

.NET Fundamentals

What kind of problems is .NET designed to solve? .NET solves problems that have plagued programmers in the past. .NET helps programmers develop the applications of the future. This chapter is designed to present an overview of Microsoft .NET by looking at a simple program rather than talking in vague generalities. While we will start discussing Microsoft .NET in detail in Chapter 6, this chapter will enable you to get a feel for the big picture right away.

Problems of Windows Development

Imagine a symphony orchestra where the violins and the percussion sections had different versions of the score. It would require a heroic effort to play the simplest musical composition. This is the life of the Windows developer. Do I use MFC? Visual Basic or C++? ODBC or OLEDB? COM interface or C style API? Even within COM: do I use IDispatch, dual, or pure vtable interfaces? Where does the Internet fit into all of this? Either the design had to be contorted by the implementation technologies that the developers understood, or the developers had to learn yet another technological approach that was bound to change in about two years.

Deployment of applications can be a chore. Critical entries have to be made in a Registry that is fragile and difficult to back up. There is no good versioning strategy for components. New releases can break existing programs often with no information about what went wrong. Given the problems with the Registry, other technologies used other configuration stores such as a metabase or SQL Server.

Security in Win32 is another problem. It is difficult to understand and difficult to use. Many developers ignored it. Developers who needed to apply security often did the best they could with a difficult programming model. The rise of Internet-based security threats transforms a bad situation into a potential nightmare.

Despite Microsoft's efforts to make development easier problems remained. Many system services had to be written from scratch, essentially providing the plumbing code that had nothing to do with your business logic. MTS/COM+ was a giant step in the direction of providing higher level services, but it required yet another development paradigm. COM made real component programming possible. Nonetheless, you either did it simply, but inflexibly in Visual Basic, or powerfully, but with great difficulty in C++, because of all the repetitive plumbing code you had to write in C++.

Applications of the Future

Even if .NET fixed all the problems of the past, it would not be enough. One of the unchanging facts of programming life is that the boundaries of customer demand are always being expanded.

The growth of the Internet has made it imperative that applications work seamlessly across network connections. Components have to be able to expose their functionality to other machines. Programmers do not want to write the underlying plumbing code, they want to solve their customers' problems.

.NET Overview

The Magic of Metadata

To solve all these problems .NET must provide an underlying set of services that is available to all languages at all times. It also has to understand enough about an application to be able to provide these services.

Serialization provides a simple example. Every programmer at some time or another has to write code to save data. Why should every programmer have to reinvent the wheel of how to persist nested objects and complicated data structures? Why should every programmer have to figure out how to do this for a variety of data stores? .NET can do this for the programmer. Programmers can also decide to do it themselves if required.

To see how this is done, look at the **Serialize** sample associated with this chapter. For the moment ignore the programming details of C# which will be covered in the next three chapters, and focus on the concepts.

```
[Serializable] class Customer
{
    public string name;
    public long id;
}
class Test
{
    static void Main(string[] args)
    {
        ArrayList list = new ArrayList();

        Customer cust = new Customer();
        cust.name = "Charles Darwin";
        cust.id = 10;
        list.Add(cust);

        cust = new Customer();
        cust.name = "Isaac Newton";
        cust.id = 20;
        list.Add(cust);

        foreach (Customer x in list)
            Console.WriteLine(x.name + ": " + x.id);

        Console.WriteLine("Saving Customer List");
        FileStream s = new FileStream("cust.txt",
            FileMode.Create);
        SoapFormatter f = new SoapFormatter();
        f.Serialize(s, list);
        s.Close();

        Console.WriteLine("Restoring to New List");
        s = new FileStream("cust.txt", FileMode.Open);
        f = new SoapFormatter();
        ArrayList list2 = (ArrayList)f.Deserialize(s);
        s.Close();

        foreach (Customer y in list2)
            Console.WriteLine(y.name + ": " + y.id);
    }
}
```

