► New Folder. Name the folder Model.
4. Right-click on the Model folder in the solution explorer window and select Add ► New Item. Under the Data node in the Add New Item window, select an ADO.NET Entity Data Model. Name the model Pubs.edmx and click Add.
5. In the Choose Model Contents screen, select the Generate from database option and click Next.
6. In the Choose Your Data Connection screen, choose an existing connection or create a new connection to the Pubs database and choose Next.
7. In the Choose Your Database Objects screen, expand the tables node; select the titles table; and then click Finish.
8. Right-click on the Activity14_2 web project node in the solution explorer window and select Add ► New Folder. Name the folder Controllers.
9. Right-click on the Controllers folder in the Solution Explorer window and select Add ► New Item. In the Add New Item window, click on the web node in the Installed Templates. Select the Web API Controller Class template, rename it to BooksController, and click the Add button.
10. Open the BooksController.cs file in the Code Editor. Add the following using statement to the top of the file.

```
using Activity14_2.Model;
```


11. Update the code so that the Get method uses LINQ to EF (Chapter 10) to retrieve an array of books and sends it back to the caller.

```
// GET api/<controller>
public IEnumerable<title> Get()
{
    var context = new pubsEntities();
    var query = from t in context.titles select t;
    var titles = query.ToArray<title>();
    return titles;
}
```

12. Right-click on the Activity14_2 web project node in the Solution Explorer window and select Add ► New Item. In the Add New Item window, click on the web node in the Installed Templates. Select the Global Application Class template. Keep it named Global.asax.

13. In the Global.asax.cs file add the following using statements.

```
using System.Web.Routing;
using System.Web.Http;
```

14. In the Application_Start event handler add the following code to map the URI verbs (Get, Put, etc.) to the BooksController methods.

```
RouteTable.Routes.MapHttpRoute(
    name: "DefaultApi",
    routeTemplate: "Activity14_2/{controller}/{id}",
    defaults: new
    {
        id= RouteParameter.Optional
    }
);
```

15. Build the solution. If there are any errors, fix them and rebuild. Figure 14-12 shows the project files in Solution Explorer.

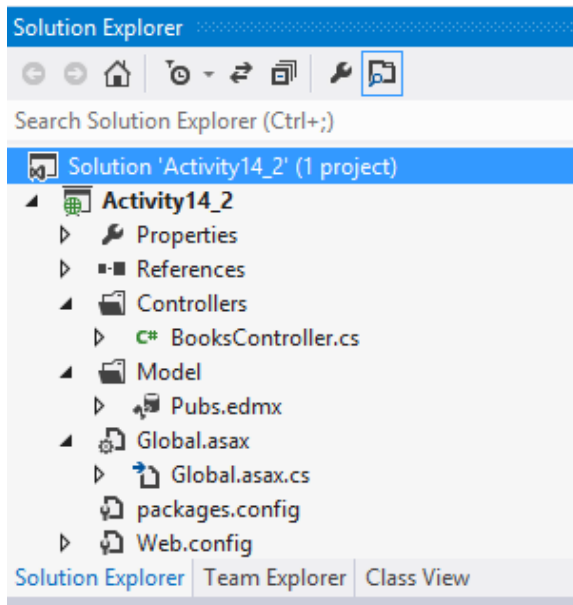


Figure 14-12. Viewing the project files

16. Launch the debugger. In the browser window, type the following URI (http://localhost:port/Activity14_2/books). Replace the port with the port number of your project. You can find the port number by right clicking the project node in Solution Explorer and selecting properties. The project URL should be listed on the Web tab (See Figure 14-13).

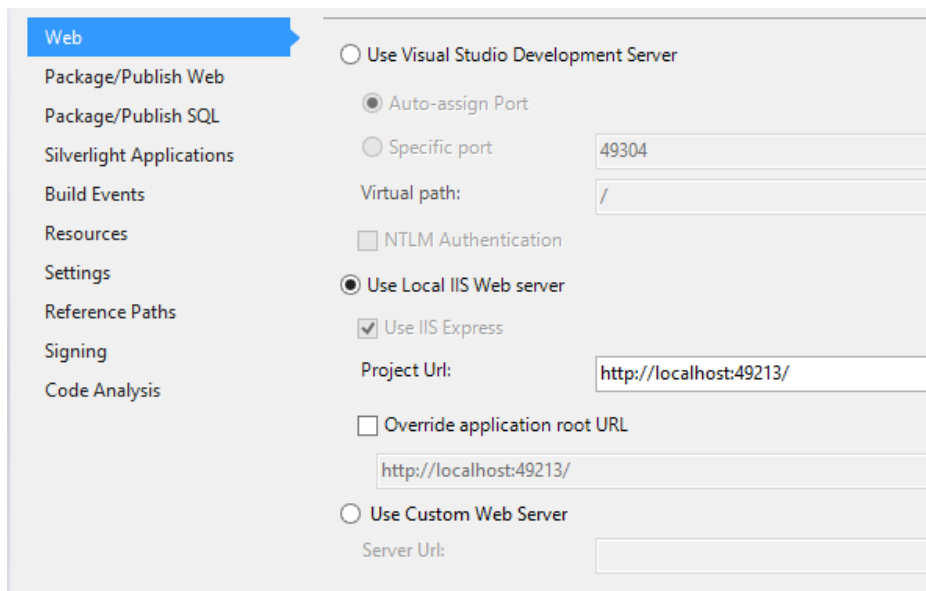


Figure 14-13. Looking up the port number

17. You should get a message asking if you want to open or save books.json. Open the file in Notepad. You should see the book information listed in JSON format (see Figure 14-14).

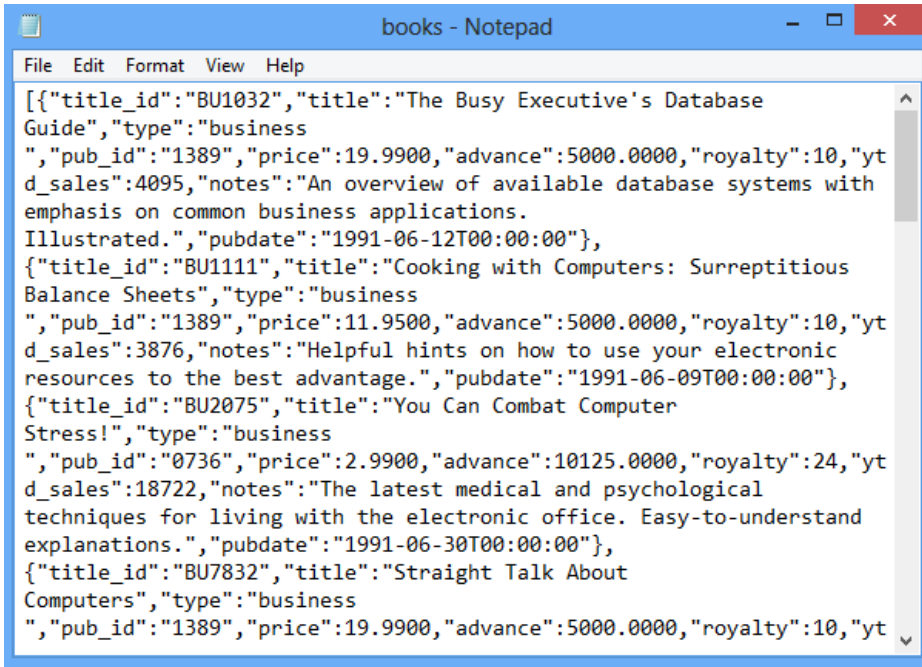


Figure 14-14. Viewing the books.json file in the Notepad

18. When finished testing, stop the debugger and exit Visual Studio.

Consuming ASP.NET Web API Services

When a client application calls an ASP.NET Web API service, it needs to make a HTTP request and wait for the response data in the form of a POX or JSON text stream. After receiving the response, the client parses the text file and loads the data in a local object for processing. Since the meta-data is not passed along with the data, Visual Studio cannot automatically generate the proxy objects for the client to work with.

The `HttpClient` class in the `System.Net.Http` namespace is used to send requests and receive responses from the Web API Service. The following code demonstrates using the `HttpClient` class to request a list of books from a web service. First an `HttpClient` object is instantiated and the base address is set.

```
private HttpClient httpClient;
httpClient = new HttpClient();
httpClient.BaseAddress = new Uri("http://localhost:49213/");
```

Once the base address is set, you set the format of the response you want returned. In this case a JSON response is expected.

```
httpClient.DefaultRequestHeaders.Accept.Add
    (new MediaTypeWithQualityHeaderValue("application/json"));
```

You can then make a request and wait for a response. Since this can take a while, you should call the web service asynchronously.

```
var response = await httpClient.GetAsync("Activity14_2/books");
```

When the response comes back, you can parse the text using the `JSONArray` class in the `Windows.Data.Json` namespace (currently shipping with Windows 8) and load it into a local object/collection to work with.

```
var bookJSON = await response.Content.ReadAsStringAsync();
var bookArray = JSONArray.Parse(bookJSON);
ObservableCollection<Book> bc = new ObservableCollection<Book>();
foreach (var b in bookArray)
{
    Book bk = new Book();
    bk.title_id=b.GetObject()["title_id"].GetString();
    bk.title = b.GetObject()["title"].GetString();
    bk.type = b.GetObject()["type"].GetString();
    bc.Add(bk);
}
```

Consuming a Web API service from the various .NET clients (WPF, ASP.NET, and a Windows Store app) is a very similar process. In the following activity, you will call the web service you created in Activity 14-2 from a Windows Store app.

ACTIVITY 14-3. CALLING A WEB API SERVICE

In this activity, you will become familiar with the following:

- calling a Web API Service

To call a Web API Service, follow these steps:

1. Start Visual Studio and open the project from Activity14-2. Add a new Windows Store Blank App named `BookServiceClient` to the solution.
2. Add a new class to the `BookServiceClient` named `Book` and add the following properties to the class.

```
class Book
{
    public string Title_id { get; set; }
    public string Title { get; set; }
    public string Type { get; set; }
    public string Notes { get; set; }
}
```

3. Open the `MainPage.xaml` file in the designer. Add the following XAML markup between the `Grid` tags to define a `GridView` to display a collection of book data.

```

<GridView Name="BooksView" ItemsSource="{Binding}" Margin="50,50,0,0">
  <GridView.ItemTemplate>
    <DataTemplate>
      <StackPanel Width="300" Height="300" Margin="10"
        Background="#FF161C8F">
        <TextBlock Margin="10,0" Text="{Binding Title_id}" />
        <TextBlock Margin="10,0" Text="{Binding Title}" />
        <TextBlock Margin="10,0" Text="{Binding Type}" />
        <TextBlock Margin="10,0" Text="{Binding Notes}"
          TextWrapping="Wrap"/>
      </StackPanel>
    </DataTemplate>
  </GridView.ItemTemplate>
</GridView>

```

4. Open the MainPage.xaml.cs file in the code editor window. Add the following using statements to the top of the file.

```

using System.Net.Http;
using System.Net.Http.Headers;
using Windows.Data.Json;
using System.Collections.ObjectModel;

```

5. Add a class level HttpClient variable to the class and update the class constructor with the following code. Replace the port number with the port number being used by your web service.

```

private HttpClient httpClient;
public MainPage()
{
    this.InitializeComponent();
    httpClient = new HttpClient();
    httpClient.BaseAddress = new Uri("http://localhost:49213/");
    httpClient.DefaultRequestHeaders.Accept.Add
        (new MediaTypeWithQualityHeaderValue("application/json"));
}

```

6. Add an asynchronous method that calls the web service, parses the result, and binds it to the GridView.

```

private async void GetBooks()
{
    var response = await httpClient.GetAsync("Activity14_2/books");
    if (response.IsSuccessStatusCode)
    {
        var bookJSON = await response.Content.ReadAsStringAsync();
        var bookArray = JsonArray.Parse(bookJSON);
        ObservableCollection<Book> bc = new ObservableCollection<Book>();
        foreach (var b in bookArray)

```

```

    {
        Book bk = new Book();
        bk.Title_id=b.GetObject()["title_id"].GetString();
        bk.Title = b.GetObject()["title1"].GetString();
        bk.Type = b.GetObject()["type"].GetString();
        bk.Notes = b.GetObject()["notes"].Stringify();
        bc.Add(bk);
    }
    this.BooksView.DataContext = bc;
}
}

```

7. In the `OnNavigatedTo` event handler call the `GetBooks` method.

```

protected override void OnNavigatedTo(NavigationEventArgs e)
{
    GetBooks();
}

```

8. Build the solution. If there are any errors, fix them and rebuild.
9. In the Solution Explorer, right-click the `Activity14_2` solution node and select `Properties`. Click on the `Startup Project` node and select `Multiple startup projects`. Make sure the `Action` of each project is set to `Start`. (See Figure 14-15).

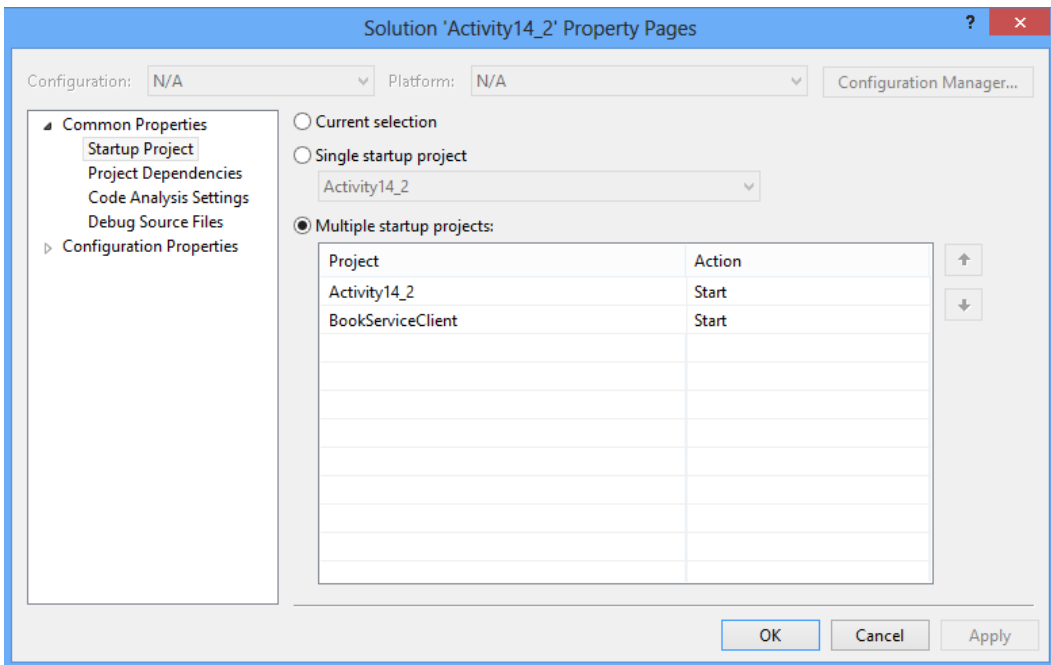


Figure 14-15. Setting the startup projects

10. Launch the debugger. You should see the book information listed in the GridView. (See Figure 14-16).

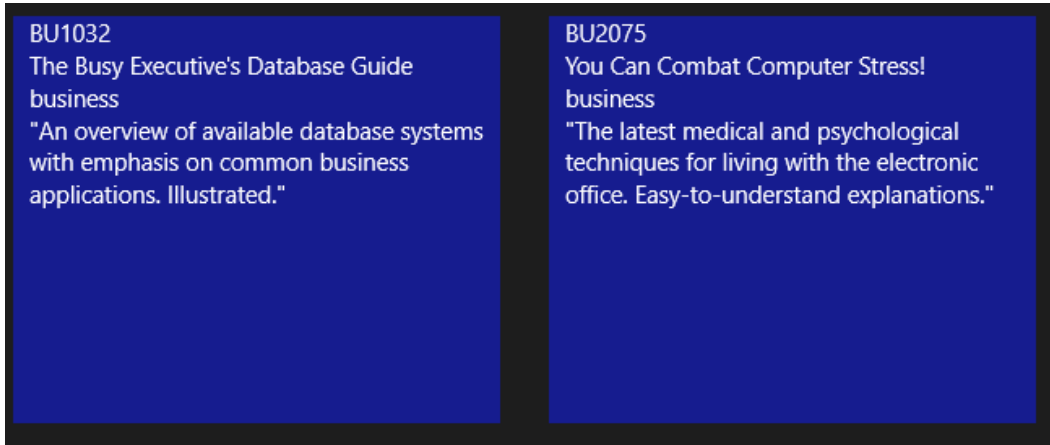


Figure 14-16. *Displaying book information*

11. When done testing, press the Alt + F4 keys to close the App. In Visual Studio stop the debugger and exit Visual Studio.
-

Summary

In this chapter, you were introduced to the fundamentals of implementing web services. In particular, you saw how to create web services using the Windows Communication Framework (WCF) and the ASP.NET Web API. You also built client applications that called these web services.

This was the final chapter in a series aimed at exposing you to the various technologies and .NET Framework classes used to build .NET applications. These chapters only scratched the surface of these technologies. As you gain experience developing .NET applications, you will need to look more deeply into each of these technologies.

Thus far in your journey you have studied UML design, object-oriented programming, the C# language, the .NET Framework, creating graphical user interfaces, and developing Web Services. You are now ready to put the pieces together and develop a working application. In Chapter 15, you will revisit the UML models you developed for the case study introduced in Chapter 4. You will transform these models into a functioning application.



Developing the Office Supply Ordering Application

In the previous chapters, you looked at three ways to develop the graphical user interface of an application. Graphical user interfaces provide human users with a way to interact with your applications and use the services they provide. You also saw how services create programmatic interfaces that other programs can call to use the services of the application without any user interaction.

