

Developing Next-Generation Web Applications

Web applications have historically been less rich and responsive than desktop applications. End users don't necessarily understand the details of how an application works, but they know that interacting with a website in the browser is distinctly different from using an application installed locally. When a development team tackles a new project, one of the first questions they are faced with is whether end users can accept the limitations of web development or whether they need to require a client desktop application to be installed. Web applications are accessible from just about any browser, just about anywhere, but they are limited by what you can do with markup and script code running in the browser.

Desktop applications, also called *fat client* applications, require that the user perform an installation on their machine, but let developers leverage the advanced mouse and graphics capabilities of the operating system that would be extremely difficult to implement in a web browser, and also take advantage of the user's machine for tasks such as offline storage. Conversely, web applications can be updated just by changing what is running on the server, and site visitors get the latest version instantaneously. However, it's much more difficult to update a desktop application, because you'd have to get users to perform yet another installation or else ensure that the application has been coded to include a clever system for doing updates automatically.

Web applications are said to use a zero-deployment model, but desktop applications use a heavy deployment and configuration model. The choice is often characterized as a tradeoff between *rich* and *reach*: Desktop applications generally offer a richer user experience than what could be offered in the browser, but with a web application you are able to reach users anywhere on any OS with almost no extra effort. Further, many companies have restrictive policies in place regarding what software can be installed on employees' machines, and they often don't allow employees to have administrative access that is required to install new applications, so web applications will be the only viable option in many situations.

Bringing Richness to Web Applications

Years ago, having a web presence was a distinguishing factor for companies. That is no longer the case. Now just having a web presence is no longer enough. Companies are distinguishing themselves further through web applications that react intuitively to customer actions and anticipate user input. This book shows you how ASP.NET AJAX addresses specific web development challenges and paves the way for taking your website to another level of user experience. In this chapter, I discuss the need for richer frameworks in web application development. I talk about the key pieces of the ASP.NET AJAX platform and highlight some other options.

The fundamental set of technologies that enable the next generation of web applications are not new. Online news articles and blogs point to Google, Flickr, and several other services as prime examples of leveraging these technologies in unique ways. The applications have some unique features, but in reality, the underlying technologies have been around and in use for nearly a decade. Take a look at how Microsoft Exchange Server provided rich access to email from a web browser in the Outlook Web Access application, and you can see that the concept of ubiquitous access from a browser while leveraging a common set of browser features for a rich user experience has been around and in practice for years. Users get a remarkably full-featured application with no local installation and are able to access e-mail from virtually any machine.

The technologies in use for building rich applications are generally referred to as AJAX (Asynchronous JavaScript and XML). While the acronym is nice, it doesn't do much to explain what is actually happening. Instead of building a web application to be just a series of page views and post backs, developers are using JavaScript to communicate asynchronously with the web server and update parts of the page dynamically. This means that the web page can dynamically adapt its appearance as the user interacts with it, and it can even post or fetch data to or from the web server in the background. Gone are the days of the ugly post back, which clears the user's screen and breaks his concentration! Instead, we need to post back now only if we want to change to a different web page.

Even that rule can be bent. Some applications are pushing this boundary and completely changing the user's view, just as though they navigated to a new page, but they do so through an asynchronous post and by changing the page content without actually navigating to a new URL.

The AJAX acronym refers to XML as the data format being exchanged between client and server, but in reality, applications are being built that retrieve simple pieces of text, XML, and JSON (JavaScript Object Notation) (which I discuss in more detail in Chapter 4). Part of the AJAX appeal isn't even covered by the acronym: In addition to communicating with the server without blocking, developers are leveraging Dynamic HTML (DHTML) and Cascading Style Sheets (CSS) to create truly amazing user interfaces. JavaScript code running on the client communicates asynchronously with the server and then uses DHTML to dynamically modify the page, which supports rich animations, transitions, and updates to the content while the user continues interacting with the page. In many cases, users won't even realize they are using a web application!

Who Benefits from AJAX?

AJAX offers benefits to both end users and developers. For end users, it reduces the "rich or reach" conflict; for developers, it helps in overcoming the constraints raised by HTTP.

Why End Users Want AJAX Applications

Users tend to view desktop applications as a sort of commitment. They install a program, usually from a disk pulled from a costly shrink-wrapped box. The program consumes hard disk space as well as a position in the program menu. The user may need to update the program periodically or perform an upgrade later on to get new features. If the program is proactive about updating itself, the user is confronted regularly with dialogs about accepting patches or downloads. In exchange for this investment of time, money, and energy, the user gets repaid by an application that is able to leverage the operating system and machine resources. It is a rich application. It has local storage capabilities, offers quick response times, and can present a compelling and intuitive graphical user interface.

