


CHAPTER 5

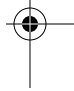

Java Servlets

Over the last few years, Java has become the predominant language for server-side programming. This is due in no small part to the Java Servlet API, which provides a standard way to extend web servers to support dynamic content generation. With the introduction of the J2EE specification for enterprise applications, servlets have taken over as the primary interface for thin-client applications. In terms of enterprise computing, servlets are a natural fit if you are using the Web as your deployment platform. You can take advantage of web browsers as universally available thin clients using the web server as middleware for running application logic. Under this model, the user makes a request of the web server, the server invokes a servlet designed to handle the request, the servlet fulfills the request, and the result is returned to the user for display in the web browser.

While this sounds like every other dynamic content technology (such as CGI, ISAPI, ASP, PHP, and the like), servlets have some major advantages. For one, servlets are persistent between invocations, which dramatically improves performance relative to CGI-style programs. Servlets are also 100% portable across operating systems and servers, unlike any of the alternatives. Finally, servlets have access to all the APIs of the Java platform, so, for example, it is easy to create a servlet that interacts with a database, using the JDBC API.

The first edition of this book, which covered Versions 2.0 and 2.1 of the Servlets API, focused on servlets as a replacement for other dynamic content technologies. During the following two years, there have been two additional revisions, the latest, Version 2.3, being finalized in September 2001 after a lengthy draft period. The new APIs integrate the Servlet API much more closely with the J2EE environment, introducing an explicit concept of a “web application.” This is a collection of static content, servlets, JavaServer pages, and configuration information that can be easily deployed as a single unit (and can easily coexist with other web applications on the same web server). Version 2.3 of the Servlet API is a required component of J2EE Version 1.3.







This chapter demonstrates the basic techniques used to write servlets using Versions 2.2 and 2.3 of the Java Servlet API, including some common web-development tasks such as cookie manipulation and session tracking. This chapter assumes that you have some experience with web development; if you are new to web development, you may want to brush up on web basics by consulting *Webmaster in a Nutshell*, by Stephen Spainhour and Robert Eckstein (O'Reilly). For a more complete treatment of servlets, check out *Java Servlet Programming* by Jason Hunter with William Crawford (O'Reilly).

Getting a Servlet Environment

You need a *servlet container* to run servlets. A servlet container uses a Java Virtual Machine* to run servlet code as requested by a web server. The servlet container is also responsible for managing other aspects of the servlet lifecycle: user sessions, class loading, servlet contexts (which we will discuss in the next session), servlet configuration information, servlet persistence, and temporary storage.

There are a few varieties of servlet containers. Some are simply add-ons to existing web servers. The most popular of these is Apache/JServ, a Servlets 2.0 container that adds servlet capability to the Apache Web Server. Since interaction with servlets occurs almost exclusively through a web browser, these servlet engines aren't useful on their own. Other servlet engines include embedded web servers. Mortbay.com's Jetty server and IBM's WebSphere product line fall into this category. Finally, there are servlet engines that can be used either as standalone web servers or connected to other servers, (for example, the Tomcat server from the Apache Jakarta Project).



Because Tomcat is the reference implementation for the Servlet API, and Tomcat 4.0 is the only 2.3-compliant container available at press time, all the examples in this chapter have been tested with it. Since Tomcat falls under the Apache umbrella, distribution is free, and you can download a copy (including, if you like, full source code) from <http://jakarta.apache.org>. Binary installations are available for Windows and several Unix flavors. Other 2.3-compatible containers should be available by press time, but since 2.2 containers (including Tomcat 3.x and all current commercially available J2EE environments) will likely be around for quite some time, the relatively small differences between 2.2 and 2.3 are noted throughout this chapter.

Servlet Basics

The Servlet API consists of two packages, `javax.servlet` and `javax.servlet.http`. The `javax` is there because servlets are a standard extension to Java, rather than a mandatory part of the API. This means that while servlets are official Java, Java virtual machine developers aren't required to include the classes for them in their Java development and execution environments. As mentioned already, however, the Servlet API is required for J2EE 1.3

* As of Servlets 2.3, a JDK 1.2 or higher JVM is required.