We have defined a **Customer** class with two fields: a **name** and an **id**. The program first creates an instance of a collection class that will be used to hold instances of the **Customer** class. We add two **Customer** objects to the collection and then print out the contents of the collection. The collection is then saved to disk. It is restored to a new collection instance and printed out. The results printed out will be identical to those printed out before the collection was saved.¹

We wrote no code to indicate how the fields of the customer object are saved or restored. We did have to specify the format (SOAP) and create the medium to which the data was saved. The .NET Framework classes are partitioned so that where you load/save, the format you use to load/save, and how you load/save can be chosen independently. This kind of partitioning exists throughout the .NET Framework.

The **Customer** class was annotated with the **Serializable** attribute in the same way the **public** attribute annotates the name field. If you do not want your objects to be serializable, do not apply the attribute to your class. If an attempt is then made to save your object, an exception will be thrown and the program will fail.²

Attribute-based programming is used extensively throughout .NET to describe how the Framework should treat code and data. With attributes you do not have to write any code; the Framework takes the appropriate action based on the attribute. Security can be set through attributes. You can use attributes to have the Framework handle multithreading synchronization. Remoting of objects becomes straightforward through the use of attributes.

The compiler adds this **Serializable** attribute to the *metadata* of the **Customer** class to indicate that the Framework should save and restore the object. Metadata is additional information about the code and data within a .NET application. Metadata, a feature of the Common Language Runtime, provides such information about the code as:

- Version and locale information
- All the types
- Details about each type, including name, visibility, and so on
- Details about the members of each type, such as methods, the signatures of methods, and the like
- Attributes

Since metadata is stored in a programming-language-independent fashion with the code, not in a central store such as the Windows Registry, it makes .NET applications self-describing. The metadata can be queried at runtime to

¹ The sample installation should have already built an instance that you can run. If not, double-click on the Visual Studio.NET solution file that has the .sln suffix. When Visual Studio comes up, hit Control-F5 to build and run the sample.

² Comment out the **Serializable** attribute in the program (you can use the C/C++/* */ comment syntax) and see what happens.

get information about the code (such as the presence or absence of the **Serializable** attribute). You can extend the metadata by providing your own custom attributes.

In our example, the Framework can query the metadata to discover the structure of the **Customer** object in order to be able to save and restore it.

Types

Types are at the heart of the programming model for the CLR. A type is analogous to a class in most object-oriented programming languages, providing an abstraction of data and behavior, grouped together. A type in the CLR contains:

- Fields (data members)
- Methods
- Properties
- Events

There are also built-in primitive types, such as integer and floating point numeric types, string, etc. We will discuss types under the guise of classes and value types when we cover C#.

.NET Framework Class Library

The **Formatter** and **FileStream** classes are just two of more than 2500 classes in the .NET Framework that provide plumbing and system services for .NET applications. Some of the functionality provided by the .NET Framework includes:

- Base class library (basic functionality such as strings, arrays, and formatting)
- Networking
- Security
- Remoting
- Diagnostics
- I/O
- Database
- XML
- Web services that allow us to expose component interfaces over the Internet
- Web programming
- Windows User Interface

Interface-Based Programming

Suppose you want to encrypt your data and therefore do not want to rely on the Framework's serialization. Your class can inherit from the **ISerializable** interface and provide the appropriate implementation. (We will discuss how to do this in a later chapter.) The Framework will then use your methods to save and restore the data.

How does the Framework know that you implemented the **ISerializable** interface? It can query the metadata related to the class to see if it implements the interface! The Framework can then use either its own algorithm or the class's code to serialize or deserialize the object.

Interface-based programming is used in .NET to allow your objects to provide implementations to standard functionality that can be used by the Framework. Interfaces also allow you to program using methods on the interface rather than methods on the objects. You can program without having to know the exact type of the object. For example, the formatters (such as the SOAP formatter used here) implement the **IFormatter** interface. Programs can be written using the **IFormatter** interface and thus are independent of any particular current (binary, SOAP) or future formatter and still work properly.