More and more applications are accessible from the web browser, where the full resources of the hardware and OS are not available, but the user commitment of a desktop application is not required. Over the years, interacting with a web application has meant a predictable pattern for users. They click a link in the page and the browser flashes while the user waits until the screen is repainted (the dreaded post back). This cycle is repeated again and again. The user looks at what is presented on the page, interacts with it, and clicks somewhere. The browser produces an audible click for feedback and begins to post back to the server. The screen of the web browser flashes blank and some icon spins or flashes while the user waits for a new version of the page to be returned from the server. Many times, the new version of the page is almost exactly the same as the previous version, with only part of the page being updated. And then the cycle begins all over again. This has a sluggish feeling even when the user has a high-speed network connection.

The AJAX set of technologies has changed what users expect from web applications. JavaScript code running in the browser works to exchange data with the web server asynchronously. There is no click sound and the browser does not flash. The request to the server is nonblocking, which means the user is able to continue viewing the page and interacting with it. The script gets the updated data from the server and modifies the page dynamically using the DHTML coding methodology. The user is able to continue looking at the page while parts of it are updated in the background. AJAX is used to provide a more responsive experience, making web applications behave more like desktop installations. JavaScript is used to provide a richer experience with support for drag-and-drop, modal dialogs, and seemingly instantaneous updates to various parts of the page based on user inputs.

A big part of successfully leveraging AJAX technologies is in the perceived performance increase. Users appreciate web applications that anticipate their actions. If you also use JavaScript code in the background to pre-fetch images and data that may be needed, users can get a speedy response without the usual pause that accompanies their actions. Nobody wants to wait for data exchanges between client and server; studies have shown that a time lag between user input and subsequent UI changes can significantly reduce their productivity and give them the frustrating feeling that they are fighting the application. Users want web applications to behave like desktop installations but without the overhead associated with an installation. As more applications employ smart caching, anticipate user actions, and provide richer UI, the difference between web and desktop applications is becoming blurred. Expectations of web applications are rising. The end user has now seen that it is possible to avoid the commitment of installing a desktop application and still have a rich and responsive experience.

Why Developers Want AJAX

Often, the first question to arise when starting a new development project is what type of application it will be. Should it be a desktop application or a web application? This is a key decision because it has

Chapter 1: Developing Next-Generation Web Applications

historically dictated a lot about the nature of the application and the development problem space. Many developers are now choosing to build web applications by default unless something about the application dictates that it must be a desktop install. If it must run offline or if it requires a user interface that is complex to achieve in HTML, targeting the web browser may be ruled out, and the choice to write a standalone application is forced.

Developers have a difficult job writing modern web applications due to the inherent worldwide-web functionality constraints imposed by the use of the Hypertext Transfer Protocol (HTTP) and the way browsers use it. HTTP is a stateless protocol. The web browser requests a page, possibly carrying some querystring or form input parameters, and the server processes the request and sends a response that includes HTML-rendered content. The server can only react to the information supplied in the current request and doesn't know, based on the information in the request itself, details about the path the user took to get to the current view. When the response is rendered, the connection may be broken and the server won't have any information to preserve for the next request. From the server's perspective, it is simply listening for requests to come in from any browser anywhere and then reacting. The browser issues a request to the page and receives an HTML page in response. It uses the HTML it receives to render the user interface. The user interacts with the page, and, in response, the browser clears the screen and submits a new request to the server, carrying some information about user input or actions. Again, a complete HTML page is returned. The browser then presents the new version of HTML. Fundamentally, the HTTP protocol is stateless. The server gets a request and responds to it. The request carries limited information about the ongoing conversation that is happening between client and server.

AJAX makes this much better. AJAX breaks this pattern by updating portions of the page separately, via *partial page rendering*. Figure 1-1 shows a typical non-AJAX series of browser and server interactions. Each request results in a full page rendering, and, in response, the browser updates the user's entire view.