The Servlet Lifecycle

When a client makes a request involving a servlet, the server loads and executes the appropriate Java classes. Those classes generate content, and the server sends the content back to the client. In most cases, the client is a web browser, the server is a web server, and the servlet returns standard HTML. From the web browser's perspective, this isn't any different from requesting a page generated by a CGI script, or, indeed, standard HTML. On the server side, however, there is an important difference: persistence.* Instead of shutting down at the end of each request, the servlet can remain loaded, ready to handle subsequent requests. Figure 5-1 shows how this all fits together.

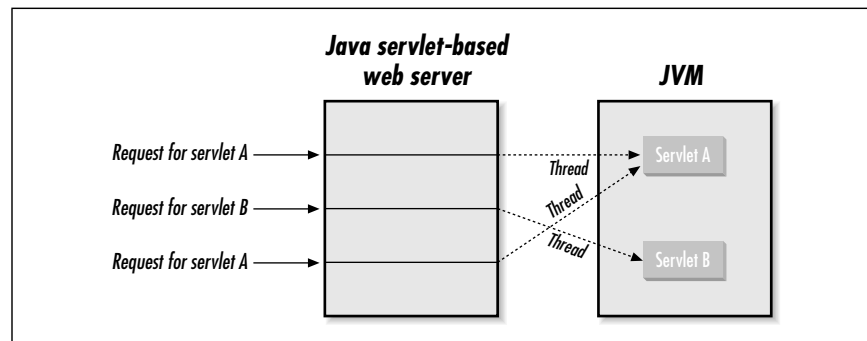


Figure 5-1: The servlet lifecycle

The request-processing time for a servlet can vary, but it is typically quite fast when compared to a similar CGI program. The real performance advantage of a servlet, however, is that you incur most of the startup overhead only once. When a servlet loads, its `init()` method is called. You can use `init()` to create I/O-intensive resources, such as database connections, for use across multiple invocations. If you have a high-traffic site, the performance benefits can be quite dramatic. Instead of putting up and tearing down a hundred thousand database connections, the servlet just needs to create a connection once.

After the `init()` method runs, the servlet container marks the servlet as available. For each incoming connection directed at a particular servlet, the container calls the `service()` method on the servlet to process the request. The `service()` method can have access to all the resources created in the `init()` method. The servlet's `destroy()` method is called to clean up resources when the server shuts down.

Because servlets are persistent, you can actually remove a lot of filesystem and/or database accesses altogether. For example, to implement a page counter, you can simply store a number in a static variable, rather than consulting a file (or database) for every request. Using this technique, you need to read and write to the disk only occasionally to preserve state. Since a servlet remains active, it can

* Note that we use persistent to mean "enduring between invocations," not "written to permanent storage."

perform other tasks when it is not servicing client requests, such as running a background processing thread (i.e., where clients connect to the servlet to view a result) or even acting as an RMI host, enabling a single servlet to handle connections from multiple types of clients. For example, if you write an order processing servlet, it can accept transactions from both an HTML form and an applet using RMI.

The Servlet API includes numerous methods and classes for making application development easier. Most common CGI tasks require a lot of fiddling on the programmer's part; even decoding HTML form parameters can be a chore, to say nothing of dealing with cookies and session tracking. Libraries exist to help with these tasks, but they are, of course, decidedly nonstandard. You can use the Servlet API to handle most routine tasks, thus cutting development time and keeping things consistent for multiple developers on a project.

Writing Servlets

The three core elements of the Servlet API are the `javax.servlet.Servlet` interface, the `javax.servlet.GenericServlet` class, and the `javax.servlet.http.HttpServlet` class. Normally, you create a servlet by subclassing one of the two classes, although if you are adding servlet capability to an existing object, you may find it easier to implement the interface.

The `GenericServlet` class is used for servlets that don't implement any particular communication protocol. Here's a basic servlet that demonstrates servlet structure by printing a short message:

```
import javax.servlet.*;
import java.io.*;

public class BasicServlet extends GenericServlet {

    public void service(ServletRequest req, ServletResponse resp)
        throws ServletException, IOException {

        resp.setContentType("text/plain");
        PrintWriter out = resp.getWriter();
        // We won't use the ServletRequest object in this example
        out.println("Hello.");
    }
}
```

`BasicServlet` extends the `GenericServlet` class and implements one method: `service()`. Whenever a server wants to use the servlet, it calls the `service()` method, passing `ServletRequest` and `ServletResponse` objects (we'll look at these in more detail shortly). The servlet tells the server what type of response to expect, gets a `PrintWriter` from the response object, and transmits its output.