In this chapter, you come full circle, back to the office supply ordering application (called OSO for short) that you designed in Chapter 4. This chapter is one big activity, and a final exam of sorts. You will create a functional application incorporating the concepts you have learned in the previous chapters. As you work through creating the application, you should be able to identify these concepts and relate them to the concepts covered previously. The application will contain a data access layer, a business logic layer, and a user interface layer.

After reading this chapter, you will understand why applications are split into different layers and how to construct them.

Revisiting Application Design

When you design an application, you can typically proceed in three distinct phases. First, you complete a conceptual design, then a logical design, and then a physical design.

The *conceptual design*, as explained in Chapter 4, constitutes the discovery phase of the process. The conceptual design phase involves a considerable amount of collaboration and communication between the users of the system and the system designers. The system designers must gain a complete understanding of the business processes that the proposed system will encompass. Using scenarios and use cases, the designers define the functional requirements of the system. A common understanding and agreement on system functionality and scope among the developers and users of the system is the required outcome of this phase.

The second phase of the design process is the *logical design*. During the logical design phase, you work out the details about the structure and organization of the system. This phase consists of the development and identification of the business objects and classes that will compose the system. UML class diagrams identify the system objects for which you identify and document the attributes and behaviors. You also develop and document the structural interdependencies of these objects using the class diagrams. Using sequence and collaboration diagrams, you discover and identify the interactions and behavioral dependencies between the various system objects. The outcome of this phase, the application object model, is independent of any implementation-specific technology and deployment architecture.

The third phase of the design process is the *physical design*. During the physical design phase, you transform the application object model into an actual system. You evaluate and decide upon specific technologies and infrastructures, do cost analysis, and determine any constraints. Issues such as programmer experience and knowledge base, current implementation technologies, and legacy system integration will all influence your decisions

during the physical design phase. You must also analyze security concerns, network infrastructure, and scalability requirements.

When designing a distributed application, you normally separate its logical architectural structure from its physical architectural structure. By separating the architectural structure in this way, you will find it much easier to maintain and update the application. You can make any physical architectural changes (to increase scalability, for example) with minimal impact. The logical architectural design typically separates the various parts of an application into tiers. Users interact with the *presentation tier*, which presents data to the user and gives the user ways to initiate business service requests. The *business logic tier* encapsulates and implements the business logic of an application. It is responsible for performing calculations, processing data, and controlling application logic and sequencing. The *data tier* is responsible for managing access to and storage of information that must be persisted and shared among various users and business processes. Figure 15-1 shows the different logical tiers of a typical 3-tier application.




Figure 15-1. Logical tiers of a 3-tiered application

When you create the physical tiers of an application, each logical tier would ideally correspond to a distinct physical tier on its own dedicated server. In reality, the physical layers of the application are influenced by such factors as available hardware and network infrastructure. You may have all the logical tiers on one physical server or spread across a web and database server. What is important is that you create applications that implement clear separation of duties among the classes.

Figure 15-2 shows the proposed layout of the OSO application. For simplicity's sake, you will create the business logic classes and the data access classes in the same assembly. This assembly is referred to as the business logic layer (BLL), while the user interface layer is contained in its own assembly user interface layer (UI). Both assemblies are contained on the same server (your computer). The important thing to remember is that because there is a clear separation of duties between the business logic classes and the data access classes, as the application grows in features and users, it can easily be refactored into separate assemblies hosted on separate servers.

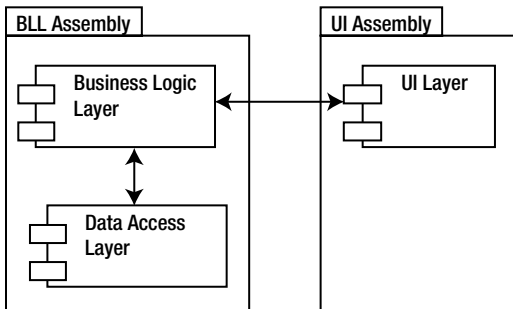


Figure 15-2. Physical tiers of the OSO application

Building the OSO Application's Data Access Layer

In order to develop the business logic and data access layers of the application, you need to review the OSO class diagram you created in Chapter 4 (shown in Figure 15-3).

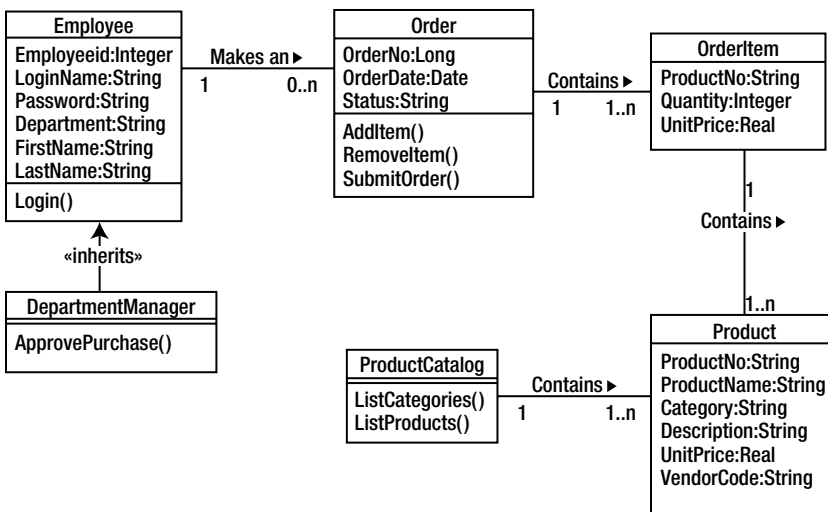


Figure 15-3. OSO application class diagram

As discussed in Chapter 4, you need to create an `Employee` class that implements a login method (`Login()`). The login method will interact with the database to verify login information. To accomplish this, you will create two employee classes: one for the business logic layer (`Employee`) and one for the data access layer (`DALEmployee`). The `Employee` class will pass the request to login from the User Interface (UI) to the `DALEmployee` class, which in turn will interact with the database to retrieve the requested information. Figure 15-4 is the database schema for the Office Supply database. This database is hosted in a SQL Server database.

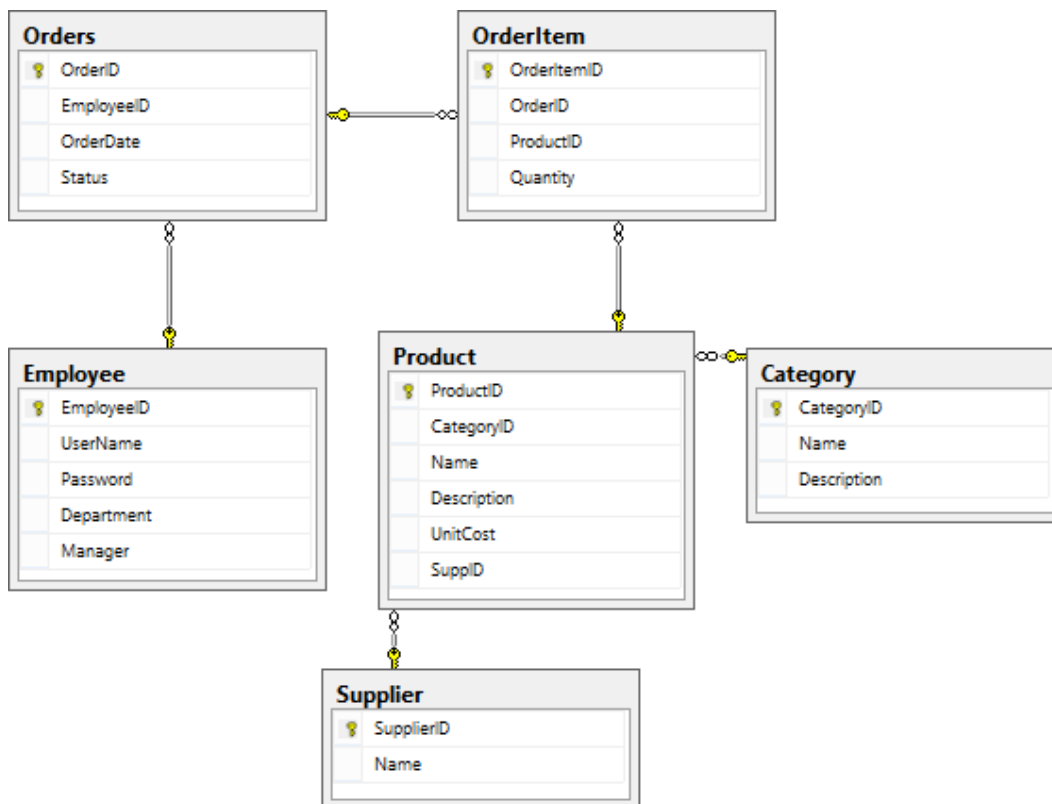


Figure 15-4. Office Supply database diagram

■ **Note** If you did not install the Office Supply database, see Appendix C for instructions.

First, you'll begin with the data access layer and then implement the business logic layer.

In Visual Studio, create a Class Library application and name it OfficeSupplyBLL and name the solution OfficeSupply as in Figure 15-5; this application will contain the classes for the business logic layer and data access layer of the OSO application.

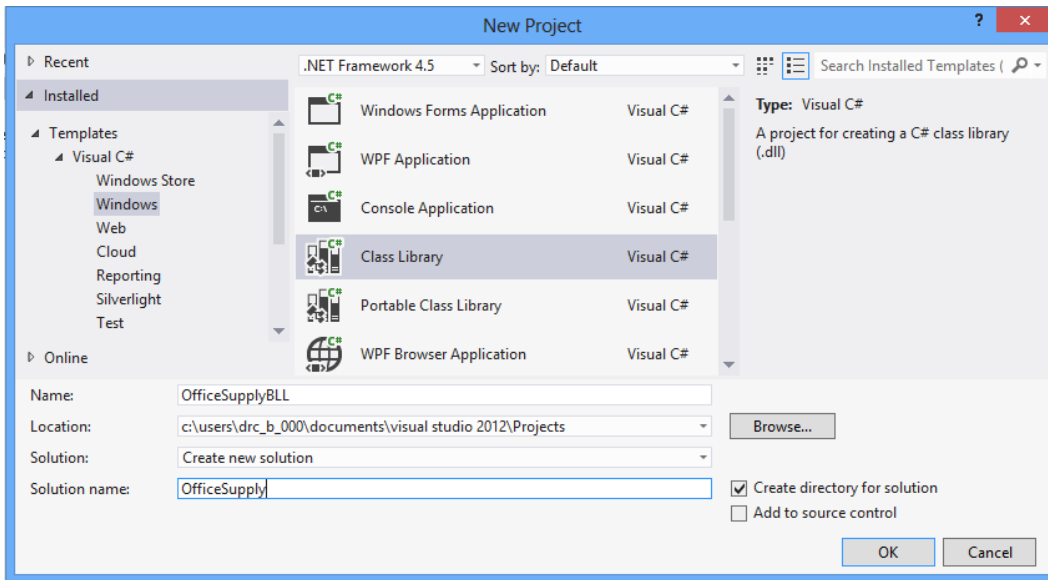


Figure 15-5. Creating the *OfficeSupplyBLL* class library

■ **Note** If you don't want to code the Office Supply Ordering application from scratch, you can download it from the Apress web site. See Appendix C for details.

Next, you'll create a static class (*DALUtility*) that implements retrieving the database connection string from the *App.config* file. The other classes will call its *GetSqlConnection* to retrieve the connection string.

Add an *App.config* file to the project. In the *App.config* file, add the xml code below to set the connection information. Remember you may have to change the data source depending on your server setup.

```
<configuration>
  <connectionStrings>
    <add name="OSConnection"
      providerName="System.Data.SqlClient"
      connectionString="data source=localhost;initial catalog=OfficeSupply;
        integrated security=True;"/>
  </connectionStrings>
</configuration>
```

In the Solution Explorer window, right-click on the References folder and select Add Reference. In the framework assemblies, select the *System.Configuration* assembly as shown in Figure 15-6.

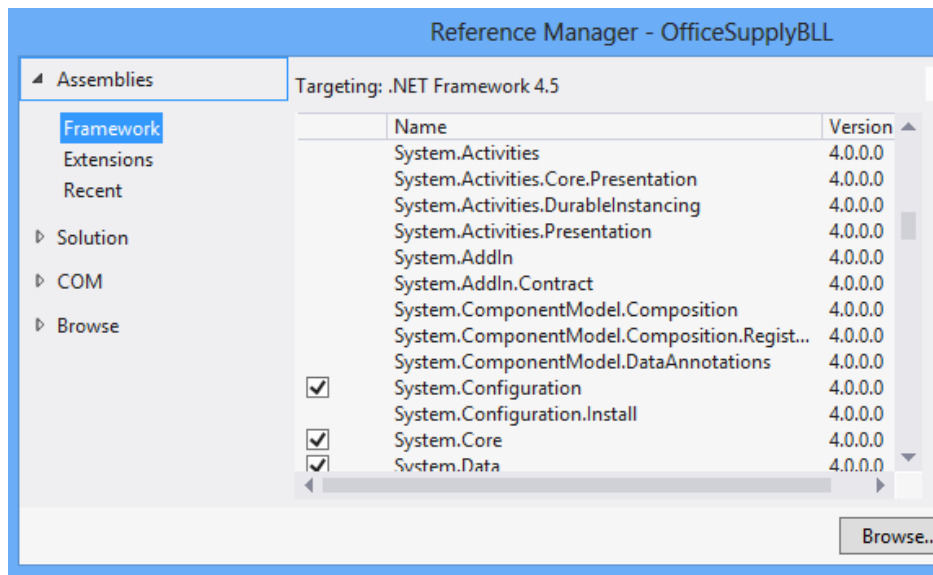


Figure 15-6. Adding a reference to the application

Add a class to the project and name it DALUtility. Add the following code to the class file:

```
using System;
using System.Configuration;

namespace OfficeSupplyBLL
{
    public static class DALUtility
    {
        public static string GetSQLConnection(string name)
        {
            // Assume failure.
            string returnValue = null;

            // Look for the name in the connectionStrings section.
            ConnectionStringSettings settings
                = ConfigurationManager.ConnectionStrings[name];

            // If found, return the connection string.
            if (settings != null)
            {
                returnValue = settings.ConnectionString;
            }
            return returnValue;
        }
    }
}
```

The next class to add is the `DALEmployee` class. This class contains a `Login` method that checks the user name and password supplied to the values in the database. It uses a `SqlCommand` object to execute a SQL statement against the database. If a match is found, it returns the employee's `userId`. If no match is found, it returns `-1`. Since a single value is returned by the SQL statement, you can use the `ExecuteScalar` method of the `SqlCommand` object.

Add a class named `DALEmployee` to the application and insert the following code into the class file:

```
using System;
using System.Data.SqlClient;

namespace OfficeSupplyBLL
{
    class DALEmployee
    {
        public int LogIn(string userName, string password)
        {
            string connString = DALUtility.GetSQLConnection("OSConnection");
            using( SqlConnection conn = new SqlConnection(connString))

            {
                using (SqlCommand cmd = new SqlCommand())
                {
                    cmd.Connection = conn;
                    cmd.CommandText = "Select EmployeeID from Employee where "
                        + " UserName = @UserName and Password = @Password ";
                    cmd.Parameters.AddWithValue("@UserName", userName);
                    cmd.Parameters.AddWithValue("@Password", password);
                    int userId;
                    conn.Open();
                    userId = (int)cmd.ExecuteScalar();
                    if (userId > 0)
                    {
                        return userId;
                    }
                    else
                    {
                        return -1;
                    }
                }
            }
        }
    }
}
```

The next class to construct is the `DALProductCatalog` class, the purpose of which is to encapsulate the functionality the application needs to retrieve and list the available products in the database. You also want to be able to view the products based on the category to which they belong. The information you need is in two database tables: the catalog table and the products table. These two tables are related through the `CatID` field.

When a client requests the product catalog information, a dataset is created and returned to the client. This service is provided in the DALProductCatalog class's GetProductInfo method. Add a class named DALProductCatalog to the project and insert the code shown in here:

```
using System;
using System.Collections.Generic;
using System.Data;
using System.Data.SqlClient;
using System.Linq;
using System.Text;
using System.Threading.Tasks;

namespace OfficeSupplyBLL
{
    public class DALProductCatalog
    {
        public DataSet GetProductInfo()
        {
            DataSet _dsProducts;
            string connString = DALUtility.GetSQLConnection("OSConnection");
            using (SqlConnection _conn = new SqlConnection(connString))
            {
                _dsProducts = new DataSet("Products");
                //Get category info
                String strSQL = "Select CategoryId, Name, Description from Category";
                using (SqlCommand cmdSelCategory = new SqlCommand(strSQL, _conn))
                {
                    using (SqlDataAdapter daCategory = new SqlDataAdapter(cmdSelCategory))
                    {
                        daCategory.Fill(_dsProducts, "Category");
                    }
                }

                //Get product info
                String strSQL2 = "Select ProductID, CategoryID, Name," +
                    "Description, UnitCost from Product";
                using (SqlCommand cmdSelProduct = new SqlCommand(strSQL2, _conn))
                {
                    using (SqlDataAdapter daProduct = new SqlDataAdapter(cmdSelProduct))
                    {
                        daProduct.Fill(_dsProducts, "Product");
                    }
                }
            }

            //Set up the table relation
            DataRelation drCat_Prod = new DataRelation("drCat_Prod",
                _dsProducts.Tables["Category"].Columns["CategoryID"],
                _dsProducts.Tables["Product"].Columns["CategoryID"], false);
            _dsProducts.Relations.Add(drCat_Prod);
            return _dsProducts;
        }
    }
}
```

When a client is ready to submit an order, it will call the `PlaceOrder` method of the `Order` class, which you will define shortly in the business logic classes. The client will pass the employee ID into the method and receive an order number as a return value. The `PlaceOrder` method of the `Order` class will pass the order information in the form of an XML string to the `DALOrder` class for processing. The `DALOrder` class contains the `PlaceOrder` method that receives an XML order string from the `Order` class and passes it into a stored procedure in the SQL Server database. The stored procedure updates the database and passes back the order number. This order number is then returned to the `Order` class, which in turn passes it back to the client.