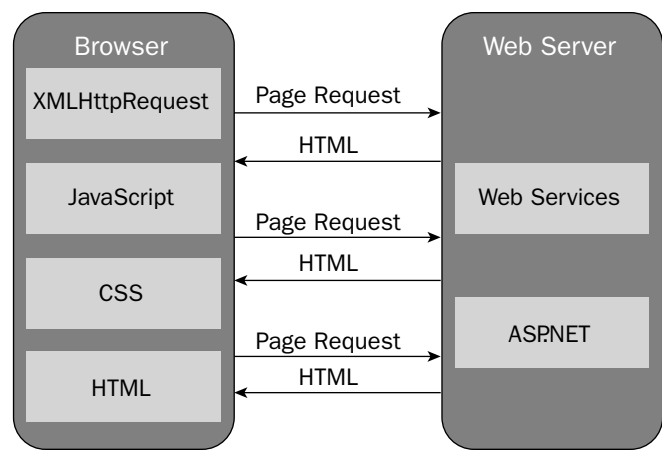


Figure 1-1

In Figure 1-2, AJAX is employed to improve the user's experience. A request is made for the initial page rendering. After that, asynchronous requests to the server are made. An asynchronous request is a background request to send or receive data in an entirely nonvisual manner. They are asynchronous because

the user interface isn't frozen during this time, and users can continue interacting with the page while the data transfer is taking place. These calls get just an incremental update for the page instead of getting an entirely new page. JavaScript running on the client reacts to the new data and updates various portions of the page, as desired. The number of requests to the server may be no different, or in some cases there may actually be more calls to the server, but the user perception is that the application feels more responsive. They aren't forced to pause, even if it's only a slight pause, and wait for the server while staring at a blank browser screen.

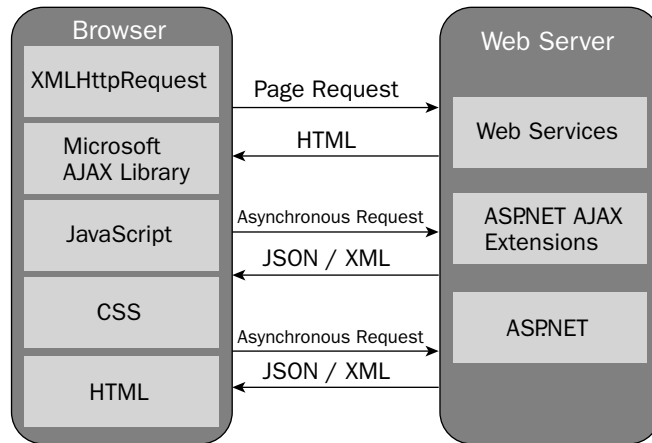


Figure 1-2

Almost a decade ago, the Microsoft Exchange Server team created an ActiveX control called `XmlHttpRequest` that could be instantiated from JavaScript and used to communicate with the server. This can occur without clearing the screen to paint a whole new page. Using the `XmlHttpRequest` object, you could send information to the server and get data back without requiring a whole new HTML page. JavaScript code could then manipulate the HTML dynamically on the client, avoiding the annoying flash and the wait that users associate with web browsing. This functionality was not limited to Internet Explorer for long. Soon, other browsers included `XmlHttpRequest` objects as well. Developers could now write richer applications with reach extending across various operating systems.

The browsers also created an advanced DOM (Document Object Model) to represent the browser, the window, the page, and the HTML elements it contained. The DOM exposed events and responded to input, allowing the page to be manipulated with script. Dynamic HTML (DHTML) opened the door to writing rich interfaces hosted within the web browser. Developers started writing hundreds and even thousands of lines of JavaScript code to make rich and compelling applications that would not require any client installation and could be accessed from any browser anywhere. Web applications began to move to a whole new level of richness. Without AJAX libraries, you would be faced with writing lots and lots of JavaScript code and debugging the sometimes subtle variations in different browsers to reach this new level of richness.

JavaScript Libraries and AJAX

Developers have had access to AJAX technologies for years, and many have been leveraging AJAX to push the limits of what can be done in the web browser. But what is really making AJAX more compelling now are the comprehensive script libraries and integration with server technologies that make it easier to write rich web applications and avoid the need to become an expert on the different versions of JavaScript. A JavaScript library is referenced within the HTML of a page using the `<script>` tag:

```
<html>
  <head>
    <script src="http://theGibbs.us/someScript.js"
type="text/javascript"></script>
  </head>
  ...
```

The script is downloaded and cached by the browser. Other pages within the application can reference the same URL for script, and the browser won't even bother to go back to the server to get it. The functionality of that script file is available for use from within the page rendered to the browser. A script library sent to the browser and then leveraged for writing a richer UI and a more responsive application is at the heart of all AJAX libraries.

What Is ASP.NET AJAX?