The `GenericServlet` class can also implement a *filtering servlet* that takes output from an unspecified source and performs some kind of alteration. For example, a filter servlet might be used to prepend a header, scan servlet output or raw HTML files for `<DATE>` tags and insert the current date, or remove `<BLINK>` tags. A more advanced filtering servlet might insert content from a database into HTML

templates. We'll talk a little more about filtering later in this chapter, as well as discuss additional content filtering support available with 2.3 containers.

Although most servlets today work with web servers, there's no requirement for that in `GenericServlet`; the class implements just that, a generic servlet. As we'll see in a moment, the `HttpServlet` class is a subclass of `GenericServlet` that is designed to work with the HTTP protocol. It is entirely possible to develop other subclasses of `GenericServlet` that work with other server types. For example, a Java-based FTP server might use servlets to return files and directory listings or perform other tasks, although this capability has in general been underutilized. Later versions of the API have increased the coupling between servlets and HTTP.

HTTP Servlets

The `HttpServlet` class is an extension of `GenericServlet` that includes methods for handling HTTP-specific data.* `HttpServlet` provides a number of methods, such as `doGet()`, `doPost()`, and `doPut()`, to handle particular types of HTTP requests (GET, POST, and so on). These methods are called by the default implementation of the `service()` method, which figures out what kind of request is being made and then invokes the appropriate method. Here's a simple `HttpServlet`:

```
import javax.servlet.*;
import javax.servlet.http.*;
import java.io.*;

public class HelloWorldServlet extends HttpServlet {

    public void doGet(HttpServletRequest req, HttpServletResponse resp)
        throws ServletException, IOException {


        resp.setContentType("text/html");
        PrintWriter out = resp.getWriter();

        out.println("<HTML>");
        out.println(
            "<HEAD><TITLE>Have you seen this before?</TITLE></HEAD>");
        out.println(
            "<BODY><H1>Hello, World!</H1><H2>Again.</H2></BODY></HTML>");
    }
}
```

`HelloWorldServlet` demonstrates many essential servlet concepts. First, `HelloWorldServlet` extends `HttpServlet`. This is standard practice for an HTTP servlet. `HelloWorldServlet` defines one method, `doGet()`, which is called whenever anyone requests a URL that points to this servlet.† The `doGet()` method is actually called by the default `service()` method of `HttpServlet`. The `service()`


* `HttpServlet` is an abstract class, implemented by the provider of the servlet container.

† In a standard Java Web Server installation, with the servlet installed in the standard `servlets` directory, this URL is `http://site:8080/servlet/HelloWorldServlet`. Note that the name of the directory (`servlets`) is unrelated to the use of "servlet" in the URL.




method is called by the web server when a request is made of `HelloWorldServlet`; the method determines what kind of HTTP request is being made and dispatches the request to the appropriate `doXXX()` method (in this case, `doGet()`). `doGet()` is passed two objects, `HttpServletRequest` and `HttpServletResponse`, that contain information about the request and provide a mechanism for the servlet to produce output, respectively.

The `doGet()` method itself does three things. First, it sets the output type to `text/html`, which indicates that the servlet produces standard HTML as its output. Second, it calls the `getWriter()` method of the `HttpServletResponse` parameter to get a `java.io.PrintWriter` that points to the client. Finally, it uses the stream to send some HTML back to the client. This isn't really a whole lot different from the `BasicServlet` example, but it gives us all the tools we'll need later on for more complex web applications. We do have to explicitly set the content type, as there is no default setting, even for HTTP servlets where one might reasonably expect `text/html`.