Everything Is an Object

So if a type has metadata, the runtime can do all kinds of wonderful things. But does everything in .NET have metadata? Yes! Every type, whether it is user defined (such as **Customer**) or part of the Framework (such as **FileStream**), is a .NET object. All .NET objects have the same base class, the system's **Object** class. Hence everything that runs in .NET has a type and therefore has metadata.

In our example, the serialization code can walk through the **ArrayList** of customer objects and save each one as well as the array it belongs to, because the metadata allows it to understand the object's type and its logical structure.

Common Type System

The .NET Framework has to make some assumptions about the nature of the types that will be passed to it. These assumptions are the *Common Type System* (CTS). The CTS defines the rules for the types and operations that the Common Language Runtime will support. It is the CTS that limits .NET classes to single implementation inheritance. Since the CTS is defined for a wide range of languages, not all languages need to support all features of the CTS.

The CTS makes it possible to guarantee type safety, which is critical for writing reliable and secure code. As we noted in the previous section, every object has a type and therefore every reference to an object points to a

defined memory layout. If arbitrary pointer operations are not allowed, the only way to access an object is through its public methods and fields. Hence it's possible to verify an object's safety by analyzing the object. There is no need to know or analyze all the users of a class.

How are the rules of the CTS enforced? The Microsoft Intermediate Language (MSIL or IL) defines an instruction set that is used by all .NET compilers. This intermediate language is platform independent. The MSIL code can later be converted to a platform's native code. Verification for type safety can be done once based on the MSIL; it need not be done for every platform. Since everything is defined in terms of MSIL, we can be sure that the .NET Framework classes will work with all .NET languages. Design no longer dictates language choice; language choice no longer constrains design.

MSIL and the CTS make it possible for multiple languages to use the .NET Framework since their compilers produce MSIL. This one of the most visible differences between .NET and Java, which in fact share a great deal in philosophy.

ILDASM

The Microsoft Intermediate Language Disassembler (ILDASM) can display the metadata and MSIL instructions associated with .NET code. It is a very useful tool both for debugging and for increasing your understanding of the .NET infrastructure. You can use ILDASM to examine the .NET Framework code itself.³ Figure 2-1 shows a fragment of the MSIL code from the **Serialize** example, where we create two new customer objects and add them to the list.⁴ The **newobj** instruction creates a new object reference using the constructor parameter.⁵ **Stloc** stores the value in a local variable. **Ldloc** loads a local variable.⁶ It is strongly recommended that you play with ILDASM and learn its features.

³ ILDASM is installed on the Tools menu in Visual Studio.NET. It is also found in the Microsoft.NET\FrameworkSDK\Bin subdirectory. You can invoke it by double-clicking on its Explorer entry or from the command line. If you invoke it from the command line (or from VS.NET) you can use the /ADV switch to get some advanced options.

⁴ Open Serialize.exe and Click on the plus (+) sign next to Test. Double-click on Main to bring up the MSIL for the Main routine.

⁵ Technically it is not a parameter. IL is a stack-based language, and the constructor is a metadata token previously pushed on the stack.