ASP.NET AJAX is the name of Microsoft's AJAX solution, and it refers to a set of client and server technologies that focus on improving web development with Visual Studio. Other companies have their own AJAX solution, often taking a radically different approach, but Microsoft has sought to build upon the popular ASP.NET technology by developing a comprehensive set of script libraries as well as server programming enhancements. One piece, the Microsoft AJAX Library, is a JavaScript library that works on a variety of browsers and serves to simplify JavaScript development. It provides you with the capability to easily write JavaScript using an object-oriented programming approach similar to code-behind programming in ASP.NET. You can use the functions of the Microsoft AJAX library to interact with the DOM, dynamically update portions of the page, communicate with the web server asynchronously, and write rich animations on the page. You'll learn about the client features in Chapters 6 and 7, but for now think of the Microsoft AJAX Library as a type system and set of class libraries that simplify writing JavaScript to enhance the user experience, manipulate the DOM, and communicate with the web server. The huge benefit in using this library is that it abstracts the tedious low-level DHTML programming into an OOP model that is much easier to work with.

The other part of the ASP.NET AJAX release is the server-side ASP.NET 2.0 AJAX Extensions. These extensions build on top of the ASP.NET classes and controls and leverage the Microsoft AJAX Library sent to the browser. They make it easy to quickly take advantage of AJAX technologies to enrich an

application. Through a set of standard web services, ASP.NET AJAX is also able to leverage server-based application services such as authentication and user profile storage. The extensions both extend some of the existing ASP.NET controls by adding new behaviors and include a new set of server controls that make it easy to do asynchronous updates to a page and leverage server resources.

The client and server pieces of ASP.NET are shown in Figure 1-3. ASP.NET is built on top of the Microsoft Internet Information Services (IIS) web server. ASP.NET AJAX builds on top of that and the web services it includes. The Microsoft AJAX Library runs in the browser to manipulate the DOM, communicate asynchronously with the web server, and take advantage of ASP.NET services.

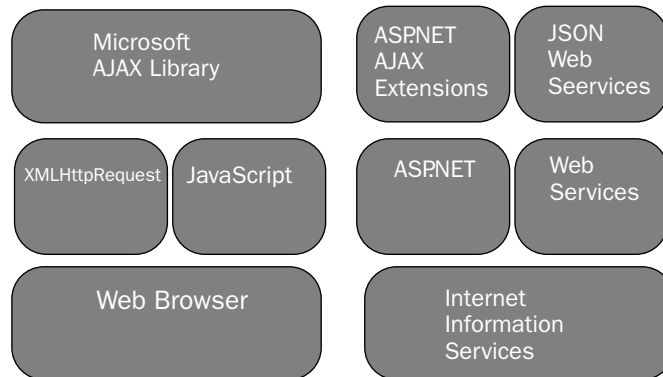


Figure 1-3

Installing ASP.NET AJAX

The ASP.NET AJAX Extensions are an add-on to ASP.NET 2.0. There is a link to the installer available at <http://ajax.asp.net/downloads/default.aspx?tabid=47>. When you run the installer, it will install the `System.Web.Extensions.dll` assembly into the Global Assembly Cache (gac) of your machine. Being in the gac allows the assembly to assert the permissions necessary to provide the full range of functionality ASP.NET AJAX offers. If Visual Studio 2005 or Visual Web Developer is present on the machine, the installer will also set up ASP.NET AJAX project templates. This allows you to create a new ASP.NET AJAX-enabled website. The `web.config` file and compilation references will have everything necessary to start building AJAX applications. The toolbox will include the `UpdatePanel` control discussed in Chapter 2, and the `ScriptManager` control covered in Chapter 5. The details of what is included in the new `web.config` file are discussed in detail in Chapter 10.

The ASP.NET Component

The original (ASP) Active Server Pages technology was released as part of Internet Information Server 3.0. It was then improved with support for transactions and access to COM objects with the release of the NT 4.0 Option Pack almost ten years ago. At the time, most websites consisted of static HTML pages. To the extent that there were any dynamic applications, they were CGI (Common Gateway Interface) or ISAPI (Internet Server API) applications typically written in C and C++. With the release of Active Server Pages (ASP), developers could use JavaScript or VBScript on the server to leverage a set of intrinsic objects provided by ASP. “Classic ASP,” as it is now called, provided a session object so that developers didn’t need to worry about the stateless nature of HTTP. It had `Request` and `Response` objects that provided for easy access to data coming from forms on the client and a way to send updated information back. The `Server` and `Application` objects were an avenue for accessing information from the web server and utilizing a common set of COM objects across the application.