If you define a `doGet()` method for a servlet, you may also want to override the `getLastModified()` method of `HttpServlet`. The server calls `getLastModified()` to find out if the content delivered by a servlet has changed. The default implementation of this method returns a negative number, which tells the server that the servlet doesn't know when its content was last updated, so the server is forced to call `doGet()` and return the servlet's output. If you have a servlet that changes its display data infrequently (such as a servlet that verifies uptime on several server machines once every 15 minutes), you should implement `getLastModified()` to allow browsers to cache responses. `getLastModified()` should return a long value that represents the time the content was last modified as the number of milliseconds since midnight, January 1, 1970, GMT. This number can be easily obtained by calling the `getTime()` method `java.util.Date`.



A servlet should also implement `getServletInfo()`, which returns a string that contains information about the servlet, such as name, author, and version (just like `getAppletInfo()` in applets). This method is called by the web server and generally used for logging purposes.

Web Applications

Now that we've seen a basic servlet, we can step back for a moment and talk about how servlets are integrated into the servlet container. Version 2.2 of the Servlet API popularized the concept of a web application installed within a web server. A web application consists of a set of resources, including servlets, static content, JSP files, and class libraries, installed within a particular path on a web server. This path is called the *servlet context*, and all servlets installed within the context are given an isolated, protected environment to operate in, without interference from (or the ability to interfere with) other software running on the server.

A servlet context directory tree contains several different types of resources. These include class files and JAR files (which aren't exposed to clients connecting via web browsers), JSP files (which are processed by the JSP servlet before being fed back to the client), and static files, such as HTML documents and JPEG images, which are served directly to the browser by the web server.

Finally, there is a virtual component to the context. For each context, the servlet container will instantiate separate copies of servlets (even if those servlets are shared) and will create a private address space that can be accessed via the `ServletContext` class. Servlets can use this class to communicate with other servlets running in the same context. We'll discuss this more later.

The simplest servlet installations will just create a single context, rooted at `/`, which is the top of the web server path tree. Servlets and static content will be installed within this context. This is the way the Servlet API 2.0 treated the entire server. More modern servlet containers allow the creation of multiple servlet contexts, rooted lower down on the directory tree. A catalog application, for example, could be rooted at `/catalog`, with all of the application paths below the context root.

If you write a web application that will be installed on multiple web servers, it isn't safe to assume the context root will be fixed. If the path of a resource within your application is `/servlet/CatalogServlet`, and it's installed within the `/catalog` context, rather than writing:

```
out.println("<a href=\"/catalog/servlet/CatalogServlet\">");
```

you should write:

```
out.println("<a href=\"" + request.getContextPath() + "/servlet/  
CatalogServlet\">");
```

This approach works regardless of the context path installed within the web server.

Structure of Web Applications

On disk, a web application consists of a directory. The directory contains a subdirectory called *WEB-INF*, and whatever other content is required for the application. The *WEB-INF* directory contains a classes directory (containing application code), a lib directory (containing application JAR files), and a file called *web.xml*. The *web.xml* file contains all of the configuration information for the servlets within the context, including names, path mappings and initialization parameters and context-level configuration information. For a detailed explanation of *web.xml*, consult your server documentation or take a look at the well-commented XML DTD itself at http://java.sun.com/dtd/web-app_2_3.dtd.

The procedure for installing a web application into a servlet container varies from product to product, but it generally consists of selecting a context root and pointing the server to the directory containing the web application.*

Mapping Requests with a Context

Servlets are installed within the servlet container and mapped to URIs. This is done either via global properties that apply to all servlets or by specific, servlet-by-servlet mappings. In the first case, a client invokes a servlet by requesting it by name. Most servers map servlets to a `/servlet/` or `/servlets/` URL. If a servlet is

* Web applications can be packaged into JAR file equivalents called WAR files. To do this, simply use the *jar* utility that comes with the JDK to pack up the web application directory (including the *WEB-INF* subdirectory) and give the resulting file a *.war* extension.

installed as `PageServlet`, then a request to `/servlet/PageServlet` would invoke it. Servlets can also be individually mapped to other URIs or to file extensions. `PageServlet` might be mapped to `/pages/page1`, or to all files with a `.page` extension (using `*.page`).

All of these mappings exist below the context level. If the web application is installed at `/app`, then the paths entered into the browser for the examples above would be `/app/servlet/PageServlet`, `/app/pages/page1`, or `/app/file.page`.