⁶ You can read all about MSIL in the ECMA documents, specifically the Partition III CIL Instruction Set.

```

[mscorlib:mscorlib.dll]
IL_0000: newobj instance void [mscorlib:System.Collections.ArrayList..ctor()
IL_0001: stloc.0
IL_0002: newobj instance void Customer..ctor()
IL_0003: stloc.1
IL_0004: ldstr "Charles Darwin"
IL_0005: stfld string Customer:Name
IL_0006: ldloc.1
IL_0007: conv.i4
IL_0008: conv.i8
IL_0009: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_000A: stloc.0
IL_000B: ldstr "Issue Newton"
IL_000C: stfld string Customer:Name
IL_000D: ldloc.1
IL_000E: conv.i4
IL_000F: conv.i8
IL_0010: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0011: pop
IL_0012: ldstr "Issue Newton"
IL_0013: stfld string Customer:Name
IL_0014: ldloc.1
IL_0015: conv.i4
IL_0016: conv.i8
IL_0017: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0018: pop
IL_0019: ldstr "Issue Newton"
IL_001A: stfld string Customer:Name
IL_001B: ldloc.1
IL_001C: conv.i4
IL_001D: conv.i8
IL_001E: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_001F: pop
IL_0020: ldstr "Issue Newton"
IL_0021: stfld string Customer:Name
IL_0022: ldloc.1
IL_0023: conv.i4
IL_0024: conv.i8
IL_0025: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0026: pop
IL_0027: ldstr "Issue Newton"
IL_0028: stfld string Customer:Name
IL_0029: ldloc.1
IL_002A: conv.i4
IL_002B: conv.i8
IL_002C: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_002D: pop
IL_002E: ldstr "Issue Newton"
IL_002F: stfld string Customer:Name
IL_0030: ldloc.1
IL_0031: conv.i4
IL_0032: conv.i8
IL_0033: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0034: pop
IL_0035: ldstr "Issue Newton"
IL_0036: stfld string Customer:Name
IL_0037: ldloc.1
IL_0038: conv.i4
IL_0039: conv.i8
IL_003A: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_003B: pop
IL_003C: ldstr "Issue Newton"
IL_003D: stfld string Customer:Name
IL_003E: ldloc.1
IL_003F: conv.i4
IL_0040: conv.i8
IL_0041: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0042: pop
IL_0043: ldstr "Issue Newton"
IL_0044: stfld string Customer:Name
IL_0045: ldloc.1
IL_0046: conv.i4
IL_0047: conv.i8
IL_0048: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0049: pop
IL_004A: ldstr "Issue Newton"
IL_004B: stfld string Customer:Name
IL_004C: ldloc.1
IL_004D: conv.i4
IL_004E: conv.i8
IL_004F: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0050: pop
IL_0051: ldstr "Issue Newton"
IL_0052: stfld string Customer:Name
IL_0053: ldloc.1
IL_0054: conv.i4
IL_0055: conv.i8
IL_0056: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0057: pop
IL_0058: ldstr "Issue Newton"
IL_0059: stfld string Customer:Name
IL_005A: ldloc.1
IL_005B: conv.i4
IL_005C: conv.i8
IL_005D: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_005E: pop
IL_005F: ldstr "Issue Newton"
IL_0060: stfld string Customer:Name
IL_0061: ldloc.1
IL_0062: conv.i4
IL_0063: conv.i8
IL_0064: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0065: pop
IL_0066: ldstr "Issue Newton"
IL_0067: stfld string Customer:Name
IL_0068: ldloc.1
IL_0069: conv.i4
IL_006A: conv.i8
IL_006B: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_006C: pop
IL_006D: ldstr "Issue Newton"
IL_006E: stfld string Customer:Name
IL_006F: ldloc.1
IL_0070: conv.i4
IL_0071: conv.i8
IL_0072: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0073: pop
IL_0074: ldstr "Issue Newton"
IL_0075: stfld string Customer:Name
IL_0076: ldloc.1
IL_0077: conv.i4
IL_0078: conv.i8
IL_0079: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_007A: pop
IL_007B: ldstr "Issue Newton"