Although Classic ASP was a big win for developers, developers soon found that developing complex applications became difficult. The lack of support for modularization meant that applications ended up with such complex interdependencies in script that developers often referred to it as “spaghetti code.” Debugging support was severely lacking, and there were no rich design tools for developing applications. Performance also suffered because ASP used a slow scripting model, and syntax errors could only be detected at runtime. ASP.NET is a big leap forward. It provides compatibility for existing applications by including the same set of intrinsic objects found in Classic ASP but moves from interpreting JavaScript and VBScript on the fly to using a compiled set of pages and modules written in C# and VB.NET. Classic ASP simply executed script in order from the top to the bottom of the page; ASP.NET is an event-driven model with a page lifecycle, making it more like writing a desktop application. Now, instead of including separate JavaScript files to represent business objects, you can create objects in any .NET language and access them directly from ASP.NET.

ASP.NET takes a set of pages that contain code and markup and generates a `Page` class that is then compiled and cached. For each request to the page, the class is instantiated and a complete page lifecycle is executed. A set of events are executed, some of which have been overridden by the generated `Page` class. Controls in the page also participate in the lifecycle, databinding to backend databases, reacting to user input, and dealing with changes to their state from the user’s previous view. To react to user actions, the developer has only to provide an event handler for the given action. For example, the button control exposes a `Click` event. When using it, you don’t need to write code to examine all form variables on a page to know if the button was clicked. Instead, you just provide code for an event handler override. The event handler code can then update the HTML for the page or the properties and data of other controls on the page.

*This book assumes some familiarity with ASP.NET as you look at how the new features of ASP.NET AJAX extend its functionality and provide ways to enhance web applications. If you are new to ASP.NET, I recommend looking at the set of QuickStart tutorials available at <http://www.asp.net/quickstart>. Other excellent resources are *Beginning ASP.NET 2.0 (VB.NET)*, *Beginning ASP.NET 2.0 with C#*, and *Professional ASP.NET 2.0*, all from Wrox Press. The ASP.NET website (<http://forums.asp.net/>) also has an extensive set of discussion forums where you can interact with fellow ASP.NET developers as well as members of the ASP.NET team.*

The JavaScript Component

AJAX technologies take advantage of the common support for JavaScript found in modern browsers. Because there is a standard that is supported across the various browsers, you can write scripts knowing that they will run. This wasn’t always the case.

Chapter 1: Developing Next-Generation Web Applications

In the mid 1990's, Netscape and Microsoft (along with others) collaborated on a standard for a scripting language that they would support in their web browsers. The standard is called EcmaScript. Microsoft's implementation is called JScript, but the language is generally referred to as JavaScript, as it was called in Netscape. (It has nothing to do with Java, but someone must have thought the association was useful for marketing purposes.) JavaScript program snippets are sent down to the browser along with the HTML, and they run inside the user's browser to affect how the page is processed on the client.

JavaScript is not compiled; it is interpreted. There is no static type-checking like you get in C++ and C#. You can declare a variable without needing to specify a type, and the type to which the variable refers can change at any time. This makes it easy to get started programming in JavaScript, but there's inevitably a certain amount of danger in allowing the data type of a variable to change dynamically at runtime. In the following snippet, notice that the variable can reference any type without difficulty:

```
var something = 1;
something = true;
something = "a string";
```

JavaScript is a dynamic language. Types can actually be extended during program execution by other code. This means that you can write code that creates types on the fly. Because there is no enforcement of type safety, your code can receive these types as parameters or return values without any problem. This provides a great degree of flexibility and coding power.

The fundamental types in JavaScript are strings, numbers, Booleans, and functions. There is also support for objects and arrays, which are collections of the fundamental types. Some additional objects are included that are considered essential for many programming tasks. This includes support for regular expressions and date and time operations.

You can use the plus operator on strings in JavaScript to concatenate them:

```
var theEnd = "THE END.";
var result = "Beginning, " + "middle, and " + theEnd;
```

In this example, the result variable is now the string: "Beginning, middle, and THE END."

JavaScript interpreters use the IEEE floating-point standard for storing numbers. Ignoring the gory details, you can assume that for most programming tasks you won't have any trouble.

The Boolean type in JavaScript is about what you would expect it to be but maybe not exactly so. The Boolean represents whether or not an expression is true, but it uses the C-style convention using integer values 0 and 1.