Add a class named `DALOrder` to the project and insert the following code into the class file:

```
using System;
using System.Data;
using System.Data.SqlClient;

namespace OfficeSupplyBLL
{
    class DALOrder
    {
        public int PlaceOrder(string xmlOrder)
        {
            string connString = DALUtility.GetSQLConnection("OSConnection");
            using (SqlConnection cn = new SqlConnection(connString))
            {
                using (SqlCommand cmd = cn.CreateCommand())
                {
                    cmd.CommandType = CommandType.StoredProcedure;
                    cmd.CommandText = "up_PlaceOrder";
                    SqlParameter inParameter = new SqlParameter();
                    inParameter.ParameterName = "@xmlOrder";
                    inParameter.Value = xmlOrder;
                    inParameter.DbType = DbType.String;
                    inParameter.Direction = ParameterDirection.Input;
                    cmd.Parameters.Add(inParameter);
                    SqlParameter ReturnParameter = new SqlParameter();
                    ReturnParameter.ParameterName = "@OrderID";
                    ReturnParameter.Direction = ParameterDirection.ReturnValue;
                    cmd.Parameters.Add(ReturnParameter);
                    int intOrderNo;
                    cn.Open();
                    cmd.ExecuteNonQuery();
                    cn.Close();
                    intOrderNo = (int)cmd.Parameters["@OrderID"].Value;
                    return intOrderNo;
                }
            }
        }
    }
}
```

Now that you have constructed the data access layer classes, you are ready to construct the business logic layer set of classes.

Building the OSO Application's Business Logic Layer

Add a class named `Employee` to the project. This class will encapsulate employee information used by the UI and pass a login request to the data access layer. Add the following code to the `Employee.cs` file:

```
using System;

namespace OfficeSupplyBLL
{
    public class Employee
    {
        int _employeeID;

        public int EmployeeID
        {
            get { return _employeeID; }
            set { _employeeID = value; }
        }

        string _loginName;

        public string LoginName
        {
            get { return _loginName; }
            set { _loginName = value; }
        }

        string _password;

        public string Password
        {
            get { return _password; }
            set { _password = value; }
        }

        Boolean _loggedIn = false;

        public Boolean LoggedIn
        {
            get { return _loggedIn; }
        }

        public Boolean LogIn()
        {
            DALEmployee dbEmp = new DALEmployee();
            int empId;
            empId = dbEmp.LogIn(this.LoginName, this.Password);
            if (empId > 0)
            {
                this.EmployeeID = empId;
                this._loggedIn = true;
                return true;
            }
        }
    }
}
```

```

        else
        {
            this._loggedIn = false;
            return false;
        }
    }
}

```

The `ProductCatalog` class provides the `Product` dataset to the UI. It retrieves the dataset from the `DALProductCatalog` class. You could perform any business logic on the `DataSet` before passing it to the UI. Add a new `ProductCatalog` class file to the project and add the following code to implement the `ProductCatalog` class.

```

using System;
using System.Collections.Generic;
using System.Data;
using System.Linq;
using System.Text;
using System.Threading.Tasks;

namespace OfficeSupplyBLL
{
    public class ProductCatalog
    {
        public DataSet GetProductInfo()
        {
            //perform any business logic befor passing to client.
            // None needed at this time.
            DALProductCatalog prodCatalog = new DALProductCatalog();
            return prodCatalog.GetProductInfo();
        }
    }
}

```

When a user adds items to an order, the order item information is encapsulated in an `OrderItem` class. This class implements the `INotifyPropertyChanged` interface. This interface is necessary to notify the UI that a property changed so that it can update any controls bound to the class. It also overrides the `ToString` method to provide an XML string containing the item information. This string will get passed to the DAL when an order is placed.

Add the `OrderItem` class to the project and insert the following code to implement it.

```

using System;
using System.ComponentModel;

namespace OfficeSupplyBLL
{
    public class OrderItem : INotifyPropertyChanged
    {
        #region INotifyPropertyChanged Members
        public event PropertyChangedEventHandler PropertyChanged;
        protected void Notify(string propName)

```

```

    {
        if (this.PropertyChanged != null)
        {
            PropertyChanged(this, new PropertyChangedEventArgs(propName));
        }
    }
#endregion

string _ProdID;
int _Quantity;
double _UnitPrice;
double _SubTotal;
public string ProdID
{
    get { return _ProdID; }
    set { _ProdID = value; }
}
public int Quantity
{
    get { return _Quantity; }
    set
    {
        _Quantity = value;
        Notify("Quantity");
    }
}
public double UnitPrice
{
    get { return _UnitPrice; }
    set { _UnitPrice = value; }
}
public double SubTotal
{
    get { return _SubTotal; }
}
public OrderItem(String productID, double unitPrice, int quantity)
{
    _ProdID = productID;
    _UnitPrice = unitPrice;
    _Quantity = quantity;
    _SubTotal = _UnitPrice * _Quantity;
}
public override string ToString()
{
    string xml = "<OrderItem";
    xml += " ProductID='" + _ProdID + "'";
    xml += " Quantity='" + _Quantity + "'";
    xml += " />";
    return xml;
}
}
}

```

The final class of the business logic layer is the `Order` class. This class is responsible for maintaining a collection of order items. It has methods for adding and deleting items as well as passing the items to the `DALOrder` class when an order is placed.

After adding an `Order` class to the project, use the following code to implement the class:

```
using System;
using System.Collections.ObjectModel;

namespace OfficeSupplyBLL
{
    public class Order
    {
        ObservableCollection<OrderItem> _orderItemList = new
            ObservableCollection<OrderItem>();

        public ObservableCollection<OrderItem> OrderItemList
        {
            get { return _orderItemList; }
        }
        public void AddItem(OrderItem orderItem)
        {
            foreach (var item in _orderItemList)
            {
                if (item.ProdID == orderItem.ProdID)
                {
                    item.Quantity += orderItem.Quantity;

                    return;
                }
            }
            _orderItemList.Add(orderItem);
        }
        public void RemoveItem(string productID)
        {
            foreach (var item in _orderItemList)
            {
                if (item.ProdID == productID)
                {
                    _orderItemList.Remove(item);
                    return;
                }
            }
        }
        public double GetOrderTotal()
        {
            if (_orderItemList.Count == 0)
            {
                return 0.00;
            }
        }
    }
}
```

```

        else
        {
            double total = 0;
            foreach (var item in _orderItemList)
            {
                total += item.SubTotal;
            }
            return total;
        }
    }
}
public int PlaceOrder(int employeeID)
{
    string xmlOrder;
    xmlOrder = "<Order EmployeeID='" + employeeID.ToString() + "'>";
    foreach (var item in _orderItemList)
    {
        xmlOrder += item.ToString();
    }
    xmlOrder += "</Order>";
    DALOrder dbOrder = new DALOrder();
    return dbOrder.PlaceOrder(xmlOrder);
}
}
}
}

```

Now that you have constructed the data access and business logic layers of the OSO application, you are ready to construct the user interface (UI). Because you created the application in distinct tiers, you have the option of using a Windows-based WPF UI, an ASP.NET web-based UI or a Windows Store app. You could also create more than one UI depending on the device users use to interact with the application. In the next section, you will construct a WPF application users will use to place office supply orders from their desktop computers.

Creating the OSO Application UI

In order to create the ordering system's WPF interface, you'll need to add a WPF project to the solution containing the OfficeSupplyBLL project.

In Visual Studio, add a WPF project to the solution and name it OfficeSupplyWPF. After the project is loaded add a reference to the OfficeSupplyBLL project as shown in Figure 15-7.

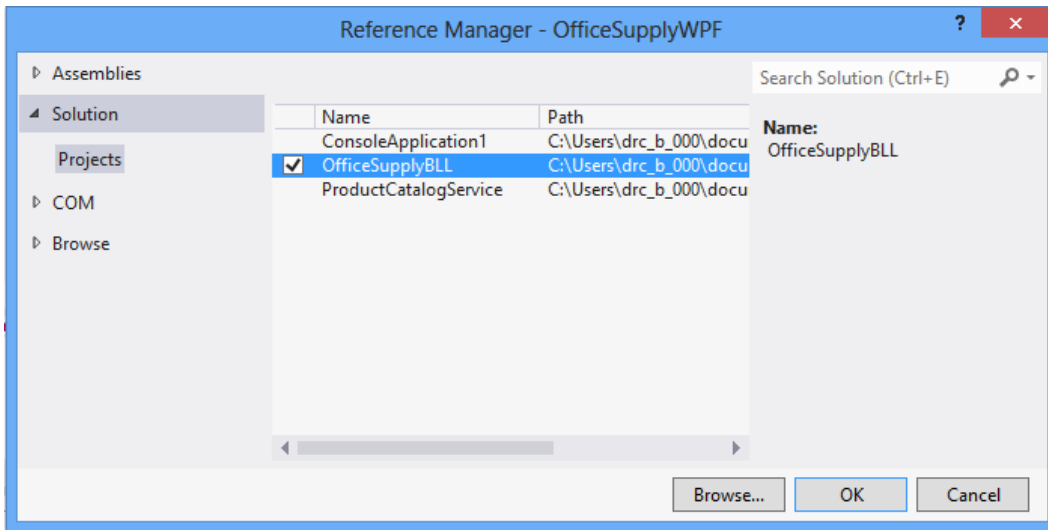
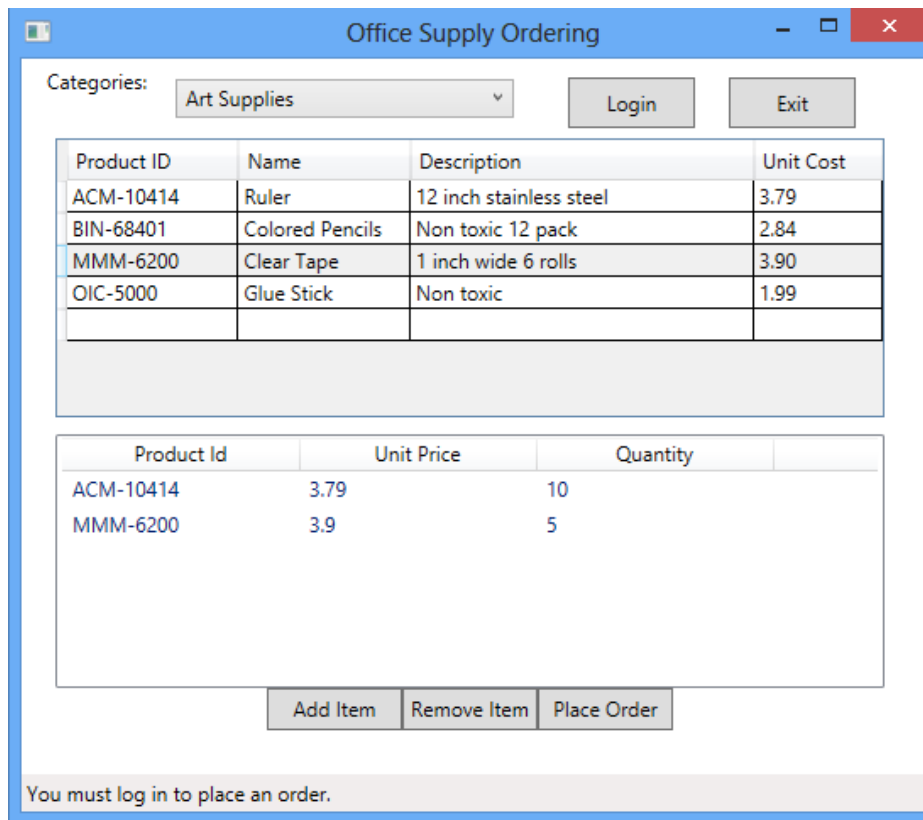


Figure 15-7. Adding a reference to the *OfficeSupplyBLL* project

In the *App.config* file, add the xml code below just before the configuration end tag. Remember you may have to change the data source depending on your server setup.

```
<connectionStrings>
  <add name="OSConnection"
    providerName="System.Data.SqlClient"
    connectionString="data source=localhost;initial catalog=OfficeSupply;
    integrated security=True;" />
</connectionStrings>
```

The first goal of the user interface is to present information about the products that can be ordered. The product information is presented in a *DataGrid* control. The user will view products in a particular category by selecting the category in a *ComboBox* control. Once products are listed, users can add products to an order. When a product is added to an order, it's displayed in a *ListView* below the *DataGrid*. Figure 15-8 shows the OSO order form with the items added to an order.



Office Supply Ordering

Categories: Art Supplies [v] Login Exit

Product ID	Name	Description	Unit Cost
ACM-10414	Ruler	12 inch stainless steel	3.79
BIN-68401	Colored Pencils	Non toxic 12 pack	2.84
MMM-6200	Clear Tape	1 inch wide 6 rolls	3.90
OIC-5000	Glue Stick	Non toxic	1.99

Product Id	Unit Price	Quantity
ACM-10414	3.79	10
MMM-6200	3.9	5

Add Item Remove Item Place Order

You must log in to place an order.

Figure 15-8. Form for adding items to an order

Add the following XAML code to the `MainWindow.xaml` file to create the OSO order form. Notice the use of data binding for the various controls. For example, the item source of the data grid is set to the data relation between the `Category` and `Product` tables (`ItemsSource="{Binding drCat_Prod}"`). This allows the data grid to show the products in the category selected in the dropdown list.

```
<Window x:Class="OfficeSupplyWPF.MainWindow"
  xmlns="http://schemas.microsoft.com/winfx/2006/xaml/presentation"
  xmlns:x="http://schemas.microsoft.com/winfx/2006/xaml"
  Title="Office Supply Ordering" Height="484" Width="550" Loaded="Window_Loaded">
  <Grid>
    <StackPanel Name="LayoutRoot" DataContext="{Binding}"
      Orientation="Vertical" HorizontalAlignment="Left" Height="auto" Width="auto">
      <StackPanel Orientation="Horizontal" HorizontalAlignment="Left">
        <Label Content="Categories:" Margin="10,0,0,0"/>
        <ComboBox ItemsSource="{Binding}" Name="categoriesComboBox"
          IsSynchronizedWithCurrentItem="True"
          DisplayMemberPath="Name" Height="23" Margin="12" Width="200" >
        <ComboBox.ItemsPanel>
          <ItemsPanelTemplate>
            <VirtualizingStackPanel />
          </ItemsPanelTemplate>
        </ComboBox.ItemsPanel>
      </StackPanel>
    </StackPanel>
  </Grid>
```

```

        </ComboBox.ItemsPanel>
    </ComboBox>
    <Button Content="Login" Height="30" Name="loginButton"
        Width="75" Margin="20,5,0,0" Click="loginButton_Click" />
    <Button Content="Exit" Height="30" Name="exitButton"
        Width="75" Margin="20,5,0,0" Click="exitButton_Click" />
</StackPanel>
<DataGrid AutoGenerateColumns="False" Height="165"
    ItemsSource="{Binding drCat_Prod}"
    Name="ProductsDataGrid" RowDetailsVisibilityMode="VisibleWhenSelected"
    Width="490" HorizontalAlignment="Left" Margin="20,0,20,10"
    SelectionMode="Single">
    <DataGrid.Columns>
        <DataGridTextColumn
            x:Name="productIDColumn" Binding="{Binding Path=ProductID}"
            Header="Product ID" Width="40*" />
        <DataGridTextColumn
            x:Name="nameColumn" Binding="{Binding Path=Name}"
            Header="Name" Width="40*" />
        <DataGridTextColumn
            x:Name="descriptColumn" Binding="{Binding Path=Description}"
            Header="Description" Width="80*" />
        <DataGridTextColumn
            x:Name="unitCostColumn" Binding="{Binding Path=UnitCost}"
            Header="Unit Cost" Width="30*" />
    </DataGrid.Columns>
</DataGrid>

<StackPanel Orientation="Vertical">
    <ListView Name="orderListView" MinHeight="150" Width="490"
        ItemsSource="{Binding}" SelectionMode="Single">
        <ListView.View>
            <GridView>
                <GridViewColumn Width="140" Header="Product Id"
                    DisplayMemberBinding="{Binding ProdID}" />
                <GridViewColumn Width="140" Header="Unit Price"
                    DisplayMemberBinding="{Binding UnitPrice}" />
                <GridViewColumn Width="140" Header="Quantity"
                    DisplayMemberBinding="{Binding Quantity}" />
            </GridView>
        </ListView.View>
    </ListView>

</StackPanel>
<StackPanel Orientation="Horizontal" HorizontalAlignment="Center">
    <Button Name="addButton" MinHeight="25" MinWidth="80"
        Content="Add Item" Click="addButton_Click" />
    <Button Name="removeButton" MinHeight="25" MinWidth="80"
        Content="Remove Item" Click="removeButton_Click"/>
    <Button Name="placeOrderButton" MinHeight="25" MinWidth="80"
        Content="Place Order" Click="placeOrderButton_Click"/>

```



```

        </StackPanel>
    </StackPanel>
    <StatusBar VerticalAlignment="Bottom" HorizontalAlignment="Stretch">
        <TextBlock Name="statusTextBlock">You must log in to place an order.</TextBlock>
    </StatusBar>
</Grid>
</Window>

```

To add an order item, the user first selects a row in the DataGrid and then selects the Add Item button. The Add Item button displays a dialog box the user uses to enter a quantity and add the item. Figure 15-9 shows the Order Item dialog.