To illustrate, imagine the following servlet mappings (all are below the context root):

<i>Mapping</i>	<i>Servlet</i>
<code>/store/furniture/*</code>	<code>FurnitureServlet</code>
<code>/store/furniture/tables/*</code>	<code>TableServlet</code>
<code>/store/furniture/chairs</code>	<code>ChairServlet</code>
<code>*.page</code>	<code>PageServlet</code>

The asterisk serves as a wildcard. URIs matching the pattern are mapped to the specified servlet, providing that another mapping hasn't already been used to deal with the URL. This can get a little tricky when building complex mapping relationships, but the servlet API does require servers to deal with mappings consistently. When the servlet container receives a request, it always maps it to the appropriate servlet in the following order:

1. By *exact path matching*. A request to `/store/furniture/chairs` is served by `ChairServlet`.
2. By *prefix mapping*. A request to `/store/furniture/sofas` is served by `FurnitureServlet`. The longest matching prefix is used. A request to `/store/furniture/tables/dining` is served by `TableServlet`.
3. By *extension*. Requests for `/info/contact.page` are served by `PageServlet`. However, requests for `/store/furniture/chairs/about.page` is served by `FurnitureServlet` (since prefix mappings are checked first, and `ChairServlet` is available only for exact matches).

If no appropriate servlet is found, the server returns an error message or attempts to serve content on its own. If a servlet is mapped to the `/` path, it becomes the default servlet for the application and is invoked when no other servlet can be found.

Context Methods

Resources within a servlet context (such as HTML files, images, and other data) can be accessed directly via the web server. If a file called `index.html` is stored at the root of the `/app` context, then it can be accessed with a request to `/app/index.html`. Context resources can also be accessed via the `ServletContext` object, which is accessed via the `getResource()` and `getResourceAsStream()` methods. A full list of available resources can be accessed via the `getResourcePaths()` method. In this case, an `InputStream` containing the contents of the `index.html` file can be retrieved by calling `getResourceAsStream("/index.html")` on the `ServletContext` object associated with the `/app` context.

The `ServletContext` interface provides servlets with access to a range of information about the local environment. The `getInitParameter()` and `getInitParameterNames()` methods allow servlets to retrieve context-wide initialization parameters. `ServletContext` also includes a number of methods that allow servlets to share attributes. The new `setAttribute()` method allows a servlet to set an attribute that can be shared by any other servlets that live in its `ServletContext`, and `removeAttribute()` allows them to be removed. The `getAttribute()` method, which previously allowed servlets to retrieve hardcoded server attributes, provides access to attribute values, while `getAttributeNames()` returns an `Enumeration` of all the shared attributes.


The servlet container is required to maintain a temporary working directory on disk for each servlet context. This directory is accessed by retrieving the `javax.servlet.context.tempdir` attribute, which consists of a `java.io.File` object pointing to the temporary directory. The temporary directory is exclusive to the context. The servlet container is not required to maintain its contents across restarts.

Version 2.1 of the Servlet API deprecated all methods related to accessing other servlets directly, due to the fact that they are inherently insecure. Thus, `getServlet()` and `getServletNames()` join the already deprecated `getServlets()`. The problem was that `getServlet()` incorrectly allowed one servlet to call another servlet's life-cycle methods, including `init()` and `destroy()`.

Servlet Requests

When a servlet is asked to handle a request, it typically needs specific information about the request so that it can process the request appropriately. Most frequently, a servlet will retrieve the value of a form variable and use that value in its output. A servlet may also need access to information about the environment in which it is running. For example, a servlet may need to find out about the actual user who is accessing the servlet, for authentication purposes.

The `ServletRequest` and `HttpServletRequest` interfaces provide access to this kind of information. When a servlet is asked to handle a request, the server passes it a request object that implements one of these interfaces. With this object, the servlet can determine the actual request (e.g., protocol, URL, type), access parts of the raw request (e.g., headers, input stream), and get any client-specific request parameters (e.g., form variables, extra path information). For instance, the `getProtocol()` method returns the protocol used by the request, while `getRemoteHost()` returns the name of the client host. The interfaces also provide methods that let a servlet get information about the server (e.g., `getServerName()`, `getServerPort()`). As we saw earlier, the `getParameter()` method provides access to request parameters such as form variables. There is also the `getParameterValues()` method, which returns an array of strings that contains all the values for a particular parameter. This array generally contains only one string, but some HTML form elements (as well as non-HTTP oriented services) do allow multiple selections or options, so the method always returns an array, even if it has a length of one.