Variables can exist in JavaScript without having a value, and a variable may simply be undefined, which can produce unexpected results. In this piece of JavaScript, three variables are declared, and all of these comparisons are designed to return a true value.

```
<script type="text/javascript">
var one = 1;
var zero = 0;
var undefinedVar;

if(one) {
```

Chapter 1: Developing Next-Generation Web Applications

```
    alert("1 is true");
}

if(!zero) {
    alert("0 is false");
}

if(!undefinedVar) {
    // this test tells us that "undefinedVar" either contains 0,
    // or is really undefined: both of which equate to false
    alert("undefinedVar is false");
}

if(one !== zero) {
    alert("one and zero are not the same");
}
</script>
```

You can check specifically to see if a variable has been defined like this:

```
if( typeof(undefinedVar) == "undefined" ) {
    alert("undefinedVar is undefined");
}
```

Variables can also have a null value, which is not the same thing as being undefined, as a null value does constitute a value.

Functions are also real types in JavaScript. They can accept arguments and return values. Functions can be passed to other functions and can be created dynamically by other script code.

Here are two equivalent definitions for a function called `Add` that will take two variables and return the result of applying the plus operator. Notice that I didn't state that it takes two numbers. Remember, JavaScript variables don't have a defined type, so I could just as easily pass two strings and get them concatenated by my `Add` function.

```
<script type="text/javascript">
function Add(x, y) {
    return x + y;
}
var AddAgain = function(x, y) { return x + y; }
</script>
```

Once either of these styles is used to create a function, it can be called from that scope and any nested scope to perform the addition. There is no advantage to one of these forms over the other. You can simply choose to use the syntax that you prefer.

```
<script type="text/javascript">
var result = Add(36, 24);
alert(result); //displays 60

var stringResult = Add("Hello ", "there.");
alert(stringResult); //displays "Hello there."
</script>
```

Objects and arrays are just collections of other types. Array types do not require that the values they hold be named; instead, you can access them by index. The values held in an object are referenced by field or property names. Objects can also hold functions (which can be accessor functions to give public visibility to local variables), which lets you create data structures that represent entities in JavaScript code. Missing from this sort of object-oriented programming is a concept of type inheritance. The Microsoft AJAX Library provides a set of classes and recommended patterns for achieving inheritance in JavaScript, making it more natural for switching between JavaScript and other high-level languages. The following code example includes a definition for an `Album` object that holds and returns the artist and album title. An array is then used to store information about several albums.

```
<script type="text/javascript">
// define an object named Album - note that this object is typeless
Album = function(title, artist) {
    var _title = title;
    var _artist = artist;

    this.get_title = function() { return _title; }
    this.get_artist = function() { return _artist; }
}

// create object instances by calling the constructor
var albumA = new Album("Rift", "Phish");
var albumB = new Album("A Picture of Nectar", "Phish");

// create an array to hold the instances (also typeless)
var albumArray = new Array();

albumArray[0] = albumA;
albumArray[1] = albumB;

// iterate over the array to show the album titles
for(var i = 0; i < albumArray.length; i++) {
    alert((albumArray[i]).get_title()); // call get_title accessor
}
</script>
```

The Web Services Component

The fundamental concept of web services is powerful and continues to evolve and advance. The original SOAP (Simple Object Access Protocol) standard is the use of the HTTP protocol to pass XML-formatted data to the server from a client and receive XML-formatted results in return. This can be from within a web browser using the `XmlHttpRequest` object or directly from a desktop application or another server. Before web services became widely adopted, it was not uncommon for developers to programmatically request a page as an HTML document and extract the desired data from it, a technique known as *screen-scraping*. This causes all sorts of frustrations as sites are continually updated and the screen-scraping clients must try to keep up by modifying their parsing code to adapt to the new HTML the target site is rendering.

This produced frustration, as sites that presented data using HTML visual pages were prone to modifying those pages and this would break the screen-scraping program, which expected to see the data in the original format. Web services were created as a nonvisual way to transfer data over the web, and they are the natural way to isolate remote method calls from the presentation layer. Now, instead of screen-scraping, you are able to call a web service and get XML-formatted data returned that is easily consumed by a program.

Chapter 1: Developing Next-Generation Web Applications

By passing plain text data formatted as XML and by eliminating the visual elements, data passed in web services is much easier to parse than HTML. And, since XML can contain an embedded schema, code can inspect the schema and use it to determine the structure and types used in the data. You can extend the schema passed with the data being returned without worrying that consuming applications will be broken, and therefore XML readers can be somewhat tolerant of modifications that would have certainly caused a screen-scraper a great deal of grief!