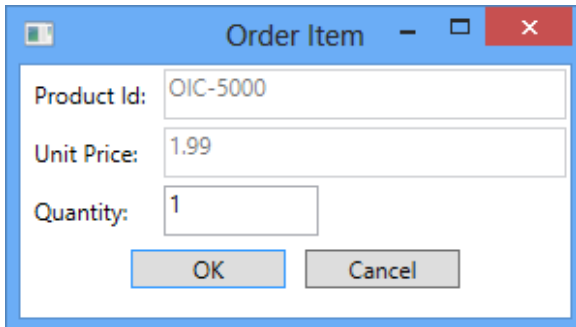


Figure 15-9. The Order Item dialog

Add a new Window to the project named OrderItemDialog.xaml. Add the following XAML code to create the OrderItemDialog form:

```

<Window x:Class="OfficeSupplyWPF.OrderItemDialog"
    xmlns="http://schemas.microsoft.com/winfx/2006/xaml/presentation"
    xmlns:x="http://schemas.microsoft.com/winfx/2006/xaml"
    WindowStartupLocation="CenterOwner"
    Title="Order Item" Height="169" Width="300">
    <Grid>
        <Grid.ColumnDefinitions>
            <ColumnDefinition Width="Auto" />
            <ColumnDefinition Width="Auto" />
            <ColumnDefinition />
        </Grid.ColumnDefinitions>

        <Grid.RowDefinitions>
            <RowDefinition Height="Auto" />
            <RowDefinition Height="Auto" />
            <RowDefinition Height="Auto" />
            <RowDefinition Height="Auto" />
            <RowDefinition />
        </Grid.RowDefinitions>
        <Label Grid.Column="0" Grid.Row="0" Margin="2">Product Id:</Label>
        <TextBox Name="productIdTextBox" Grid.Column="1"
            Grid.Row="0" Margin="2" Grid.ColumnSpan="2" IsEnabled="False"/>
    </Grid>

```

```

<Label Grid.Column="0" Grid.Row="1" Margin="2">Unit Price:</Label>
<TextBox Name="unitPriceTextBox" Grid.Column="1"
    Grid.Row="1" Margin="2" Grid.ColumnSpan="2" IsEnabled="False"/>
<Label Grid.Column="0" Grid.Row="2" Margin="2" >Quantity:</Label>
<TextBox Name="quantityTextBox" Grid.Column="1"
    Grid.Row="2" Margin="2" MinWidth="80" Text="1"/>
<StackPanel Grid.Column="0" Grid.ColumnSpan="3"
    Grid.Row="3" Orientation="Horizontal"
    HorizontalAlignment="Center">
    <Button Name="okButton" Click="okButton_Click" IsDefault="True"
        MinWidth="80" Margin="5">OK</Button>
    <Button Name="cancelButton" Click="cancelButton_Click" IsCancel="True"
        MinWidth="80" Margin="5">Cancel</Button>
</StackPanel>
</Grid>
</Window>

```

Before users can submit an order, they must log in. When they click on the Login button, they are presented with a Login Dialog window, shown in Figure 15-10.




Figure 15-10. The Login dialog

Add a new window to the project named LoginDialog.xaml. Add the following XAML code to create the LoginDialog form.

```

<Window x:Class="OfficeSupplyWPF.LoginDialog"
    xmlns="http://schemas.microsoft.com/winfx/2006/xaml/presentation"
    xmlns:x="http://schemas.microsoft.com/winfx/2006/xaml"
    Title="Login" Height="131" Width="300"
    WindowStartupLocation="CenterOwner"
    FocusManager.FocusedElement="{Binding ElementName=nameTextBox}">
<Grid>
    <Grid.ColumnDefinitions>
        <ColumnDefinition Width="Auto" />
        <ColumnDefinition />
    </Grid.ColumnDefinitions>

    <Grid.RowDefinitions>
        <RowDefinition Height="Auto" />
        <RowDefinition Height="Auto" />
    </Grid.RowDefinitions>

```

```

        <RowDefinition Height="Auto" />
        <RowDefinition />
    </Grid.RowDefinitions>
    <Label Grid.Column="0" Grid.Row="0" Margin="2">Name:</Label>
    <TextBox Name="nameTextBox" Grid.Column="1" Grid.Row="0" Margin="2"/>
    <Label Grid.Column="0" Grid.Row="1" Margin="2">Password:</Label>
    <PasswordBox Name="passwordTextBox" Grid.Column="1" Grid.Row="1" Margin="2"/>

    <StackPanel Grid.Column="0" Grid.ColumnSpan="2" Grid.Row="2"
        Orientation="Horizontal" HorizontalAlignment="Center">
        <Button Name="okButton" Click="okButton_Click" IsDefault="True"
            MinWidth="80" Margin="5">OK</Button>
        <Button Name="cancelButton" Click="cancelButton_Click" IsCancel="True"
            MinWidth="80" Margin="5">Cancel</Button>
    </StackPanel>
</Grid>
</Window>

```

Now that you have created the windows that make up the UI, you are ready to add the implementation to the windows' codebehind files.

Add the following code to the MainWindow.xaml.cs codebehind file. A discussion of the code follows.

```

using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using System.Windows;
using System.Windows.Controls;
using System.Windows.Data;
using System.Windows.Documents;
using System.Windows.Input;
using System.Windows.Media;
using System.Windows.Media.Imaging;
using System.Windows.Navigation;
using System.Windows.Shapes;
using System.Data;
using OSOBLL;
using System.Collections.ObjectModel;

namespace OfficeSupplyWPF
{
    /// <summary>
    /// Interaction logic for MainWindow.xaml
    /// </summary>
    public partial class MainWindow : Window
    {

        DataSet _dsProdCat;
        Employee _employee;
        Order _order;
    }
}

```

```

public MainWindow()
{
    InitializeComponent();
}

private void Window_Loaded(object sender, RoutedEventArgs e)
{
    ProductCatalog prodCat = new ProductCatalog();
    _dsProdCat = prodCat.GetProductInfo();
    this.DataContext = _dsProdCat.Tables["Category"];
    _order = new Order();
    _employee = new Employee();
    this.orderListView.ItemsSource = _order.OrderItemList;
}

private void loginButton_Click(object sender, RoutedEventArgs e)
{
    LoginDialog dlg = new LoginDialog();
    dlg.Owner = this;
    dlg.ShowDialog();
    // Process data entered by user if dialog box is accepted
    if (dlg.DialogResult == true)
    {
        _employee.LoginName = dlg.nameTextBox.Text;
        _employee.Password = dlg.passwordTextBox.Password;
        if (_employee.LogIn() == true)
        {
            this.statusTextBlock.Text = "You are logged in as employee number " +
                _employee.EmployeeID.ToString();
        }
        else
        {
            MessageBox.Show("You could not be verified. Please try again.");
        }
    }
}

private void exitButton_Click(object sender, RoutedEventArgs e)
{
    this.Close();
}

private void addButton_Click(object sender, RoutedEventArgs e)
{
    OrderItemDialog orderItemDialog = new OrderItemDialog();

    DataRowView selectedRow;
    selectedRow = (DataRowView)this.ProductsDataGrid.SelectedItems[0];
}

```


The `loginButton_Click` event launches an instance of the `LoginDialog` window and checks the `Dialog` result. If it comes back as `true`, the `_employee` object's values are set to the values entered in the dialog and the `Login` method of the `Employee` class is called. If the `Login` method returns `true`, the user is notified that they are logged in.

```
private void loginButton_Click(object sender, RoutedEventArgs e)
{
    LoginDialog dlg = new LoginDialog();
    dlg.Owner = this;
    dlg.ShowDialog();
    // Process data entered by user if dialog box is accepted
    if (dlg.DialogResult == true)
    {
        _employee.LoginName = dlg.nameTextBox.Text;
        _employee.Password = dlg.passwordTextBox.Password;
        if (_employee.Login() == true)
        {
            this.statusTextBlock.Text = "You are logged in as employee number " +
                _employee.EmployeeID.ToString();
        }
        else
        {
            MessageBox.Show("You could not be verified. Please try again.");
        }
    }
}
```

The `addButton_Click` event launches an instance of the `OrderItemDialog` window and fills the textboxes with information from the selected row of the `ProductsDataGrid`. If the `DialogResult` returns `true`, the information entered in the dialog is used to create an `OrderItem` object and add it to the `Order`'s `OrderItem` collection.

```
private void addButton_Click(object sender, RoutedEventArgs e)
{
    OrderItemDialog orderItemDialog = new OrderItemDialog();

    DataRowView selectedRow;
    selectedRow = (DataRowView)this.ProductsDataGrid.SelectedItems[0];
    orderItemDialog.productIdTextBox.Text = selectedRow.Row.ItemArray[0].ToString();
    orderItemDialog.unitPriceTextBox.Text = selectedRow.Row.ItemArray[4].ToString();
    orderItemDialog.Owner = this;
    orderItemDialog.ShowDialog();
    if (orderItemDialog.DialogResult == true )
    {
        string productId = orderItemDialog.productIdTextBox.Text;
        double unitPrice = double.Parse(orderItemDialog.unitPriceTextBox.Text);
        int quantity = int.Parse(orderItemDialog.quantityTextBox.Text);
        _order.AddItem(new OrderItem(productId,unitPrice,quantity));
    }
}
```

The `removeButton_Click` event checks to see if an item is selected in the `orderListView` and removes it from the order.

```
private void removeButton_Click(object sender, RoutedEventArgs e)
{
    if (this.orderListView.SelectedItem != null)
    {
        var selectedOrderItem = this.orderListView.SelectedItem as OrderItem;
        _order.RemoveItem(selectedOrderItem.ProdID);
    }
}
```

The `placeOrderButton_Click` event checks to see if the user is logged in and places the order if they are.

```
private void placeOrderButton_Click(object sender, RoutedEventArgs e)
{
    if (_employee.LoggedIn == true)
    {
        //place order
        int orderId;
        orderId = _order.PlaceOrder(_employee.EmployeeID);
        MessageBox.Show("Your order has been placed. Your order id is " + orderId.ToString());
    }
    else
    {
        MessageBox.Show("You must be logged in to place an order.");
    }
}
```

Now that the `MainWindow`'s codebehind is implemented, you are ready to add the codebehind for the dialog widows.

Add the following code to the `OrderItemDialog.xaml.cs` codebehind file. If the user clicks the OK button, the `DialogResult` is set to `true`. If the user clicks cancel, the `DialogResult` is set to `false`.

```
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using System.Windows;
using System.Windows.Controls;
using System.Windows.Data;
using System.Windows.Documents;
using System.Windows.Input;
using System.Windows.Media;
using System.Windows.Media.Imaging;
using System.Windows.Shapes;

namespace OfficeSupplyWPF
{
    /// <summary>
    /// Interaction logic for OrderItemDialog.xaml
    /// </summary>
}
```

```

public partial class OrderItemDialog : Window
{
    public OrderItemDialog()
    {
        InitializeComponent();
    }

    private void okButton_Click(object sender, RoutedEventArgs e)
    {
        this.DialogResult = true;
    }

    private void cancelButton_Click(object sender, RoutedEventArgs e)
    {
        this.DialogResult = false;
    }
}
}

```

Add the following code to the LoginDialog.xaml.cs codebehind file. It is similar to OrderItemDialog code.

```

using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using System.Windows;
using System.Windows.Controls;
using System.Windows.Data;
using System.Windows.Documents;
using System.Windows.Input;
using System.Windows.Media;
using System.Windows.Media.Imaging;
using System.Windows.Shapes;

namespace OfficeSupplyWPF
{
    /// <summary>
    /// Interaction logic for LoginDialog.xaml
    /// </summary>
    public partial class LoginDialog : Window
    {
        public LoginDialog()
        {
            InitializeComponent();
        }

        private void okButton_Click(object sender, RoutedEventArgs e)
        {
            this.DialogResult = true;
        }
    }
}

```



```
private void cancelButton_Click(object sender, RoutedEventArgs e)
{
    this.DialogResult = false;
}
}
```

Now that you have added the implementation code to the UI, you are ready to test the application. Take some time to play with the application. Add some items to the order and test removing some items from the order. Try placing an order (a valid login is JSmith for the user and js for the password). As you play with the application, think of some things you could do to improve it. One area that needs improvement is trapping for errors that may occur. For example, if you click the Add Item button without selecting an item in the grid, the program will crash. One way to deal with this is to disable the button unless a grid row is selected.

You may also want to expand the application. For example, you could implement a Windows Store App interface. Another interesting project is developing a web service that accepts the orders placed by the application.

■ **Note** Although this is a functional application, it's for demonstration purposes only and is not production ready.

Summary

In this chapter, you revisited the office supply ordering (OSO) application designed in Chapter 4. You created a functional application incorporating the concepts you learned in the previous chapters. The application contains a data access layer, a business logic layer, and a user interface layer. You learned why applications are split into different layers and how to construct a working application comprised of the various layers. Although you didn't create a web-based user interface application layer, because you created the application in distinct tiers, you could easily replace the Windows-based WPF UI with a web-based UI or a Windows Store app. You also created the data access layer using ADO.NET. Another option is to use the Entity Framework (EF). You could also create a web service to send the orders to the suppliers.

CHAPTER 16



Wrapping Up

If you've made it this far, take a moment and pat yourself on the back. You've come a long way since the day you first cracked open the cover of this book; you've gained valuable skills and learned concepts you can use to successfully program using the .NET Framework, C#, and the Visual Studio IDE. These include, but are not limited to, the following:

- the importance of the application design cycle
- the Unified Modeling Language (UML) and how it can help facilitate the analysis and design of object-oriented programs
- the common language runtime (CLR)
- the structure of the .NET Framework
- how to create and use class structures and hierarchies
- how to implement inheritance, polymorphism, and interfaces
- object interaction and collaboration
- event-driven programming
- structured error handling
- how to work with data structures and data sources using ADO.NET
- using the Entity Framework to create object relational mappings to a SQL Server database
- how to use the features of the Visual Studio IDE to increase productivity and facilitate debugging
- how to implement a Windows-based graphical user interface using the Windows Presentation Framework
- how to implement a web-based graphical user interface using ASP.NET and Web Forms
- how to create a Windows Store app
- how to create and consume services using Windows Communication Foundation and the ASP.NET Web API

Congratulations! You can now call yourself a C# programmer (albeit a neophyte). However, don't get too high on yourself. If your goal is to become a *professional C#* programmer, your journey has just begun. The next stage of your development is to gain experience. In other words, design and code, and then design and code some more. If you are designing and coding C# at work, this will be easy. (Although it will be stressful if you are expected to be an expert after that three-day course they sent you to!)

If you are learning on your own, you will have to find the time and projects on which to work. This is easier than you might think. Commit to an hour a day and come up with an idea for a program. For example, you could design a program that converts recipes into Extensible Markup Language (XML) data. The XML data could then generate a shopping list. Heck, if you really want to go all out, incorporate an inventory tracking system that tracks ingredients you have in stock. However you go about gaining experience, remember the important adage: use it or lose it!

The following sections highlight some other important things to consider as you develop your programming skills.

Improve Your Object-Oriented Design Skills

Object-oriented analysis and design are two of the hardest tasks you will perform as a programmer. These are not skills that come easily for most programmers. They are, however, two of the most important skills you should strive to master. They are what separate what I call a *programmer* from a *coder*. If you talk to most CIOs and programming managers, finding coders is easy; it is the programmer they are after.

Remember that there is no one “correct” method, rather several that are equally valid.

Investigate the .NET Framework Namespaces

The .NET Framework contains a vast number of classes, interfaces, and other types aimed at optimizing and expediting your development efforts. The various namespaces that make up the .NET Framework Class Library are organized by functionality. It’s important you take the time to become familiar with the capabilities provided by these namespaces.

Start out with the namespaces that incorporate functionality you will use most often, such as the root namespace `System` and the `System.IO`, which contains the .NET Framework classes for reading and writing to streams and files.

After you become familiar with the more common namespaces, explore some of the more obscure ones. For example, `System.Security.Cryptography` provides cryptographic services such as data encoding, hashing, and message authentication. You will be amazed at the extent of the support provided by the framework. You can find a wealth of information on the members of the various namespaces in Visual Studio’s integrated documentation.

Become Familiar with ADO.NET and the Entity Framework

Data is fundamental to programming. You store, retrieve, and manipulate data in every program you write. The data structure a program works with during execution is *nondurable* data—it is held in RAM. When the application terminates, this data is lost and has to be re-created the next time the application runs. *Durable data* is data that is maintained in a permanent data structure such as a file system or a database. Most programs need to retrieve data from and persist data to some sort of durable data storage. This is where ADO.NET steps in. ADO.NET refers to the namespaces that contain the functionality for working with durable data. (It also contains functionality for organizing and working with nondurable data in a familiar relational database or XML-type structure.) Although I have introduced you to ADO.NET and the Entity Framework, this is such an important topic that it deserves a book devoted solely to these data access technologies. (Don’t worry—there are many!) This is definitely an area to which you need to devote further study. To learn more about these technologies, visit the Data Developer Center site at <http://msdn.microsoft.com/en-us/data>.

Learn More about Creating Great User Interfaces (UI)

Although you were introduced to WPF, ASP.NET, and Windows Store Apps, I only scratched the surface of these powerful technologies. WPF, ASP.NET, and Windows Store apps are packed full of features for developing engaging, interactive user experiences on the web, desktop, and mobile devices. For more information on programming

WPF, visit the Windows Client development center at <http://windowsclient.net>. For more information about programming in ASP.NET, visit the developer center at <http://www.asp.net>. If you are interested in Windows Store App development, visit the developer center at <http://msdn.microsoft.com/en-US/windows/apps/br229512>. These sites are full of learning materials and demo applications showcasing the power of these technologies.