HttpServletRequest adds a few more methods for handling HTTP-specific request data. For instance, `getHeaderNames()` returns an enumeration of the names of all the HTTP headers submitted with a request, while `getHeader()` returns a particular header value. Other methods exist to handle cookies and sessions, as we'll discuss later.

Example 5-1 shows a servlet that restricts access to users who are connecting via the HTTPS protocol, using Digest style authentication, and coming from a government site (a domain ending in `.gov`).

Example 5-1: Checking Request Information to Restrict Servlet Access

```
import javax.servlet.*;
import javax.servlet.http.*;
import java.io.*;

public class SecureRequestServlet extends HttpServlet {

    public void doGet(HttpServletRequest req, HttpServletResponse resp)
        throws ServletException, IOException {

        resp.setContentType("text/html");
        PrintWriter out = resp.getWriter();

        out.println("<HTML>");
        out.println("<HEAD><TITLE>Semi-Secure Request</TITLE></HEAD>");
        out.println("<BODY>");

        String remoteHost = req.getRemoteHost();
        String scheme = req.getScheme();
        String authType = req.getAuthType();

        if((remoteHost == null) || (scheme == null) || (authType == null)) {
            out.println("Request Information Was Not Available.");
            return;
        }

        if(scheme.equalsIgnoreCase("https") && remoteHost.endsWith(".gov")
            && authType.equals("Digest")) {
            out.println("Special, secret information.");
        }
        else {
            out.println("You are not authorized to view this data.");
        }

        out.println("</BODY></HTML>");
    }
}
```

Forms and Interaction

The problem with creating a servlet like `HelloWorldServlet` is that it doesn't do anything we can't already do with HTML. If we are going to bother with a servlet

at all, we should do something dynamic and interactive with it. In many cases, this means processing the results of an HTML form. To make our example less impersonal, let's have it greet the user by name. The HTML form that calls the servlet using a GET request might look like this:

```
<HTML>
<HEAD><TITLE>Greetings Form</TITLE></HEAD>
<BODY>
<FORM METHOD=GET ACTION="/servlet/HelloServlet">
What is your name?
<INPUT TYPE=TEXT NAME=username SIZE=20>
<INPUT TYPE=SUBMIT VALUE="Introduce Yourself">
</FORM>
</BODY>
</HTML>
```

This form submits a form variable named `username` to the URL `/servlet/HelloServlet`. The `HelloServlet` itself does little more than create an output stream, read the `username` form variable, and print a nice greeting for the user. Here's the code:

```
import javax.servlet.*;
import javax.servlet.http.*;
import java.io.*;

public class HelloServlet extends HttpServlet {

    public void doGet(HttpServletRequest req, HttpServletResponse resp)
        throws ServletException, IOException {

        resp.setContentType("text/html");
        PrintWriter out = resp.getWriter();

        out.println("<HTML>");
        out.println("<HEAD><TITLE>Finally, interaction!</TITLE></HEAD>");
        out.println("<BODY><H1>Hello, " + req.getParameter("username") +
            "!</H1>");
        out.println("</BODY></HTML>");
    }
}
```

All we've done differently is use the `getParameter()` method of `HttpServletRequest` to retrieve the value of a form variable.* When a server calls a servlet, it can also pass a set of request parameters. With HTTP servlets, these parameters come from the HTTP request itself—in this case, in the guise of URL-encoded form variables. Note that a `GenericServlet` running in a web server also has access to these parameters using the simpler `ServletRequest` object. When the `HelloServlet` runs, it inserts the value of the `username` form variable into the HTML output, as shown in Figure 5-2.