The schema for data can be extended without requiring all consumers to be updated. Consumers can easily get the parts of the XML document they wish to process and disregard the rest. This has progressed beyond simple XML formats. Unlike previous implementations of web services, you can now define web service contracts to be built to employ arbitrary encoding and utilize any one of a number of wire protocols. What drives the web service concept is the ability to access data easily from various applications in a loosely coupled way, and the new Microsoft Windows Communication Foundation (WCF) takes this concept to a whole new level, allowing the contract to specify wire protocols, deployment strategies, and logging infrastructure, along with providing support for transactions.

ASP.NET AJAX provides a set of JavaScript proxy objects to access some new web services built-in to ASP.NET. Profile information, membership services, and role management can be easily accessed from the client. Developers don't need to create their own infrastructure to support these fundamental application services, but can include a few lines of code to take advantage of server resources from JavaScript code running in the browser, thereby dramatically extending the reach of ASP.NET to include both the client and the server. And because the JavaScript libraries are designed to be easy to use for developers already familiar with server-side .NET programming, all of this extra functionality comes in a friendly package that is easy to leverage.

The Dynamic HTML Component

Dynamic HTML is not a freestanding technology. It is the use of a set of technologies in a specific way. HTML is returned to the browser following a web server request. The browser then renders the page, and the user is able to view it. The browser also exposes the Document Object Model (DOM) that represents the structure of the HTML being displayed. The DOM can be accessed from JavaScript embedded in, or referenced by, the page. The appearance of the HTML is affected by applying CSS (Cascading Style Sheets) that control colors, fonts, position, visibility, and more. You can bind JavaScript code to events that the browser will raise when users perform certain actions, such as hovering over a particular element or entering text in a textbox. The JavaScript code can update text or manipulate the CSS settings for elements within the page. And, of course, it can communicate with the server to expand the dynamic nature of the page even further. The user will see a dynamically changing user interface that responds to his actions in real-time, which will greatly enhance his overall experience, thereby increasing his productivity and satisfaction with the application.

Other AJAX Libraries

In addition to ASP.NET AJAX, many third-party AJAX libraries are available that can be used with ASP.NET, although not all of them were specifically designed for it. Some are mostly focused on providing JavaScript libraries for use from within the browser to make manipulation of the browser DOM (Document Object Model) easier. Others include some level of server functionality for use within ASP.NET pages (where server controls will render on the client side). This section briefly highlights some of what these libraries offer. The ASP.NET AJAX Framework can coexist with script and controls

from other libraries, although given the dynamic nature of the JavaScript language, it is possible to extend types so that they conflict with each other. Mixing and matching libraries might work just fine for many uses, but you might find conflicts in other cases.

- ❑ **Ajax.NET Professional:** Michael Schwartz developed Ajax.NET Professional as a tool primarily used to simplify the data transport mechanism that enables a client JavaScript routine to communicate with a server program. The server code is simple to use: You merely need to register the control in your page and decorate some code-behind methods with attributes to designate which ones can be called from the client. Then you can leverage the script library to make the calls and pass data. This is intended for developers who are well versed with DHTML, and there aren't many prebuilt visual controls. This is a lightweight solution with very little overhead in terms of bytes transferred and processing cycles needed on the client and server. The source code is available, and the package is free (<http://www.ajaxpro.info>).
- ❑ **Anthem.NET:** Anthem.NET is a SourceForge project where users are able to download the sources to the project. It targets ASP.NET 1.1 and ASP.NET 2.0. It has a set of server controls that use their underlying JavaScript library to communicate with the server. They provide the ability to access the state of controls on the page during an asynchronous callback. At the time of writing, the Anthem.NET web page (<http://anthem-dot-net.sourceforge.net>) points out that the Anthem.NET user needs to be an experienced ASP.NET developer to get the most out of it. However, this is generally easier to use than Ajax.NET Professional, especially for developers who aren't well-versed in DHTML. This project is similar to ASP.NET Ajax in many ways but isn't as comprehensive.
- ❑ **DoJo:** The DoJo toolkit can be found at <http://dojotoolkit.com>. It is a client-side library for AJAX development without ties to any server technology. DoJo has a type system for JavaScript and a function for binding script to events from JavaScript objects or DHTML elements. One of its strengths is rich support for dynamic script loading. You can specify dependencies and ordering in the way that scripts are retrieved and processed.
- ❑ **Prototype:** The Prototype script library is available at <http://prototype.conio.net>. It does not target any server technology for integration. It has a type system for scripting in a more object-oriented way, along with some shortcut syntaxes for dealing with JavaScript arrays as well as accessing and manipulating HTML elements on the page. Prototype provides networking functionality and a method for automatically updating an HTML element with the results of an HTTP request when given a URL. The Prototype library also has functions for associating script objects and methods with DOM objects and events. The library is focused on simplifying tasks that can be cumbersome and tedious. It doesn't provide much help for producing a richer user interface but puts forth the building blocks for an improved web-scripting experience.
- ❑ **Script.aculo.us:** The Script.aculo.us library can be found at the website of the same name: <http://script.aculo.us>. Their tagline is "it's about the user interface, baby!" which accurately describes their focus. Script.aculo.us is built on top of the Prototype library and picks up where it stops. It includes functionality for adding drag-and-drop support to an application. It has a lot of effects code for fading, shrinking, moving, and otherwise animating DOM elements. Script.aculo.us also has a slider control and a library for manipulating lists of elements.
- ❑ **Rico:** The Rico library also builds on top of the Prototype system. It has support for adding drag-and-drop behavior to browser DOM elements. It also has some controls to bind JavaScript objects to DOM elements for manipulating data. Rico has constructs for revealing and hiding portions of a page using an accordion style. It also has animation, sizing, and fading effects pre-built for easier use. These UI helpers are available at <http://openrico.org>.