Move toward Component-Based Development

After you have mastered object-oriented development and the encapsulation of your programming logic in a class system, you are ready to move toward component-based development. *Components* are assemblies that further encapsulate the functionality of your programming logic. Although the OSO application's business logic tier is logically isolated from the data access tier, physically they reside in the same assembly. You can increase code maintenance and reuse by compiling each into its own assembly. You should start moving to a Lego approach of application development. This is where your application is composed of a set of independent pieces (assemblies) that can be snapped together and work in conjunction to perform the necessary services. For more information on this and other best practices, go to the Microsoft's patterns and practices web site at <http://pnp.azurewebsites.net/en-us/>.

Find Help

An enormous amount of information is available on the .NET Framework and the C# programming language. The help system provided with Visual Studio is an excellent resource for programmers. Get in the habit of using this resource religiously.

Another extremely important resource is <http://msdn.microsoft.com>. This web site, provided by Microsoft for developers, contains a wealth of information including white papers, tutorials, and webcast seminars; quite honestly, it's one of the most informative sites in the industry. If you are developing your programming skills using Microsoft technologies, visiting this site should be as much of a routine as reading the daily paper. There are also a number of independent web sites dedicated to the various .NET programming languages. One good site is C# Corner (<http://www.c-sharpcorner.com/>), which contains tons of articles on all aspects of programming in C#. You can use your favorite search engine to discover other good sites on the Web dedicated to C# programming.

Join a User Group

Microsoft provides a great deal of support for the development of local .NET user groups. User groups consist of members with an interest in .NET programming. These groups provide a great avenue for learning, mentoring, and networking. There is a list of .NET user groups available at <http://msdn.microsoft.com>. The International .NET Association (INETA) also provides support for .NET user groups; you can find a listing of INETA affiliated user groups at <http://www.ineta.org>.

If you can't find a .NET user group in your area, heck, why not start one?

Please Provide Feedback

Although every effort has been made to provide you with an error-free text, it is inevitable that some mistakes will make it through the editing process. I am committed to providing updated errata at the Apress Web site (<http://www.apress.com>), but I can't do this without your help. If you have come across any mistakes while reading this text, please report them to me through the Apress site.

Thank You, and Good Luck!

I sincerely hope you found working your way through this text to be an enjoyable and worthwhile experience. I want to thank you for allowing me to be your guide on this journey. Just as your skills as a developer increased as a result of reading this book, my skills as a developer have increased immensely as a result of writing it. My experience of teaching and training for the past two decades has been that you really don't fully comprehend a subject until you can teach it to someone else. So, again, thank you and good luck!



Fundamental Programming Concepts

The following information is for readers who are new to programming and need a primer on some fundamental programming concepts. If you have programmed in another language, chances are the concepts presented in this appendix are not new to you. You should, however, review the material briefly to become familiar with the C# syntax.

Working with Variables and Data Types

Variables in programming languages store values that can change while the program executes. For example, if you wanted to count the number of times a user tries to log in to an application, you could use a variable to track the number of attempts. A variable is a memory location where a value is stored. Using the variable, your program can read or alter the value stored in memory. Before you use a variable in your program, however, you must declare it. When you declare a variable, the compiler also needs to know what kind of data will be stored at the memory location. For example, will it be numbers or letters? If the variable will store numbers, how large can a number be? Will the variable store decimals or only whole numbers? You answer these questions by assigning a data type to the variable. A login counter, for example, only needs to hold positive whole numbers. The following code demonstrates how you declare a variable named `counter` in C# with an integer data type:

```
int counter;
```

Specifying the data type is referred to as *strong typing*. Strong typing results in more efficient memory management, faster execution, and compiler type checking, all of which reduces runtime errors. Runtime errors are errors that can occur while a program is running and are not caught by the compiler when you build the application. For example, incorrect rounding can occur as a result of using the wrong data types in a calculation.

Once you declare the variable, you need to assign it an initial value, either in a separate statement or within the declaration statement itself. For instance, the following code:

```
int counter = 1;
```

is equivalent to this:

```
int counter;  
counter = 1;
```

Understanding Elementary Data Types

C# supports elementary data types such as numeric, character, and date types.

Integral Data Types

Integral data types represent whole numbers only. Table A-1 summarizes the integral data types used in C#.

Table A-1. *Integral Data Types*

Data Type	Storage Size	Value Range
Byte	8-bit	0 through 255
Short	16-bit	-32,768 through 32,767
Integer	32-bit	-2,147,483,648 through 2,147,483,647
Long	64-bit	-9,223,372,036,854,775,808 through 9,223,372,036,854,775,807

Obviously, memory size is important when choosing a data type for a variable. A less obvious consideration is how easily the compiler works with the data type. The compiler performs arithmetic operations with integers more efficiently than the other types. Often, it's better to use integers as counter variables even though a byte or short type could easily manage the maximum value reached.

Non-Integral Data Types

If a variable will store numbers that include decimal parts, then you must use a non-integral data type. C# supports the non-integral data types listed in Table A-2.

Table A-2. *Non-Integral Data Types*

Data Type	Storage Size	Value Range
Single	32-bit	-3.4028235E+38 through -1.401298E-45 for negative values; 1.401298E-45 through 3.4028235E+38 for positive values
Double	64-bit	1.79769313486231570E+308 through -4.94065645841246544E-324 for negative values; 4.94065645841246544E-324 through 1.79769313486231570E+308 for positive values
Decimal	128-bit	0 through +/-79,228,162,514,264,337,593,543,950,335 with no decimal point; 0 through +/-7.9228162514264337593543950335 with 28 places to the right of the decimal

The decimal data type holds a larger number of significant digits than either the single or the double data types and it is not subject to rounding errors. Decimal data types are usually reserved for financial or scientific calculations that require a higher degree of precision.

Character Data Types

Character data types are for variables that hold characters used in human languages. For example, a character data type holds letters such as "a" or numbers used for display and printing such as "2 apples." The character data types in C# are based on Unicode, which defines a character set that can represent the characters found in every language from English to Arabic and Mandarin Chinese. C# supports two character data types: `char` and `string`. The `char` data type holds single (16-bit) Unicode character values such as a or B. The `string` data type holds a sequence of Unicode characters. It can range from zero up to about two billion characters.

Boolean Data Type

The Boolean data type holds a 16-bit value that is interpreted as true or false. It's used for variables that can be one of only two values, such as yes or no, or on or off.

Date Data Type

Dates are held as 64-bit integers where each increment represents a period of elapsed time from the start of the Gregorian calendar (1/1/0001 at 12:00 A.M.).

Object Data Type

An object data type is a 32-bit address that points to the memory location of another data type. It is commonly used to declare variables when the actual data type they refer to can't be determined until runtime. Although the object data type can be a catch-all to refer to the other data types, it is the most inefficient data type when it comes to performance and should be avoided unless absolutely necessary.

Nullable Types

By default, value types such as the Boolean, integer, and double data types can't be assigned a null value. This can become problematic when retrieving data from data structures such as a database that does allow nulls. When declaring a value type variable that may be assigned a null, you make it a nullable type by placing a question mark symbol (?) after the type name, like so:

```
double salary = null; // Not allowed.
double? salary = null; // allowed.
```

Introducing Composite Data Types

Combining elementary data types creates composite data types. Structures, arrays, and classes are examples of composite data types.

Structures

A structure data type is useful when you want to organize and work with information that does not need the overhead of class methods and constructors. It's well suited for representing lightweight objects such as the coordinates of a point or rectangle. A single variable of type structure can store such information. You declare a structure with the `struct` keyword. For example, the following code creates a structure named `Point` to store the coordinates of a point in a two-dimensional surface:

```
public struct Point
{
    public int _x, _y;
```



```

public Point(int x, int y)
{
    _x = x;
    _y = y;
}

```

Once you define the structure, you can declare a variable of the structure type and create a new instance of the type, like so:

```
Point p1 = new Point(10,20);
```

Arrays

Arrays are often used to organize and work with groups of the same data type; for example, you may need to work with a group of names, so you declare an array data type by placing square brackets ([]) immediately following the variable name, like so:

```
string[] name;
```

The new operator is used to create the array and initialize its elements to their default values. Because the elements of the array are referenced by a zero-based index, the following array holds five elements:

```
string[] name = new string[4];
```

To initialize the elements of an array when the array is declared, you use curly brackets ({}) to list the values. Since the size of the array can be inferred, you do not have to state it.

```
string[] name = {"Bob", "Bill", "Jane", "Judy"};
```

C# supports multidimensional arrays. When you declare the array, you separate the size of the dimensions by commas. The following declaration creates a two-dimensional array of integers with five rows and four columns:

```
string[,] name = new string[4,3];
```

To initialize the elements of a two-dimensional array when the array is declared, you use curly brackets inside curly brackets to list the array elements.

```
int[,] intArray = {{1,2}, {3,4}, {5,6}, {7,8}};
```

You access elements of the array using its name followed by the index of the element in brackets. Index numbering starts at 0 for the first element. For example, `name[2]` references the third element of the names array declared previously and has a value of Jane.

Classes

Classes are used extensively in object-oriented programming languages. Most of this book is devoted to their creation and use. At this point, it suffices to say that classes define a complex data type for an object. They contain information about how an object should behave, including its name, methods, properties, and events. The .NET Framework contains many predefined classes with which you can work. You can also create your own class type definitions. A variable defined as a class type contains an address pointer to the memory location of the object. The following code declares an object instance of the `StringBuilder` class defined in the .NET Framework:

```
StringBuilder sb = new StringBuilder();
```

Looking at Literals, Constants, and Enumerations

Although the values of variables change during program execution, literals and constants contain items of data that do not change.

Literals

Literals are fixed values implicitly assigned a data type and are often used to initialize variables. The following code uses a literal to add the value of 2 to an integer value:

```
Count = Count + 2;
```

By inspecting the literal, the compiler assigns a data type to the literal. Numeric literals without decimal values are assigned the integer data type; those with a decimal value are assigned as double data type. The keywords `true` and `false` are assigned the Boolean data type. If the literal is contained in quotes, it is assigned as a string data type. In the following line of code, the two string literals are combined and assigned to a string variable:

```
FullName = "Bob" + "Smith";
```

It's possible to override the default data type assignment of the literal by appending a type character to the literal. For example, a value of `12.25` will be assigned the double data type but a value of `12.25f` will cause the compiler to assign it a single data type.

Constants

Many times you have to use the same constant value repeatedly in your code. For example, a series of geometric calculations may need to use the value of pi. Instead of repeating the literal 3.14 in your code, you can make your code more readable and maintainable by using a declared constant. You declare a constant using the `const` keyword followed by the data type and the constant name:

```
const Single pi = 3.14159265358979323846;
```

The constant is assigned a value when it is declared and this value can't be altered or reassigned.

Enumerations

You often need to assign the value of a variable to one of several related predefined constants. In these instances, you can create an enumeration type to group the values. Enumerations associate a set of integer constants to names that can be used in code. For example, the following code creates an enum type of `Manager` used to define three related manager constants with names of `DeptManager`, `GeneralManager`, and `AssistantManager` with values of 0, 1, and 2, respectively:

```
enum Manager
{
    DeptManager,
    GeneralManager,
    AssistantManager,
}
```

A variable of the enum type can be declared and set to one of the enum constants.

```
Manager managerLevel = Manager.DeptManager;
```

■ **Note** The .NET Framework provides a variety of intrinsic constants and enumerations designed to make your coding more intuitive and readable. For example, the `StringAlignment` enumeration specifies the alignment of a text string relative to its layout rectangle.

Exploring Variable Scope

Two important aspects of a variable are its scope and lifetime. The scope of a variable refers to how the variable can be accessed from other code. The lifetime of a variable is the period of time when the variable is valid and available for use. A variable's scope and lifetime are determined by where it is declared and the access modifier used to declare it.

Block-Level Scope

A code block is a set of grouped code statements. Examples of code blocks include code organized in `if-else`, `do-loop`, or `for-next` statements. Block-level scope is the narrowest scope a variable can have. A variable declared within a block of code is available only within the block in which it is declared. In the following code, the variable `blockCount` can only be accessed from inside the `if` block. Any attempt to access the variable outside the block will generate a compiler error.

```
if (icount > 10)
{
    int blockCount;
    blockCount = icount;
}
```

Procedure Scope

Procedures are blocks of code that can be called and executed from other code. There are two types of procedures supported in C#: method and property. Variables declared outside of a code block but within a procedure have procedure-level scope. Variables with procedure scope can be accessed by code within the same procedure. In the following code, the counter `iCount` is declared with procedure scope and can be referenced from anywhere within the procedure block of the `Counter` method:

```
void Counter()
{
    int iCount = 0;
    do
    {
        iCount = iCount + 2;
    }
    while (iCount < 10);
}
```

The lifetime of a procedure scope variable is limited to the duration of the execution of the procedure.

Module Scope

Variables with module scope are available to any code within the class or structure. To have module scope, the variable is declared in the general declaration section (outside of any procedure blocks) of the class or structure. To limit the accessibility to the module where it is declared, you use the `private` access modifier keyword. In the following code, the `iCount` variable can be accessed by both procedures defined in the class:

```
public class Class1
{
    private int _iCount;
    public void IncrementCount()
    {
        int iCount = 0;
        do
        {
            iCount = iCount + 2;
        }
        while (iCount < 10);
    }
    public void ReadCount()
    {
        Console.WriteLine(_iCount.ToString());
    }
}
```

The lifetime of the variable declared with module scope is the same as the lifetime of the object instance of the class or structure in which it is declared.

■ **Note** For a further discussion of scope see Chapter 6.

Understanding Data Type Conversion

During program execution, there are many times when a value must be converted from one data type to another. The process of converting between data types is referred to as *casting* or *conversion*.

Implicit Conversion

The C# compiler will perform some data type conversions for you automatically. For numeric types, an implicit conversion can be made when the value to be stored can fit into the variable without being truncated or rounded off. For example, in the following code, an integer data type is implicitly converted to a long data type:

```
int i1 = 373737373;
long l1 = i1;
l1 *= l1;
```

Explicit Conversion

Explicit conversion is referred to as *casting*. To perform a cast, you specify the type that you are casting to in parentheses in front of the value or variable to be converted. The following code uses a cast to explicitly convert the double type `n1` to an integer type:

```
double n1 = 3.73737373;
int i1 = (int)n1;
```

Widening and Narrowing Conversions

Widening conversions occur when the data type being converted to can accommodate all the possible values contained in the original data type. For example, an integer data type can be converted to a double data type without any data loss or overflow. Data loss occurs when the number gets truncated. For example, 2.54 gets truncated to 2 if it is converted to an integer data type. Overflow occurs when a number is too large to fit in the new data type. For example, if the number 50000 is converted to a short data type, the maximum capacity of the short data type is exceeded, causing the overflow error. Narrowing conversions, on the other hand, occur when the data type being converted to can't accommodate all the values that can be contained in the original data type. For example, when the value of a double data type is converted to a short data type, any decimal values contained in the original value will be lost. In addition, if the original value is more than the limit of the short data type, a runtime exception will occur. You should be particularly careful to trap for these situations when implementing narrowing conversions in your code.

Working with Operators

An operator is a code symbol that tells the compiler to perform an operation on a value. The operation can be arithmetic, comparative, or logical.

Arithmetic Operators

Arithmetic operators perform mathematical manipulation to numeric types. Table A-3 lists the commonly used arithmetic operators available in C#.

Table A-3. Arithmetic Operators

Operator	Description
=	Assignment
*	Multiplication
/	Division
+	Addition
-	Subtraction

The following code increments the value of an integer data type by the number one:

```
Count = Count + 1;
```

C# also supports shorthand assignment operators that combine the assignment with the operation. The following code is equivalent to the previous code:

```
Count += 1;
```

If you are going to increment by one, you can also use the shorthand assignment ++. The following code is equivalent to the previous code:

```
Count ++;
```

Comparison Operators

A comparison operator compares two values and returns a Boolean value of true or false. Table A-4 lists the common comparison operators used in C#.

Table A-4. Comparison Operators

Operator	Description
<	Less than
<=	Less than or equal to
>	Greater than
>=	Greater than or equal to
==	Equal to
!=	Not equal to

You use comparison operators in condition statements to decide when to execute a block of code. The following `if` block checks to see if the number of invalid login attempts is greater than three before displaying a message:

```
if (_loginAttempts > 3)
{
    MessageBox.Show("Invalid login.");
}
```

Logical Operators

Logical operators combine the results of conditional operators. The three most commonly used logical operators are the AND, OR, and NOT operators. The AND operator (`&&`) combines two expressions and returns `true` if both expressions are true. The OR operator (`||`) combines two expressions and returns `true` if either one is true. The NOT operator (`!`) switches the result of the comparison: a value of `true` returns `false` and a value of `false` returns `true`. The following code checks to see whether the logged-in user is a department manager or assistant manager before running a method:

```
if (currentUserLevel == Manager.AssistantManager ||
    currentUserLevel == Manager.DeptManager)
{
    ReadLog();
}
```

Ternary Operator

The ternary operator evaluates a Boolean expression and returns one of two values depending on the result of the expression. The following shows the syntax of the ternary operator:

```
condition ? first_expression : second_expression;
```

If the condition evaluates to `true`, the result of the first expression is returned. If the condition evaluates to `false`, the result of the second expression is returned. The following code checks to see if the value of `x` is zero. If it is, it returns `0`; if not, it divides `y` by `x` and returns the result.

```
return x == 0.0 ? 0 : y/x;
```

Introducing Decision Structures

Decision structures allow conditional execution of code blocks depending on the evaluation of a condition statement. The `if` statement evaluates a Boolean expression and executes the code block if the result is `true`. The `switch` statement checks the same expression for several different values and conditionally executes a code block depending on the results.