IL_007C: stfld string Customer:Name
IL_007D: ldloc.1
IL_007E: conv.i4
IL_007F: conv.i8
IL_0080: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0081: pop
IL_0082: ldstr "Issue Newton"
IL_0083: stfld string Customer:Name
IL_0084: ldloc.1
IL_0085: conv.i4
IL_0086: conv.i8
IL_0087: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0088: pop
IL_0089: ldstr "Issue Newton"
IL_008A: stfld string Customer:Name
IL_008B: ldloc.1
IL_008C: conv.i4
IL_008D: conv.i8
IL_008E: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_008F: pop
IL_0090: ldstr "Issue Newton"
IL_0091: stfld string Customer:Name
IL_0092: ldloc.1
IL_0093: conv.i4
IL_0094: conv.i8
IL_0095: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_0096: pop
IL_0097: ldstr "Issue Newton"
IL_0098: stfld string Customer:Name
IL_0099: ldloc.1
IL_009A: conv.i4
IL_009B: conv.i8
IL_009C: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_009D: pop
IL_009E: ldstr "Issue Newton"
IL_009F: stfld string Customer:Name
IL_00A0: ldloc.1
IL_00A1: conv.i4
IL_00A2: conv.i8
IL_00A3: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00A4: pop
IL_00A5: ldstr "Issue Newton"
IL_00A6: stfld string Customer:Name
IL_00A7: ldloc.1
IL_00A8: conv.i4
IL_00A9: conv.i8
IL_00AA: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00AB: pop
IL_00AC: ldstr "Issue Newton"
IL_00AD: stfld string Customer:Name
IL_00AE: ldloc.1
IL_00AF: conv.i4
IL_00B0: conv.i8
IL_00B1: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00B2: pop
IL_00B3: ldstr "Issue Newton"
IL_00B4: stfld string Customer:Name
IL_00B5: ldloc.1
IL_00B6: conv.i4
IL_00B7: conv.i8
IL_00B8: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00B9: pop
IL_00BA: ldstr "Issue Newton"
IL_00BB: stfld string Customer:Name
IL_00BC: ldloc.1
IL_00BD: conv.i4
IL_00BE: conv.i8
IL_00BF: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00C0: pop
IL_00C1: ldstr "Issue Newton"
IL_00C2: stfld string Customer:Name
IL_00C3: ldloc.1
IL_00C4: conv.i4
IL_00C5: conv.i8
IL_00C6: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00C7: pop
IL_00C8: ldstr "Issue Newton"
IL_00C9: stfld string Customer:Name
IL_00CA: ldloc.1
IL_00CB: conv.i4
IL_00CC: conv.i8
IL_00CD: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00CE: pop
IL_00CF: ldstr "Issue Newton"
IL_00D0: stfld string Customer:Name
IL_00D1: ldloc.1
IL_00D2: conv.i4
IL_00D3: conv.i8
IL_00D4: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00D5: pop
IL_00D6: ldstr "Issue Newton"
IL_00D7: stfld string Customer:Name
IL_00D8: ldloc.1
IL_00D9: conv.i4
IL_00DA: conv.i8
IL_00DB: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00DC: pop
IL_00DD: ldstr "Issue Newton"
IL_00DE: stfld string Customer:Name
IL_00DF: ldloc.1
IL_00E0: conv.i4
IL_00E1: conv.i8
IL_00E2: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00E3: pop
IL_00E4: ldstr "Issue Newton"
IL_00E5: stfld string Customer:Name
IL_00E6: ldloc.1
IL_00E7: conv.i4
IL_00E8: conv.i8
IL_00E9: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00EA: pop
IL_00EB: ldstr "Issue Newton"
IL_00EC: stfld string Customer:Name
IL_00ED: ldloc.1
IL_00EE: conv.i4
IL_00EF: conv.i8
IL_00F0: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00F1: pop
IL_00F2: ldstr "Issue Newton"
IL_00F3: stfld string Customer:Name
IL_00F4: ldloc.1
IL_00F5: conv.i4
IL_00F6: conv.i8
IL_00F7: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00F8: pop
IL_00F9: ldstr "Issue Newton"
IL_00FA: stfld string Customer:Name
IL_00FB: ldloc.1
IL_00FC: conv.i4
IL_00FD: conv.i8
IL_00FE: call void [mscorlib:System.Collections.ArrayList:Add(object)
IL_00FF: pop