Balancing Client and Server Programming with ASP.NET AJAX

Without the advanced use of JavaScript running in the browser, web applications have their logic running on the server. This means a lot of page refreshes for potentially small updates to the user's view. With AJAX, much of the logic surrounding user interactions can be moved to the client. This presents its own set of challenges. Some examples of AJAX use include streaming large datasets to the browser that are managed entirely in JavaScript. While JavaScript is a powerful language, the debugging facilities and options for error handling are very limited. Putting complex application logic on the client can take a lot of time, effort, and patience. ASP.NET AJAX allows you to naturally migrate some parts of the application processing to the client while leveraging partial page rendering to let the server control some aspects of the page view.

Some websites make an application run entirely from a single page request, where JavaScript and AJAX will do a great deal of work. This presents some tough challenges. Users generally expect that the Back button will take them to the state of the application they were just viewing, but with AJAX applications this is not necessarily the case. The client may be sending some information to the server for persistent state management (perhaps in server memory or a database), but this requires extra code and special attention to error handling and recovery.

The richest, most maintainable applications seem to be those that balance client and server resources to provide quick response times, easy access to server resources, and a minimum of blocking operations while new page views are fetched.

ASP.NET AJAX provides a mix of client and server programming features. The Microsoft AJAX Library is aimed at client development. It provides a type system for an object-oriented approach to JavaScript development. It makes it easy to register code to respond to events. It provides useful functions to simplify common tasks like finding elements on the page, attaching event handlers, and accessing the server. The server features include functionality for managing JavaScript code to be sent to the client, declaring regions of the page to be updated asynchronously, creating timers for continuous updates, and accessing ASP.NET services such as user profile data and authentication.

Summary

The web has evolved over the last decade from providing a static presence to being the default choice for developers writing applications. With web applications, you get reach without having to deal with deployment and servicing issues that accompany desktop applications. But the bar continues to move higher for web applications as users come to expect more. AJAX technologies are driving web applications to rival rich desktop apps. You can use the results of asynchronous communication with the web server to update portions of the page without forcing the user to stop his work and wait for the page to post back and repaint. Dynamic HTML allows you to create a rich GUI with transitions and animations leveraging CSS for colors, fonts, positioning, and more.

ASP.NET AJAX includes the Microsoft AJAX Library that makes writing browser-based JavaScript easier and simplifies many common client programming tasks. It is easy to attach code to DOM events, write JavaScript in an object-oriented way, and access the server for persistent storage, authentication, and updated data.

ASP.NET AJAX also includes extensions to version 2.0 of the .NET Framework that can greatly improve your web application. There is built-in support for returning data in the JSON format that is easily consumed by JavaScript in the browser.

In this book, you will see how the client and server features of ASP.NET AJAX make it easier to push the limits of what you can do with a web application! You will learn how to update portions of a page asynchronously and how to manage the scripts that are used in the browser. You'll find out how to use the networking facilities, with a dive into accessing ASP.NET services like authentication and profile storage. You'll get a closer look at the JavaScript language and how the Microsoft AJAX Library builds on top of the language to simplify programming tasks. You'll also see what ASP.NET AJAX offers for adding richer UI to web applications and look at how to debug and test web applications. And you'll learn about some details of deploying and debugging applications and look at other resources available for use with ASP.NET AJAX.