If Statements

To execute a code block if a condition is true, use the following structure:

```
if (condition1)
{
    //code
}
```

To execute a code block if a condition is true and an alternate code block if it is false, add an else block.

```
if (condition1)
{
    //code
}
else
{
    //code
}
```

To test additional conditions if the first evaluates to false, add an else-if block:

```
if (condition1)
{
    //code
}
else if (condition2)
{
    //code
}
else
{
    //code
}
```

An if statement can have multiple else-if blocks. If a condition evaluates to true, the corresponding code statements are executed, after which execution jumps to the end of the statements. If a condition evaluates to false, the next else-if condition is checked. The else block is optional, but if included, it must be the last. The else block has no condition check and executes only if all other condition checks have evaluated to false. The following code demonstrates using the if statement to evaluate a series of conditions. It checks a performance rating to determine what bonus to use and includes a check to see if the employee is a manager to determine the minimum bonus.

```
if (performance ==1)
{
    bonus = salary * 0.1;
}
else if (performance == 2)
{
    bonus = salary * 0.08;
}
```



```

else if (employeeLevel == Manager.DeptManager)
{
    bonus = salary * 0.05;
}
else
{
    bonus = salary * 0.03;
}

```

Switch Statements

Although the switch statement is similar to the if-else statement, it's used to test a single expression for a series of values. The structure of the switch statement is as follows:

```

switch (expression)
{
case 1:
    Console.WriteLine("Case 1");
    break;
case 2:
    Console.WriteLine("Case 2");
    break;
default:
    Console.WriteLine("Default case");
    break;
}

```

A switch statement can have multiple case blocks. If the test expression value matches the case expression, the code statements in the case block execute. After the case block executes, you need a break statement to bypass the rest of the case statements. If the test expression doesn't match the case expression, execution jumps to the next case block. The default block doesn't have an expression. It executes if no other case blocks are executed. The default block is optional, but if used, it must be last. The following example uses a switch to evaluate a performance rating to set the appropriate bonus rate:

```

switch(performance)
{
    case 1:
        bonus = salary * 0.1;
        break;
    case 2:
        bonus = salary * 0.08;
        break;
    case 3:
        bonus = salary * 0.03;
        break;
    default:
        bonus = salary * 0.01;
        break;
}

```

Using Loop Structures

Looping structures repeat a block of code until a condition is met. C# supports the following looping structures.

While Statement

The while statement repeats the execution of code while a Boolean expression remains true. The expression gets evaluated at the beginning of the loop. The following code executes until a valid login variable evaluates to true:

```
while (validLogin == false)
{
    //code statements...
}
```

Do-While Statement

The do-while loop is similar to the while loop except the expression is evaluated at the end of the loop. The following code will loop until the maximum login attempts are met:

```
do
{
    //code statements...
}
while (iCount < maxLoginAttempts);
```

For Statement

A for statement loops through a code block a specific number of times based on the value stored in a counter. For statements are a better choice when you know the number of times a loop needs to execute at design time. In the parentheses that follow a for statement, you initialize a counter, define the evaluation expression, and define the counter increment amount.

```
for (int i = 0; i < 10; i++)
{
    //Code statements...
}
```

For Each Statement

The for-each statement loops through code for each item in a collection. A *collection* is a group of ordered items; for example, the controls placed on a Windows Form are organized into a Controls collection. To use the for-each statement, you first declare a variable of the type of items contained in the collection. This variable is set to the current item in the collection. The following for-each statement loops through the employees in an employee list collection:

```
foreach (Employee e in employeeList)
{
    //Code statements
}
```

If you need to conditionally exit a looping code block, you can use the `break` statement. The following code shows breaking out of the `for-each` loop:

```
foreach (Employee e in employeeList)
{
    //Code statements
    if (e.Name == "Bob")
    {
        break;
    }
}
```

Introducing Methods

Methods are blocks of code that can be called and executed from other code. Breaking an application up into discrete logical blocks of code greatly enhances code maintenance and reuse. C# supports methods that return values and methods that do not. When you declare a method, you specify an access modifier, a return type, and a name for the method. The following code declares a method with no return type (designated by the keyword `void`) used to record logins to the event log:

```
public void RecordLogin(string userName)
{
    EventLog appLog = new EventLog();
    appLog.Source = "OSO App";
    appLog.WriteEntry(userName + " has logged in.");
}
```

You can declare methods with a parameter list that defines arguments that must be passed to the method when it is called. The following code defines a method that encapsulates the assignment of a bonus rate. The calling code passes an integer type value to the method and receives a double type value back.

```
public double GetBonusRate(int performanceRating)
{
    double bonusRate;
    switch (performanceRating)
    {
        case 1:
            bonusRate = 0.1;
            break;
        case 2:
            bonusRate = 0.08;
            break;
        case 3:
            bonusRate = 0.03;
            break;
        default:
            bonusRate = 0.01;
            break;
    }
    return bonusRate;
}
```

The following code demonstrates how the method is called:

```
double salary;  
int performance = 0;  
double bonus = 0;  
// Get salary and performance data from data base...  
bonus = GetBonusRate(performance) * salary;
```

If the access modifier of the method is private, it is only accessible from code within the same class. If the method needs to be accessed by code in other classes, then the `public` access modifier is used.

The information in this appendix was designed to get you up to speed on some fundamental programming concepts. You should now be more familiar with C# syntax, and ready to expand and build on these concepts by working through the rest of this book.



Exception Handling in C#

The topics discussed here extend the discussion of exception handling found in Chapter 8, so this discussion assumes that you have first thoroughly reviewed Chapter 8. The purpose of this appendix is to review Microsoft's recommendations for exception management and present a few of the exception classes provided by the .NET Framework.

Managing Exceptions

Exceptions are generated when the implicit assumptions made by your programming logic are violated. For example, when a program attempts to connect to a database, it assumes that the database server is up and running on the network. If the server can't be located, an exception is generated. It's important that your application gracefully handles any exceptions that may occur. If an exception is not handled, your application will terminate.

You should incorporate a systematic exception handling process in your methods. To facilitate this process, the .NET Framework makes use of structured exception handling through the `try`, `catch`, and `finally` code blocks. The first step is to detect any exceptions that may be thrown as your code executes. To detect any exceptions thrown, place the code within the `try` block. When an exception is thrown in the `try` block, execution transfers to the `catch` block. You can use more than one `catch` block to filter for specific types of exceptions that may be thrown. The `finally` block performs any cleanup code that you wish to execute. The code in the `finally` block executes regardless of whether an exception is thrown. The following code demonstrates reading a list of names from a file using the appropriate exception handling structure:

```
public ArrayList GetNames(string file)
{
    try
    {
        using( StreamReader stream = new StreamReader())
        {
            ArrayList names = new ArrayList();

            while (stream.Peek() > -1)
            {
                names.Add(stream.ReadLine());
            }
        }
    }
}
```

```

    catch (FileNotFoundException e)
    {
        //Could not find file
    }
    catch (FileLoadException e)
    {
        //Could not open file
    }
    catch (Exception e)
    {
        //Some kind of error occurred. Report error.
    }
    finally
    {
        stream.Close();
    }
    return names;
}

```

After an exception is caught, the next step in the process is to determine how to respond to it. You basically have two options: either recover from the exception or pass the exception to the calling procedure. The following code demonstrates how to recover from a `DivideByZeroException` by setting the result to zero:

```

...
try
{
    Z = x / y;
}
catch (DivideByZeroException e)
{
    Z = 0;
}
...

```

An exception is passed to the calling procedure using the `throw` statement. The following code demonstrates throwing an exception to the calling procedure where it can be caught and handled:

```

catch (FileNotFoundException e)
{
    throw e;
}

```

As exceptions are thrown up the calling chain, the relevance of the original exception can become less obvious. To maintain relevance, you can wrap the exception in a new exception containing additional information that adds relevancy to the exception. The following code shows how to wrap a caught exception in a new one and then pass it up the calling chain:

```

catch (FileLoadException e)
{
    throw new Exception("GetNames function could not open file", e);
}

```

You preserve the original exception by using the `InnerException` property of the `Exception` class.

```
catch (FileLoadException e)
{
    throw new Exception("Could not open file. " + e.InnerException, e);
}
```

Implementing this exception management policy consistently throughout the various methods in your application will greatly enhance your ability to build highly maintainable, flexible, and successful applications.

Using the .NET Framework Exception Classes

The Common Language Runtime (CLR) has a set of built-in exception classes. The CLR will throw an object instance of the appropriate exception type if an error occurs while executing code instructions. All .NET Framework exception classes derive from the `SystemException` class, which in turn derives from the `Exception` class. These base classes provide functionality needed by all exception classes.

Each namespace in the framework contains a set of exception classes that derive from the `SystemException` class. These exception classes handle common exceptions that may occur while implementing the functionality contained in the namespace. To implement robust exception handling, it's important for you to be familiar with the exception classes provided by the various namespaces. For example, Table B-1 summarizes the exception classes in the `System.IO` namespace.

Table B-1. *Exception Classes in the System.IO Namespace*

Exception	Description
<code>IOException</code>	The base class for exceptions thrown while accessing information using streams, files, and directories.
<code>DirectoryNotFoundException</code>	Thrown when part of a file or directory can't be found.
<code>EndOfStreamException</code>	Thrown when reading is attempted past the end of a stream.
<code>FileLoadException</code>	Thrown when a file is found but can't be loaded.
<code>FileNotFoundException</code>	Thrown when an attempt to access a file that does not exist on disk fails.
<code>PathTooLongException</code>	Thrown when a path or filename is longer than the system-defined maximum length.

Every exception class in the .NET Framework contains the properties listed in Table B-2. These properties help identify where the exception occurred and its cause.

Table B-2. *Exception Class Properties*

Property	Description
<code>Message</code>	Gets a message that describes the current exception.
<code>Source</code>	Gets or sets the name of the application or the object that causes the error.
<code>StackTrace</code>	Gets a string representation of the frames on the call stack at the time the current exception was thrown.
<code>InnerException</code>	Gets the exception instance that caused the current exception.
<code>HelpLink</code>	Gets or sets a link to the help file associated with this exception.

In addition, the `ToString` method of the exception classes provides summary information about the current exception. It combines the name of the class that threw the current exception, the message, the result of calling the `ToString` method of the inner exception, and the stack trace information of the current exception.

You will find that the exception classes in the .NET Framework provide you with the capabilities to handle most exceptions that may occur in your applications. In cases where you may need to implement custom error handling, you can create your own exception classes. These classes need to inherit from `System.ApplicationException`, which in turn inherits from `System.Exception`. The topic of creating custom exception classes is an advanced one and thus beyond the scope of this text; for more information, consult the .NET Framework documentation at <http://msdn.microsoft.com/en-us/library/>.

The Importance of Using

As you write applications, there are many times when you have to access unmanaged resources from the C# managed types. For instance you may have to read or write to a text file or open a connection to a database. It is important that you release these resources as soon as you are finished using them. If you wait around for the garbage collector to clean up the resources, they can be inaccessible for long periods of time, blocking other programs from gaining access to the resources. One scenario in which this can occur is connecting to a database. There are a limited number of connections that can be made. If these connections are not managed correctly, then performance will suffer greatly.

In order to safely dispose of any unmanaged resources, classes that access these resources must implement the `IDisposable` interface and a `Dispose` method. Users of the class are guaranteed that calling the `Dispose` method will properly clean up the resources. For example, when you are done with a database connection, you should call its `Dispose` method in a `finally` block as shown in the following code.

```
SqlConnection pubConnection = new SqlConnection();
string connString;
try
{
    connString = "Data Source=drcsv01;Initial Catalog=pubs;Integrated Security=True";
    pubConnection.ConnectionString = connString;
    pubConnection.Open();
    //work with data
}
catch (SqlException ex)
{
    throw ex;
}
finally
{
    pubConnection.Dispose();
}
```

This pattern is so important for writing reliable code that Microsoft has implemented a convenient syntax that ensures the correct use of `IDisposable` objects. When you use the `using` statement, you ensure that the `Dispose` method is properly called when the object goes out of scope or if an exception occurs while you are calling methods of the object. The following code is equivalent to the code shown previously for creating and working with a `SqlConnection`.


```
string connString = "Data Source=drcsrv01;Initial Catalog=pubs;Integrated Security=True";
using(SqlConnection pubConnection = new SqlConnection())
{
    pubConnection.ConnectionString = connString;
    pubConnection.Open();
    //work with data
}
```

You should get in the habit of employing the using pattern for any class that implements the `IDisposable` interfaces. The following code uses this pattern to open and read from a file using the `StreamReader` class.

```
string path = @"c:\temp\MyTest.txt";
using (StreamReader sr = File.OpenText(path))
{
    string s = "";
    while ((s = sr.ReadLine()) != null)
    {
        // work with text
    }
}
```

This discussion has provided you with a supplement to Chapter 8, covering methods for effectively handling exceptions and using the .NET Framework exception classes. Please see Chapter 8 for the basic discussion of this topic.



Installing the Required Software

I have included many learning activities throughout this book. In order to get the most out of the topics I discuss, you should complete these activities. This is where the theory becomes concrete. It is my hope that you will take these activities seriously and work through them thoroughly, and even repeatedly.

The UML modeling activities in Part 1 are meant for someone using UMLet. I chose this program because it is a good diagramming tool to learn on. It enables you to create UML diagrams without adding a lot of advanced features. UMLet is a free open source tool and can be downloaded from <http://www.umlet.com>. But you don't need a tool to complete these activities; a paper and pencil will work just fine.

The activities in Part 2 require Visual Studio 2012 with C# installed. You can use either the free version, Visual Studio 2012 Express, or a trial version of Visual Studio 2012 Professional. These versions are available at <http://msdn.microsoft.com/en-us/vstudio/>. I encourage you to install the help files and make abundant use of them while you're completing the activities.

The activities in Part 3 require Microsoft SQL Server 2008 R2 or SQL Server 2012. You can use either the free version SQL Server Express or a trial version of SQL Server available at <http://msdn.microsoft.com/en-us/sqlserver/>. When you install SQL Server, be sure you add yourself as an administrator.

Installing the Sample Databases

The scripts to install the sample database used in this book are available for download at <http://www.apress.com/9781430249351> (scroll down and click on the Source Code tab). In order to install the scripts, follow these steps:

1. Open a command prompt window.
2. From the command prompt, use the `cd` command to navigate to the folder containing the sample database scripts.

```
cd c:\SampleDatabases
```

3. Run `SQLCmd.exe` specifying `inst0SODB.sql` as the input file.
4. To install the database on a default instance, use

```
SQLCmd.exe -E -i inst0SODB.sql
```

5. To install the database on a named instance, use

```
SQLCmd.exe -E -S ComputerName\InstanceName -i inst0SODB.sql
```

6. Repeat the procedure for the `instpubs.sql` and `instnwnd.sql` files.

■ **Note** You can also use SQL Server Management Studio to create the databases.

Verifying the Database Installs

To verify the database installs:

1. Start Visual Studio. If you don't see the Server Explorer window shown in Figure C-1, open it by choosing Server Explore on the View menu.




Figure C-1. The Server Explorer window

2. In the Server Explorer window, right-click the Data Connections node and select Add Connection. In the Add Connections dialog box shown in Figure C-2, fill in the name of your server, select the Northwind database, and click OK.

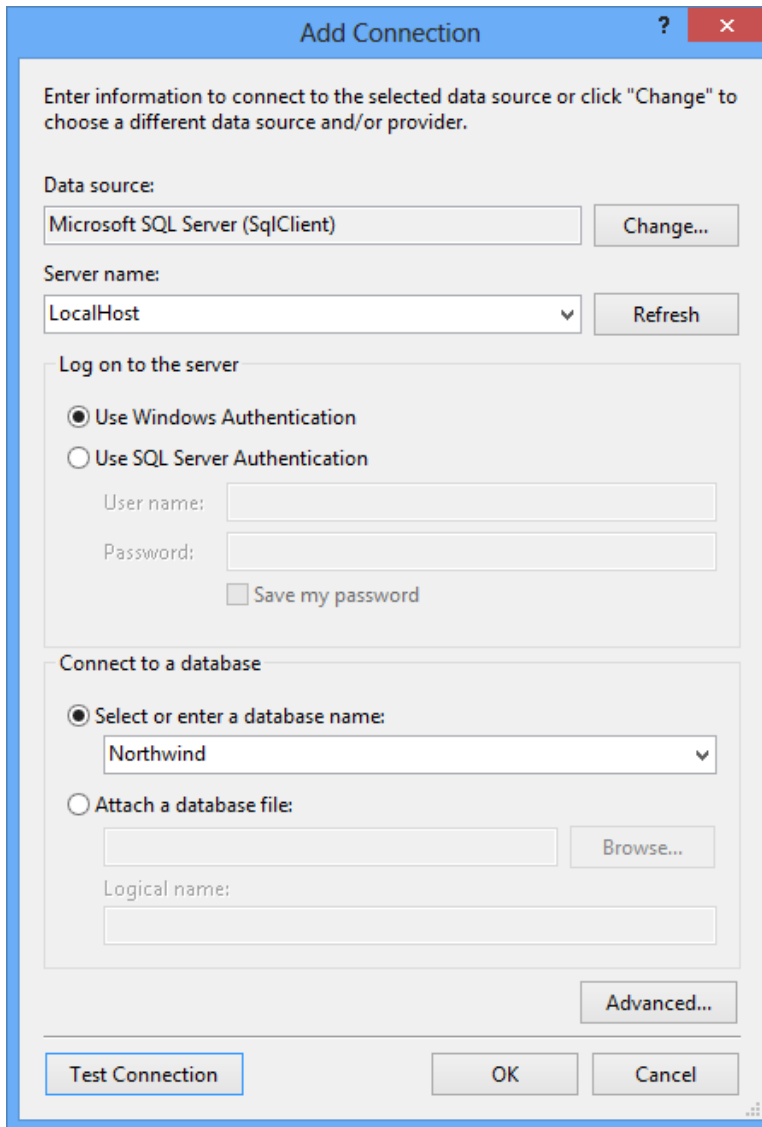


Figure C-2. The Add Connections dialog box

3. Expand the Northwind database node and the Tables node in the Database Explorer window, as shown in Figure C-3.

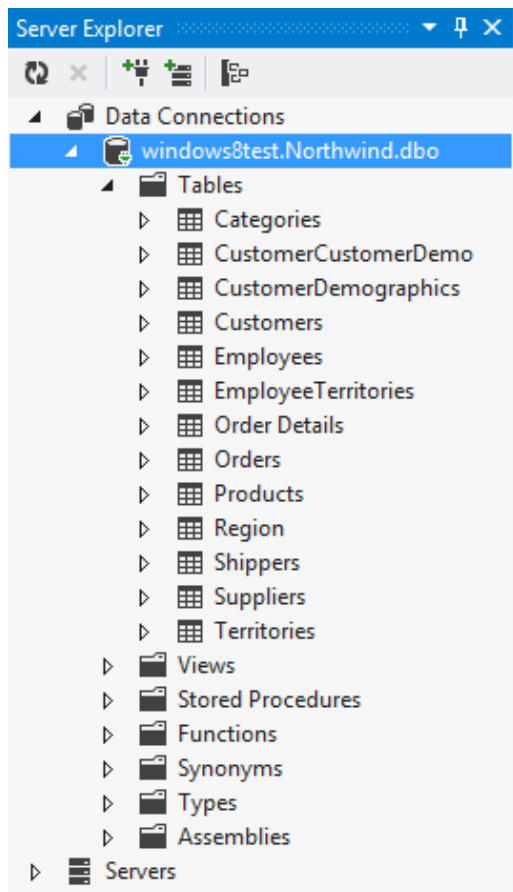


Figure C-3. Expanding the Tables node

■ **Note** If you are using SQLExpress you need to use the server name and the instance name (LocalHost\SQLExpress) to connect.