```

Figure 2-1

Code fragment from *Serialize* example.

Language Interoperability

Having all language compilers use a common intermediate language and common base class make it *possible* for languages to interoperate. But since all languages need not implement all parts of the CTS, it is certainly possible for one language to have a feature that another does not.

The *Common Language Specification* (CLS) defines a subset of the CTS representing the basic functionality that all .NET languages should implement if they are to interoperate with each other. This specification enables a class written in Visual Basic.NET to inherit from a class written in COBOL.NET or C#, or to make interlanguage debugging possible. An example of a CLS rule is that method calls need not support a variable number of arguments, even though such a construct can be expressed in MSIL.

CLS compliance applies only to publicly visible features. A class, for example, can have a private member that is non-CLS compliant and still be a base class for a class in another .NET language. For example, C# code should not define public and protected class names that differ only by case-sensitivity, since languages such as VB.NET are not case-sensitive. Private fields could have case-sensitive names.

Microsoft itself is providing several CLS-compliant languages: C#, Visual Basic.NET, and C++ with Managed Extensions. Third parties are providing

additional languages (there are over a dozen so far). ActiveState is implementing Perl and Python. Fujitsu is implementing COBOL.

Managed Code

In the serialization example a second instance of the Customer object was assigned to the same variable (**cust**) as the first instance without freeing it. None of the allocated storage in the example was ever deallocated. .NET uses automatic garbage collection to reclaim memory. When memory allocated on the heap becomes orphaned, or passes out of scope, it is placed on a list of memory locations to be freed. Periodically, the system runs a garbage collection thread that returns the memory to the heap.

By having automatic memory management the system has eliminated memory leakage, which is one of the most common programming errors. In most cases, memory allocation is much faster with garbage collection than with classic heap allocation schemes. Note that variables such as **cust** and **list** are object references, not the objects themselves. This makes the garbage collection possible.

Garbage collection is one of several services provided by the *Common Language Runtime* (CLR) to .NET programs.⁷ Data that is under the control of the CLR garbage collection process is called managed data. Managed code is code that can use the services of the CLR. .NET compilers that produce MSIL can produce managed code.

Managed code is not automatically type safe. C++ provides the classic example. You can use the `__gc` attribute to make a class garbage collected. The C++ compiler will prevent such classes from using pointer arithmetic. Nonetheless, C++ cannot be reliably verified.⁸

Code is typically verified for type safety before compilation. This step is optional and can be skipped for trusted code. One of the most significant differences between verified and unverified code is that verified code cannot

⁷ Technically, metadata, the CTS, the CLS, and the Virtual Execution System (VES) are also part of the CLR. We are using CLR here in the sense that it is commonly used. The VES loads and runs .NET programs and supports late binding. For more details refer to the Common Language Infrastructure (CLI) Partition I: Concepts and Architecture document submitted to ECMA. This document is loaded with the .NET Framework SDK.

⁸ The most immediate reason for this is that the C Runtime Library (CRT) that is the start-up code for C++ programs was not converted to run under .NET because of time constraints. Even if this were to be done, however, there are two other obstacles to verifying C++ code. First, to ensure that the verification process can complete in a reasonable amount of time, the CLR language specifications require certain IL language patterns to be used and the managed C++ compiler would have to be changed to accommodate this. Second, after disallowing the C++ constructs that inhibit verification (like taking the address of a variable on the stack, or pointer arithmetic), you would wind up with a close approximation to the C# language.

use pointers.⁹ Code that used pointers could subvert the Common Type System and access any memory location.