4. Right-click the Suppliers table node and select Show Table Data. The Suppliers table data should display as shown in Figure C-4.

	SupplierID	CompanyName	ContactName	ContactTitle	Address	City
▶	1	Exotic Liquids	Charlotte Cooper	Purchasing Ma...	49 Gilbert St.	London
	2	New Orleans C...	Shelley Burke	Order Administ...	P.O. Box 78934	New Orleans
	3	Grandma Kelly' ...	Regina Murphy	Sales Represent...	707 Oxford Rd.	Ann Arbor
	4	Tokyo Traders	Yoshi Nagase	Marketing Man...	9-8 Sekimai Mu...	Tokyo
	5	Cooperativa de ...	Antonio del Val...	Export Adminis...	Calle del Rosal 4	Oviedo
	6	Mayumi's	Mayumi Ohno	Marketing Repr...	92 Setsuko Chu...	Osaka
	7	Pavlova, Ltd.	Ian Devling	Marketing Man...	74 Rose St. Mo...	Melbourne
	8	Specialty Biscui...	Peter Wilson	Sales Represent...	29 King's Way	Manchester
	9	PB Knäckebröd ...	Lars Peterson	Sales Agent	Kaloadagatan 13	Göteborg
	10	Refrescos Amer...	Carlos Diaz	Marketing Man...	Av. das Americ...	Sao Paulo
	11	Heli Süßwaren ...	Petra Winkler	Sales Manager	Tiergartenstraß...	Berlin
	12	Plutzer Lebens...	Martin Bein	International M...	Bogenallee 51	Frankfurt
	13	Nord-Ost-Fisch...	Sven Petersen	Coordinator Fo...	Frahmredder 11...	Cuxhaven
	14	Formaggi Forti...	Elio Rossi	Sales Represent...	Viale Dante, 75	Ravenna
	15	Norske Meierier	Beate Vileid	Marketing Man...	Hatlevegen 5	Sandvika
	16	Bigfoot Breweries	Cheryl Saylor	Regional Accou...	3400 - 8th Aven...	Bend
	17	Svensk Sjöföda ...	Michael Björn	Sales Represent...	Brovallavägen 2...	Stockholm
	18	Aux joyeux eccl...	Guyène Nodier	Sales Manager	203, Rue des Fr...	Paris

Figure C-4. Viewing the table data

5. Repeat these steps to test the pubs and the OfficeSupply databases. After testing, exit Visual Studio.

Index

■ A

Abstraction, 3
ADO.NET, 161, 322
 interoperability, 162
 scalability, 162
Aggregation, 5
AND operator (&&), 334
Arithmetic operators, 332–333
Arrays, 328

■ B

Block-level scope, 330
Boolean data types, 327
Byte data type, 326

■ C

Casting, 332
Character data types, 326
C#, history of, 5
Classes, 329
 class methods, 85
 constructors and overloading
 methods, 89–90
 class method, 94
 creation steps, 92
 employee class constructors, testing, 93
 update method, testing, 95
 definition, 84
 employee class, 86
 definition, 87
 testing, 88
 encapsulation, 85
 instance variables, 84
 and objects, 83
 human resource application, 83
 properties and methods, 83

Class hierarchies, 97
 base class method
 from derived class, 106
 hiding, 107
 derived class method from base class, 105
 inheritance and polymorphism (*see* Inheritance;
 Polymorphism)
 interfaces, 111
 overloading methods, 106
 Account class, 107
 base modifier, 110
 using base qualifier, 109
 Withdraw methods, testing, 108
 overriding methods
 abstract base class, 104
 Deposit method, 104
 override keyword, 104
 virtual keyword, 104
Class structure, 7
Collections, 337
 arrays and array lists, 144
 Array Lists, 152
 command line arguments, 149
 console output, 150
 creation, 149
 DOS command, 144
 foreach loop, 146
 index values, 145
 multidimensional arrays, 151
 numeric types, 145
 one-dimensional array, 147
 static methods, 145
 two-dimensional array, 145, 147
 generic collections, 153
 extend, 156
 IComparer interface, 154
 implementation, 155
 unsorted and sorted by date, 157
 weakly typed, 153

Collections (*cont.*)
 .NET Framework, 143
 class interfaces, 144
 collection classes, 144
 System.Collections namespace, 143
 stacks and queues, 157
 GetLastMove method, 157
 implementation, 158
 peek method, 157
 pop method, 157
 RecordMove method, 157
 Common Language Runtime (CLR), 343
 Comparison operators, 333–334
 Component-based development, 323
 Composite data type
 arrays, 328
 classes, 329
 structures, 327–328
 Computer-Aided Software Engineering
 (CASE) tool, 13
 Conceptual schema definition language
 (CSDL), 182
 Constants, 329
 Conversion of data type. *See* Data type conversion

D

Data access layer
 ADO.NET, 161
 Command object, 164
 CommandText method, 164
 ExecuteReader method, 164
 ExecuteScalar method, 164
 using stored procedure, 165
 Connection object, 162
 data providers, 162
 Data Tables and DataSets, 172
 DataRelation object, 175
 editing and updating data, 178
 establish relationships, 179
 Fill method, 174
 Load method, 173
 from SQL Server database, 176
 System.Data namespace, 172
 Update Command, 175
 Entity Framework (EF), 181
 CSDL, 182
 disadvantages, 181
 Entity Data Model, 186
 Entity Data Model Wizard, 181
 LINQ to EF, 185, 189
 mapping schema, 181
 model designer, 183
 MSL, 183
 navigation properties, 185
 ObjectContext class, 184

SaveChanges method, 186
 SSDL, 182
 retrieve data
 DataAdapter, 167
 DataReader object, 166
 execute stored procedure,
 Command object, 171
 from SQL server database, 168–169
 using DataReader object, 170
 Data binding
 DataGrid to DataTable, 209
 loading DataGrid, 211
 OneWay binding, 207
 TwoWay binding, 207
 updating data, 213
 using DataContext, 207
 Window Layout, 210
 Data type(s), 325
 Boolean, 327
 character, 326
 composite, 327
 arrays, 328
 classes, 329
 structures, 327–328
 date, 327
 integral
 byte, 326
 integer, 326
 long, 326
 short, 326
 non-integral
 decimal, 326
 double, 326
 single, 326
 nullable, 327
 object, 327
 Data type conversion, 332
 explicit, 332
 implicit, 332
 narrowing, 332
 widening, 332
 Date data type, 327
 Decimal data type, 326
 Decision structures, 334
 if statement, 335–336
 switch statement, 336
 Dialog boxes, WPF, 204
 custom dialog box, 206
 MessageBox
 complex MessageBox, 206
 Show method, 205
 DirectoryNotFoundException, 343
 Distributed Component Object Model (DCOM)
 technology, 63
 Double data type, 326
 do-while statement, 337

E

- else-if blocks, 335
- Encapsulation, 4
- EndOfStreamException, 343
- Entity Framework, 322
- Enumerations, 330
- Exception classes, .NET framework
 - DirectoryNotFoundException, 343
 - EndOfStreamException, 343
 - FileLoadException, 343
 - FileNotFoundException, 343
 - IOException, 343
 - PathTooLongException, 343
 - properties, 343
- Exception handling in C#
 - finally block, 341
 - InnerException property, 343
 - .NET framework exception classes, 343
 - DirectoryNotFoundException, 343
 - EndOfStreamException, 343
 - FileLoadException, 343
 - FileNotFoundException, 343
 - IOException, 343
 - PathTooLongException, 343
 - properties, 343
 - recovery from DivideByZeroException, 342
 - throw statement, 342
 - try-catch block, 341
 - unmanaged resource access, 344–345
- Explicit type conversion, 332
- Extensible Application Markup Language (XAML)
 - Button control, 194
 - Grid.Row and Grid.Column attribute, 194
 - property element, 194

F

- FileLoadException, 343
- FileNotFoundException, 343
- for-each statement, 337
- for statement, 337

G

- Global Assembly Cache (GAC), 64

H

- Help system, 323

I, J, K

- if statement, 335–336
- Implicit type conversion, 332

Inheritance, 5, 97

- abstract class, 99
- access modifiers, 99
- Account base class, 97
- derived class, 98
- purpose of, 97
- sealed/final class, 99
- using base class and derived class
 - abstract class, accessibility of, 103
 - creation, 100
 - GetBalance method, protect, 102
 - testing, 101
 - Withdraw method, testing, 103
- InnerException property, 343
- Integer data type, 326. *See also* Data types, integral
- IOException, 343

L

- Layout control, WPF
 - Canvas control, 196
 - DockPanel, 196
 - fixed positioning, 195
 - Grid control, 195
 - StackPanel, 196
 - WrapPanel, 196
- Literals, 329
- Logical operators, 334
- Long data type, 326
- Loop structures, 337
 - do-while statement, 337
 - for-each statement, 337
 - for statement, 337
 - while statement, 337

M

- Mapping specification language (MSL), 183
- Methods, 338–339
- Model Binding
 - ASP.NET GridView control (displaying data)
 - data objects, 247–248
 - EmployeeGridViewPage.aspx.cs, 249
 - SelectMethod, 248
 - steps, 247
 - ASP.NET GridView control (updating data)
 - EmployeeGridViewPage.aspx.cs, 249
 - Page_Load event handler method, 250
 - UpdateMethod attribute, 249
 - ASP.NET Repeater control (displaying data)
 - AutherListPage.aspx, 247
 - Eval method, 247
 - select statement, 246
 - SqlDataSource control, 247
 - SqlDataSource Tasks pane, 245

Model Binding (*cont.*)

SQL server data source, 245–246

steps, 245

Module scope, 331

Multidimensional arrays, 328

■ N

Narrowing type conversion, 332

.NET Framework, 59, 322

assemblies and manifests, 64

compile and executing managed code, 65

components, 61

application services, 63

base class library, 62

CLR, 61

data classes, 62

web applications, 63

windows applications, 62

Windows Store app, 63

goals

application deployment, 60

extensibility, 60

industry standards and practices, 59

memory management, 61

security models, 61

unified programming models, 60

system namespaces and assemblies, 64

New operator, 328

Non-integral data types. *See* Data types,

non-integral

NOT operator (!), 334

Nullable data type, 327

■ O

Object(s), 3

Object collaboration, 119. *See also* Object communication

Object communication

asynchronous messaging, 136

async method, 137

call method asynchronously, 140

call method synchronously, 138

TAP, 137

delegation, 121

event-driven programs, 121

event handler methods, 123

event messages

add and raise, 124

handle multiple events, 126

receive events, 125

events, 122

messaging, 119

method signatures, 120

passing parameters, 120

static properties and methods, 131

creation, 133

filter exceptions, 136

Logger Form and Control Properties, 133

structured exception handler, 135

System.String class, Compare method, 132

TaxRate property, 132

UserLog, 132

structured exception handling

benefits of, 128

finally block, 129

nesting exception handling, 131

throwing exceptions, 130

try-catch block, 128

subscription-based messaging, 121

Object data type, 327

Object interaction

activity diagrams

branching condition, 38

decision points and guard conditions, 34

generic diagram, 34

objects and activities, 36

ownership, 35

parallel processing, 35

partitions, 37

using UMLet, 36

GUI design, 39

activity diagrams, 39

application prototyping, 41

interface flow diagram, 41

interface prototyping, 40

message branching, 29

message constraints, 29

message iteration, 28

message types, 27

recursive messages, 28

scenarios, 25

sequence diagrams, 26

Object-oriented analysis and design, 322

Object-oriented programming (OOP)

benefits, 3

characteristics

abstraction, 3

aggregation, 5

encapsulation, 4

inheritance, 4

objects, 3

polymorphism, 4

definition, 1

history of, 2

software design, 7

UML, 8

uses, 2

- Office-supply ordering (OSO) application, 43
 - business logic layer
 - Employee.cs file, 304-305
 - Order class, 307
 - OrderItem class, 305-306
 - ProductCatalog class, 305
 - class model
 - activity diagram, 52
 - behavior model, 52
 - class associations, 51
 - class attributes, 49
 - class identification, 48
 - preliminary class diagram, 49
 - sequence diagram, 53
 - User Interface model design, 55
 - data access layer
 - class diagram, 297
 - DALEmployee class, 301
 - DALOrder, 303
 - DALProductCatalog class, 302
 - DALUtility, 300
 - database diagram, 297-298
 - OfficeSupplyBLL class
 - library, 298-299
 - PlaceOrder method, 303
 - reference, 299
 - static class, 299
 - design pitfalls, 58
 - design process
 - conceptual design, 295
 - logical design, 295
 - logical tiers, 296
 - physical design, 295
 - physical tiers, 296-297
 - 3-tiered application, 296
 - SRS creation, 44
 - UI
 - addButton_Click event, 317
 - App.config file, 309
 - ListView control, 316
 - loginButton_Click, 317
 - LoginDialog form, 313-314
 - MainWindow.xaml.cs, 314-315
 - order form, 309, 310
 - order item dialog, 312
 - OrderItemDialog code, 319
 - OrderItemDialog form, 312-313
 - OrderItemDialog.xaml.cs, 318
 - placeOrderButton_Click event, 318
 - reference (OfficeSupplyBLL project), 308
 - XAML code, 310-311
 - use case development
 - actors, 45
 - includes and extends relationship, 47
 - Login use case, 47

- Purchase request use case, 47
 - with UMLet, 46
 - OLEDB, 62
 - OOP solution
 - class model
 - attributes, 49
 - behavior model, 52
 - class associations, 51
 - classes identification, 48
 - design pitfalls, 58
 - system requirements, 43
 - use cases, 45
 - User Interface model design, 55
 - Operators
 - arithmetic, 332-333
 - comparison, 333-334
 - logical, 334
 - ternary, 334
 - OR operator (||), 334

■ P, Q

- Paid time off (PTO), 119
- PathTooLongException, 343
- Polymorphism, 4, 97, 111
 - using inheritance, 113
 - using interface, 115
- Procedure scope, 331
- Process Movie Rental use case
 - scenarios, 25
 - sequence diagram, 27
- Programming
 - constants, 329
 - data type(s), 325
 - Boolean, 327
 - byte, 326
 - character, 326
 - composite, 327-329
 - date, 327
 - decimal, 326
 - double, 326
 - integer, 326
 - long, 326
 - nullable, 327
 - object, 327
 - short, 326
 - single, 326
 - data type conversion, 332
 - explicit, 332
 - implicit, 332
 - narrowing, 332
 - widening, 332
 - decision structures, 334
 - If statement, 335-336
 - switch statement, 336

Programming (*cont.*)

- enumerations, 330
- literals, 329
- loop structures, 337
 - do-while statement, 337
 - for-each statement, 337
 - for statement, 337
 - while statement, 337
- methods, 338-339
- operators, 332
 - arithmetic, 332-333
 - comparison, 333-334
 - logical, 334
 - ternary, 334
- strong typing, 325
- variables, 325
- variable scope, 330
 - block-level scope, 330
 - module scope, 331
 - procedure scope, 331

■ R

- RESTful web services
 - ASP.NET Web API service, 285
 - clients and servers, 285
 - consume (ASP.NET Web API service), 290
 - POX, 285
 - Web API service, 291
 - asynchronous method, 292
 - book information, 294
 - books.json file, 290
 - Get method, 288
 - Global.asax.cs file, 288
 - HttpClient variable, 292
 - MainPage.xaml file, 291
 - OnNavigatedTo event handler, 293
 - port number, 289
 - project window, 286
 - startup projects, 293
 - steps, 286, 291
 - view (project files), 288-289

■ S

- Scenarios, 25
- Scope of variable. *See* Variable scope
- Short data type, 326
- Shorthand assignment operators, 333
- Simula, 2
- Single data type, 326
- SQLCmd.exe, 347
- SQL Server database
 - free versions, 347
 - Sample database

- installation, 347-348
 - installation verification, 348-349, 351
- Store schema definition language (SSDL), 182
- Strong typing, 325
- Structures, 327-328
- switch statement, 334, 336

■ T

- Task-Based Asynchronous
 - Pattern (TAP), 137
- Ternary operator, 334
- Throw statement, 342
- ToString method, 344

■ U

- UMLet, 13, 347
- Unified Modeling Language (UML)
 - activity diagram, 8
 - advantages, 8
 - class diagram, 8, 16, 20
 - association shape, 21
 - attributes and operations, 16
 - classes and attributes, 20
 - class shape, 21
 - creation of, 21
 - definition, 16
 - Flight class, 17
 - generalization arrow, 22
 - collaboration diagram, 8
 - object relationships, 17
 - association classes, 19
 - class associations, 17
 - depicting aggregations, 18
 - documenting inheritance, 18
 - sequence diagram, 8
 - Software Requirement Specification, 8-9, 12
 - textual and graphical models, 8
 - UMLet, 347
 - use case, 8, 16
 - actor template, 14
 - CASE tool, 13
 - Communications Link shape, 15
 - creation of, 12-13
 - extension, 11
 - flight-booking application, 10-11
 - generic diagram, 10
 - issues of, 12
 - preconditions and postconditions, 11
 - in system boundary, 15
 - textual description, 11
 - UMLet, 13
- User groups, 323
- User Interfaces (UI), 322

V

- Variables, 325
- Variable scope
 - block-level scope, 330
 - module scope, 331
 - procedure scope, 331
- Visual Studio (VS)
 - assembly build and execute, 76
 - debugging features
 - conditional breakpoints, 79
 - locate and fix build errors, 81
 - through code, 77
 - features, 65
 - free versions, 347
 - IDE customization, 66
 - new project creation, 67
 - solution explorer and class view, 68
 - toolbox and properties window, 73
- Void keyword, 338

W, X, Y, Z

- Web applications
 - chapter overview, 221
 - control events, 228
 - data-bound web controls
 - DataSources, 243
 - model binding, 244–250
 - SQLDataSource control, 243
 - life cycle (web page)
 - page load event, 228
 - stages, 227
 - maintaining view state
 - code behind class, 240
 - debugger, 240
 - steps, 240
 - Textbox control, 240
 - reading and writing cookies, 241
 - server controls and web pages
 - absolute positioning, 233
 - control events, 236–237
 - controls, 234
 - div tag, 232
 - lblMessage controls, 237
 - Label control, 233
 - list items, 235–236
 - non-postback server control events, 235
 - properties window, 232
 - server-side control event handlers, 235
 - steps, 231
 - web control properties, 233
 - web site dialog box, 231
 - session and application state, 242–243
 - session events and application, 229
 - storing and sharing state
 - cookies, 238
 - query strings, 238
 - session and application state, 239
 - stateless protocol, 237
 - view state, 238
- Visual Studio web page designer
 - Split view, 226–227
 - Toolbox, 226
- Web Pages and Forms, 221
- web server control inheritance hierarchy
 - hierarchy chain (Page class), 224–225
 - Page class, 224
 - page directive, 223
 - System.Web.UI namespace, 224
 - TemplateControl class, 224
 - TextBox control, 225
- Web server controls, 223
- Web Pages and Web Forms
 - advantages, 221
 - .aspx file, 222
 - code-behind file links, 222
 - IE, 222
 - IIS, 223
 - programming logic, 221
 - visual interface, 221
- while statement, 337
- Widening type conversion, 332
- Windows Communication Foundation (WCF), 63
- Windows Communication Foundation (WCF) web services. *See also* RESTful web services
 - chapter overview, 273
 - consume
 - console window, 279
 - endpoint configuration, 278
 - service reference, 278
 - TaxService, 279
 - creation
 - contract, 276
 - exposed methods, 276
 - interface and class files, 275
 - request and response messages, 276–277
 - template, 274
 - testing, 277
 - data contracts
 - types, 279–280
 - XSD, 280
 - .NET client creation, 282
 - async method, 283
 - service reference, 282–283
 - startup projects, 284
 - testing, 284
 - service-oriented applications, 274
 - services, 273
 - WCF service
 - code replacement, 281
 - creation, 280

Windows Presentation Foundation (WPF)

- control and data templates, [214](#)
 - binding a ListBox, [216](#)
 - template creation, [217](#)
 - control events, [202](#)
 - custom dialog box, [206](#)
 - data binding, [207](#)
 - OneWay binding, [207](#)
 - TwoWay binding, [207](#)
 - using DataContext, [207](#)
 - data bound controls, [209](#)
 - dialog boxes, [204](#)
 - display controls, [196](#)
 - event handler method, [198](#)
 - delegation object, [198](#)
 - parameters, [199](#)
 - fundamentals, [193](#)
 - layout controls
 - Canvas control, [196](#)
 - DockPanel, [196](#)
 - fixed positioning, [195](#)
 - Grid control, [195](#)
 - StackPanel, [196](#)
 - WrapPanel, [196](#)
 - memo viewer interface, [200](#)
 - MessageBox, [205](#)
 - complex MessageBox, [206](#)
 - Show method, [205](#)
 - Visual Studio Designer, [197](#)
 - XAML, [194](#)
 - Button control, [194](#)
 - Grid.Row and Grid.Column attribute, [194](#)
 - property element, [194](#)
- Windows store applications
- application interface
 - data entry page, [258](#)
 - Grid.RowDefinitions, [257](#)
 - grid XAML, [256](#)
 - interaction mode, [259](#)
 - MainPage.xaml, [256](#)
 - Simulator, [258](#)
 - Steps, [256](#)
 - TextBox controls, [257](#)
 - chapter overview, [251](#)

- coding control events
 - Code Editor window, [260](#)
 - MainPage.xaml, [259](#)
 - StackPanel tag, [260](#)
 - steps, [259](#)
- data binding controls
 - Author class, [265](#)
 - author data, [261](#)
 - author information, [267](#)
 - AuthorTemplate, [266](#)
 - code file, [267](#)
 - CollectionViewSource, [260](#)
 - CollectionViewSource View methods, [261](#)
 - DataTemplate, [265–266](#)
 - full-data updating scenarios, [260](#)
 - GridView, [261–262](#)
 - INotifyPropertyChanged interface, [263](#)
 - MainPage.xaml file, [265](#)
 - modes, [268](#)
 - OnNavigatedTo, [267](#)
 - PropertyChanged event, [263–264](#)
 - steps, [263](#)
 - WPF, [260](#)
- handling control events
 - event handler method, [255](#)
 - parameters, [255](#)
 - properties window, [254–255](#)
 - WPF and ASP.NET apps, [254](#)
 - XAML editor, [254–255](#)
- page navigation, [269](#)
- passing data and page navigation
 - AuthorDetailPage, [270–272](#)
 - AuthorsGridView, [271](#)
 - btnBack_Click event handler, [271](#)
 - steps, [269](#)
 - TextBlock, [270](#)
- style sheets, [254](#)
- user interface
 - button control, [251](#)
 - Grid control, [253](#)
 - grid rows and columns, [252](#)
 - input form page, [253](#)
 - StackPanel control, [254](#)
 - visual interfaces, [251](#)
 - XAML, [252](#)

Beginning C# Object-Oriented Programming

Second Edition



Dan Clark

Apress®

Beginning C# Object-Oriented Programming

Copyright © 2013 by Dan Clark

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

ISBN-13 (pbk): 978-1-4302-4935-1

ISBN-13 (electronic): 978-1-4302-4936-8

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image, we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

President and Publisher: Paul Manning

Lead Editor: Gwenan Spearing

Technical Reviewer: Todd Meister

Editorial Board: Steve Anglin, Ewan Buckingham, Gary Cornell, Louise Corrigan, Morgan Ertel,

Jonathan Gennick, Jonathan Hassell, Robert Hutchinson, Michelle Lowman, James Markham,

Matthew Moodie, Jeff Olson, Jeffrey Pepper, Douglas Pundick, Ben Renow-Clarke, Dominic Shakeshaft,

Gwenan Spearing, Matt Wade, Tom Welsh

Coordinating Editor: Kevin Shea

Copy Editor: Larissa Shmailo

Compositor: SPi Global

Indexer: SPi Global

Artist: SPi Global

Cover Designer: Anna Ishchenko

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com.

For information on translations, please e-mail rights@apress.com, or visit www.apress.com.

Apress and friends of ED books may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Special Bulk Sales-eBook Licensing web page at www.apress.com/bulk-sales.

Any source code or other supplementary materials referenced by the author in this text is available to readers at www.apress.com. For detailed information about how to locate your book's source code, go to www.apress.com/source-code.

*This edition is dedicated to my father, whose technical
prowess is an inspiration to me every day!*

—Your Loving Son, Dan

Contents

About the Author	xvii
About the Technical Reviewer	xix
Acknowledgments	xxi
Introduction	xxiii
■ Chapter 1: Overview of Object-Oriented Programming	1
What is OOP?.....	1
The History of OOP	2
Why Use OOP?.....	2
The Characteristics of OOP.....	3
Objects.....	3
Abstraction	3
Encapsulation	4
Polymorphism.....	4
Inheritance.....	4
Aggregation	5
The History of C#.....	5
Summary.....	6
■ Chapter 2: Designing OOP Solutions: Identifying the Class Structure	7
Goals of Software Design	7
Understanding the Unified Modeling Language	8
Developing a SRS	9
Introducing Use Cases.....	10

Understanding Class Diagrams	16
Modeling Object Relationships.....	17
Association	17
Inheritance.....	18
Aggregation	18
Association Classes.....	19
Summary.....	24
■ Chapter 3: Designing OOP Solutions: Modeling the Object Interaction	25
Understanding Scenarios	25
Introducing Sequence Diagrams	26
Message Types	27
Recursive Messages	28
Message Iteration.....	28
Message Constraints.....	29
Message Branching.....	29
Understanding Activity Diagrams	34
Decision Points and Guard Conditions.....	34
Parallel Processing	35
Activity Ownership.....	35
Exploring GUI Design	39
GUI Activity Diagrams	39
Interface Prototyping	40
Interface Flow Diagrams	41
Application Prototyping	41
Summary.....	42
■ Chapter 4: Designing OOP Solutions: A Case Study.....	43
Developing an OOP Solution.....	43
Creating the System Requirement Specification	43
Developing the Use Cases	45
Diagramming the Use Cases	46

Developing the Class Model	48
Avoiding Some Common OOP Design Pitfalls	58
Summary	58
■ Chapter 5: Introducing the .NET Framework and Visual Studio	59
Introducing the .NET Framework.....	59
Goals of the .NET Framework	59
Components of the .NET Framework.....	61
Working with the .NET Framework.....	64
Using the Visual Studio Integrated Development Environment	65
Summary	81
■ Chapter 6: Creating Classes	83
Introducing Objects and Classes	83
Defining Classes.....	84
Creating Class Properties	84
Creating Class Methods.....	85
Using Constructors	89
Overloading Methods	90
Summary.....	96
■ Chapter 7: Creating Class Hierarchies.....	97
Understanding Inheritance	97
Creating Base and Derived Classes.....	97
Creating a Sealed Class.....	99
Creating an Abstract Class	99
Using Access Modifiers in Base Classes	99
Overriding the Methods of a Base Class	104
Calling a Derived Class Method from a Base Class	105
Calling a Base Class Method from a Derived Class	106
Overloading Methods of a Base Class	106
Hiding Base Class Methods.....	107

Implementing Interfaces	111
Understanding Polymorphism	111
Summary.....	117
■ Chapter 8: Implementing Object Collaboration.....	119
Communicating Through Messaging.....	119
Defining Method Signatures.....	119
Passing Parameters	120
Understanding Event-Driven Programming.....	121
Understanding Delegation	121
Implementing Events.....	122
Responding To Events	123
Windows Control Event Handling	123
Handling Exceptions in the .NET Framework	128
Using the Try-Catch Block	128
Adding a Finally Block	129
Throwing Exceptions	130
Nesting Exception Handling.....	130
Static Properties and Methods	131
Using Asynchronous Messaging.....	136
Summary.....	141
■ Chapter 9: Working with Collections	143
Introducing the .NET Framework Collection Types.....	143
Working with Arrays and Array Lists	144
Using Generic Collections.....	153
Programming with Stacks and Queues.....	157
Summary.....	160
■ Chapter 10: Implementing the Data Access Layer.....	161
Introducing ADO.NET	161
Working with Data Providers.....	162

Establishing a Connection	162
Executing a Command	164
Using Stored Procedures	165
Using the DataReader Object to Retrieve Data.....	166
Using the DataAdapter to Retrieve Data.....	167
Working with DataTables and DataSets	172
Populating a DataTable from a SQL Server Database.....	173
Populating a DataSet from a SQL Server Database.....	174
Establishing Relationships between Tables in a DataSet	174
Working with the Entity Framework.....	181
Querying Entities with LINQ to EF.....	184
Updating Entities with the Entity Framework.....	186
Summary.....	191
■ Chapter 11: Developing WPF Applications	193
Windows Fundamentals.....	193
Introducing XAML.....	194
Using Layout Controls	195
Adding Display Controls	196
Using the Visual Studio Designer	197
Handling Control Events.....	198
Creating and Using Dialog Boxes	204
Presenting a MessageBox to the User	205
Creating a Custom Dialog Box.....	206
Data Binding in Windows-Based GUIs.....	207
Binding Controls Using a DataContext.....	207
Creating and Using Control and Data Templates	214
Summary.....	219

■ Chapter 12: Developing Web Applications.....	221
Web Pages and Web Forms.....	221
Web Server Control Fundamentals.....	223
Understanding Web Page and Web Server Control Inheritance Hierarchy	223
Using the Visual Studio Web Page Designer.....	226
The Web Page Life Cycle	227
Control Events	228
Understanding Application and Session Events	229
Creating Server-Side Control Event Handlers.....	235
Storing and Sharing State in a Web Application.....	237
Maintaining View State.....	237
Using Query Strings.....	238
Using Cookies	238
Maintaining Session and Application State	239
Data-Bound Web Controls	243
Model Binding.....	244
Summary.....	250
■ Chapter 13: Developing Windows Store Applications	251
Building the User Interface.....	251
Using Style Sheets	254
Handling Control Events	254
Data Binding Controls.....	260
Page Navigation	269
Summary.....	272
■ Chapter 14: Developing and Consuming Web Services	273
What Are Services?	273
WCF Web Services.....	274
Creating a WCF Web Service	274

Consuming a WCF Web Service	278
Using Data Contracts	279
RESTful Data Services.....	285
Creating an ASP.NET Web API Service	285
Consuming ASP.NET Web API Services	290
Summary.....	294
■ Chapter 15: Developing the Office Supply Ordering Application	295
Revisiting Application Design	295
Building the OSO Application’s Data Access Layer	297
Building the OSO Application’s Business Logic Layer	304
Creating the OSO Application UI	308
Summary.....	320
■ Chapter 16: Wrapping Up.....	321
Improve Your Object-Oriented Design Skills.....	322
Investigate the .NET Framework Namespaces.....	322
Become Familiar with ADO.NET and the Entity Framework	322
Learn More about Creating Great User Interfaces (UI).....	322
Move toward Component-Based Development	323
Find Help	323
Join a User Group.....	323
Please Provide Feedback	323
Thank You, and Good Luck!	324
■ Appendix A: Fundamental Programming Concepts	325
Working with Variables and Data Types	325
Understanding Elementary Data Types.....	325
Integral Data Types	326
Non-Integral Data Types	326
Character Data Types.....	326
Boolean Data Type	327

Date Data Type.....	327
Object Data Type.....	327
Nullable Types.....	327
Introducing Composite Data Types	327
Structures.....	327
Arrays	328
Classes	329
Looking at Literals, Constants, and Enumerations	329
Literals.....	329
Constants.....	329
Enumerations.....	330
Exploring Variable Scope.....	330
Block-Level Scope	330
Procedure Scope	331
Module Scope.....	331
Understanding Data Type Conversion.....	332
Implicit Conversion	332
Explicit Conversion	332
Widening and Narrowing Conversions.....	332
Working with Operators.....	332
Arithmetic Operators	332
Comparison Operators.....	333
Logical Operators	334
Ternary Operator.....	334
Introducing Decision Structures.....	334
If Statements	335
Switch Statements	336
Using Loop Structures	337
While Statement	337
Do-While Statement	337

For Statement.....	337
For Each Statement.....	337
Introducing Methods	338
■ Appendix B: Exception Handling in C#.....	341
Managing Exceptions	341
Using the .NET Framework Exception Classes.....	343
The Importance of Using	344
■ Appendix C: Installing the Required Software.....	347
Installing the Sample Databases.....	347
Verifying the Database Installs	348
Index.....	353

About the Author



Dan Clark is a senior IT consultant specializing in .NET and SQL Server technology. He is particularly interested in C# programming and SQL Server Business Intelligence development. For over a decade, he has been developing applications and training others to develop applications using Microsoft technologies. Dan has published several books and numerous articles on .NET programming. He is a regular speaker at various developer conferences and user group meetings, and he conducts workshops in object-oriented programming and database development. He finds particular satisfaction in turning new developers on to the thrill of developing and designing object-oriented applications. In a previous life, he was a physics teacher and still loves the wonders and awe inspired by the study of the Universe and why things behave the way they do. You can reach Dan at clark.drc@gmail.com.

About the Technical Reviewer



Todd Meister has been working in the IT industry for over fifteen years. He has been the technical editor of over 75 titles on topics ranging from SQL Server to the .NET Framework. He is the senior IT architect at Ball State University in Muncie, Indiana. He lives in central Indiana with his wife, Kimberly, and their five entertaining children.

Acknowledgments

Thanks to the team at Apress for once again making the writing of this book as painless as possible while still keeping me on task. A special shout out goes to Gwenan and Kevin for their perseverance and considerable help with this project. And last but not least, to my technical reviewer Todd, thank you for your attention to detail and excellent suggestions while reviewing this book.

—Dan Clark