Type safe code cannot be subverted. A buffer overwrite is not able to corrupt other data structures or programs. Methods can only start and end at well-defined entry and exit points. Security policy can be applied to type safe code.¹⁰ For example, access to certain files or user interface features can be allowed or denied. You can prevent the execution of code from unknown sources. You can prevent access to unmanaged code to prevent subversion of .NET security. Type safety also allows paths of execution of .NET code to be isolated from one another.¹¹

Assemblies

Another function of the CLR is to load and run .NET programs.

.NET programs are deployed as assemblies. An *assembly* is one or more EXEs or DLLs with associated metadata information. The metadata about the entire assembly is stored in the assembly's manifest. The manifest contains, for example, a list of the assemblies upon which this assembly is dependent.

In our `Serialize` example there is only file in the assembly, **serialize.exe**. That file contains the metadata as well as the code. Since the manifest is stored in the assembly and not in a separate file (like a type library or registry), the manifest cannot get out of sync with the assembly. Figure 2–2 shows the metadata in the manifest for this example.¹² Note the **assembly extern** statements that indicate the dependencies on the Framework assemblies **mscorlib** and **System.Runtime.Formatters.SOAP**. These statements also indicate the version of those assemblies that `serialize.exe` depends on.

Assemblies can be versioned, and the version is part of the name for the assembly. To version an assembly it needs a unique name. Public/private encryption keys are used to generate a unique (or strong) name.

Assemblies can be deployed either privately or publicly. For private deployment all the assemblies that an application needs are copied to the same directory as the application. If an assembly is to be publicly shared, an entry is made in the *Global Assembly Cache* (GAC) so that other assemblies can locate it. For assemblies put in the GAC a strong name is required. Since the version is part of the assembly name, multiple versions can be deployed

⁹ It would not be correct to say that code written in MSIL is managed code. The CTS permits MSIL to have unmanaged pointers in order to work with unmanaged data in legacy code. The reverse is not true; unmanaged code cannot access managed data. The CLS prohibits unmanaged pointers.

¹⁰ This is discussed in more detail in Chapter 12.

¹¹ See the discussion of Application Domains in Chapter 8.

¹² Open **serialize.exe** in ILDASM and double-click on the MANIFEST item.

```

assembly extern mscorlib
{
    publickeytoken = (B7 7A 5C 56 19 34 8E 89 )
    var 1:0:24110
}
assembly extern System.Runtime.Serialization.Formatters.Soap
{
    publickeytoken = (B0 3F 5F 7F 11 D8 3A 3A )
    var 1:0:24110
}
assembly Serialize
{
    custom instance void [mscorlib:System.Reflection.AssemblyKeyTokenAttribute::ctor(string) = ( 01 00 00 00 00 )
    custom instance void [mscorlib:System.Reflection.AssemblyKeyFileAttribute::ctor(string) = ( 03 00 00 00 00 )
    custom instance void [mscorlib:System.Reflection.AssemblyDelaySignAttribute::ctor(bool) = ( 01 00 00 00 00 )
    custom instance void [mscorlib:System.Reflection.AssemblyTrademarkAttribute::ctor(string) = ( 04 00 00 00 00 )
    custom instance void [mscorlib:System.Reflection.AssemblyCopyrightAttribute::ctor(string) = ( 05 00 00 00 00 )
    custom instance void [mscorlib:System.Reflection.AssemblyProductAttribute::ctor(string) = ( 06 00 00 00 00 )
    custom instance void [mscorlib:System.Reflection.AssemblyCompanyAttribute::ctor(string) = ( 07 00 00 00 00 )
    custom instance void [mscorlib:System.Reflection.AssemblyConfigurationAttribute::ctor(string) = ( 08 00 00 00 00 )
    custom instance void [mscorlib:System.Reflection.AssemblyDescriptionAttribute::ctor(string) = ( 09 00 00 00 00 )
    custom instance void [mscorlib:System.Reflection.AssemblyTitleAttribute::ctor(string) = ( 0A 00 00 00 00 )
    // -- The following custom attribute is added automatically, do not uncomment! -----
    // custom instance void [mscorlib:System.Diagnostics.DelegatableAttribute::ctor(bool,
    // bool) = ( 01 00 01 01 00 00 )
    //
    hash algorithm (00000004)
    var 1:0x5423974
}
module Serialize.exe
{
    MVID: (10206554-C296-419C-8E2A-8144E38FC4B9)
    language C#00000000
    subsystem 0x00000001
    file alignment 512
}

```

Figure 2-2

Manifest for the Serialize assembly.

side by side on the same machine without interfering with each other. Whether you use public or private deployment there is no more “DLL Hell.”¹³

Assembly deployment with language interoperability makes component development almost effortless.

JIT Compilation

Before executing on the target machine, MSIL has to be translated into the machine’s native code. This can either be done before the application is called, or at runtime. At runtime, the translation is done by a just-in-time (JIT) compiler. The Native Image Generator (Ngen.exe) translates MSIL into native code so that it is already translated when the program is started.

The advantage of pretranslation is that optimizations can be performed. Optimizations are generally impractical with JIT because the time it takes to do the optimization can be longer than it takes to compile the code. Start-up time is also faster with pretranslation because no translation has to be done when the application starts.

¹³ This is discussed in much more detail in Chapter 7.

The advantage of JIT is that it knows what the execution environment is when the program is run and can make better assumptions, such as register assignments, when it generates the code. Only the code that is actually executed is translated, code that never gets executed is never translated.

In the first release of .NET, the Native Image Generator and the JIT compiler use the same compiler. No optimizations are done for Ngen, its only current advantage is faster start-up. For this reason we do not discuss Ngen in this book.

Performance

You may like the safety and ease-of-use features of managed code but you might be concerned about performance. Early assembly language programmers had similar concerns when high-level languages came out.

The CLR is designed with high performance in mind. With JIT compilation, the first time a method is encountered, the CLR performs verifications and then compiles the method into native code (which will contain safety features, such as array bounds checking). The next time the method is encountered, the native code executes directly. Memory management is designed for high performance. Allocation is almost instantaneous, just taking the next available storage from the managed heap. Deallocation is done by the garbage collector, which has an efficient multiple-generation algorithm.

You do pay a penalty when security checks have to be made that require a stack walk as we will explain in the Security chapter.

Web pages use compiled code, not interpreted code. As a result ASP.NET is much faster than ASP.

For 98% of the code that programmers write, any small loss in performance is far outweighed by the gains in reliability and ease of development. High performance server applications might have to use technologies such as ATL Server and C++.

Summary

.NET solves the problems that have plagued Windows development in the past. There is one development paradigm for all languages. Design and programming language choices are no longer in conflict. Deployment is more rational and includes a versioning strategy. While we will talk more about it in later chapters, metadata, attribute-based security, code verification, and type-safe assembly isolation make developing secure applications much easier. The plumbing code for fundamental system services is provided, yet you can extend or replace it if you must.

The Common Language Runtime provides a solid base for developing applications of the future. The CLR is the foundation whose elements are the Common Type System, metadata, the Common Language Specification, and the Virtual Execution System (VES) that executes managed code.¹⁴ As we shall see in future chapters, .NET makes it easier to develop Internet applications for both service providers and customer-based solutions. With the unified development platform .NET provides, it will be much easier than in the past for Microsoft or others to provide extensions.

All this is made possible by putting old technologies together in the CLR creatively: intermediate languages, type-safe verification, and of course, metadata. As you will see, metadata is used in many features in .NET.

We shall expand on these topics in the course of the book. We next cover the C# language. Depending on your knowledge of C#, you might be able to skim Chapter 3 and skim Chapters 4 and 5. Chapter 4 introduces the Acme Travel Agency case study, which is used throughout the book. Chapter 5 covers important topics about the interaction of C# and the .NET Framework.

¹⁴ The Base Class Libraries classes (BCL) are also part of the CLR.