* In the Java Web Server 1.1, the `getParameter()` method was deprecated in favor of `getParameterValues()`, which returns a `String` array rather than a single string. However, after an extensive write-in campaign, Sun took `getParameter()` off the deprecated list for Version 2.0 of the Servlet API, so you can safely use this method in your servlets. This is not an issue with later versions of the API.

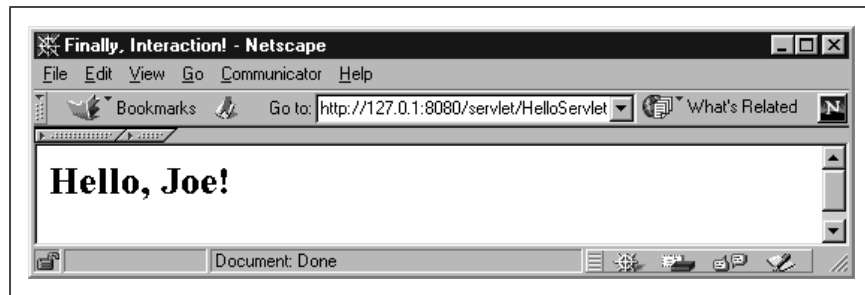


Figure 5-2: Output from HelloServlet

POST, HEAD, and Other Requests

As mentioned earlier, `doGet()` is just one of a collection of enabling methods for HTTP request types. `doPost()` is the corresponding method for POST requests. The POST request is designed for posting information to the server, although in practice it is also used for long parameterized requests and larger forms, to get around limitations on the length of URLs.

If your servlet is performing database updates, charging a credit card, or doing anything that takes an explicit client action, you should make sure this activity is happening in a `doPost()` method. That's because POST requests aren't *idempotent*, which means that they aren't safely repeatable, and web browsers treat them specially. For example, a browser can't bookmark or, in some cases, reload a POST request. On the other hand, GET requests are idempotent, so they can safely be bookmarked, and a browser is free to issue the request repeatedly without necessarily consulting the user. You can see why you don't want to charge a credit card in a GET method!

To create a servlet that can handle POST requests, all you have to do is override the default `doPost()` method from `HttpServlet` and implement the necessary functionality in it. If necessary, your application can implement different code in `doPost()` and `doGet()`. For instance, the `doGet()` method might display a postable data entry form that the `doPost()` method processes. `doPost()` can even call `doGet()` at the end to display the form again.

The less common HTTP request types, such as HEAD, PUT, TRACE, and DELETE, are handled by other `doXXX()` dispatch methods. A HEAD request returns HTTP headers only, PUT and DELETE allow clients to create and remove resources from the web server, and TRACE returns the request headers to the client. Since most servlet programmers don't need to worry about these requests, the `HttpServlet` class includes a default implementation of each corresponding `doXXX()` method that either informs the client that the request is unsupported or provides a minimal implementation. You can provide your own versions of these methods, but the details of implementing PUT or DELETE functionality go rather beyond our scope.

Servlet Responses

In order to do anything useful, a servlet must send a response to each request that is made to it. In the case of an HTTP servlet, the response can include three components: a status code, any number of HTTP headers, and a response body.

The `ServletResponse` and `HttpServletResponse` interfaces include all the methods needed to create and manipulate a servlet's output. We've already seen that you specify the MIME type for the data returned by a servlet using the `setContentType()` method of the response object passed into the servlet. With an HTTP servlet, the MIME type is generally `text/html`, although some servlets return binary data: a servlet that loads a GIF file from a database and sends it to the web browser should set a content type of `image/gif` while a servlet that returns an Adobe Acrobat file should set it to `application/pdf`.

`ServletResponse` and `HttpServletResponse` each define two methods for producing output streams, `getOutputStream()` and `getWriter()`. The former returns a `ServletOutputStream`, which can be used for textual or binary data. The latter returns a `java.io.PrintWriter` object, which is used only for textual output. The `getWriter()` method examines the content-type to determine which charset to use, so `setContentType()` should be called before `getWriter()`.

`HttpServletResponse` also includes a number of methods for handling HTTP responses. Most of these allow you to manipulate the HTTP header fields. For example, `setHeader()`, `setIntHeader()`, and `setDateHeader()` allow you to set the value of a specified HTTP header, while `containsHeader()` indicates whether a certain header has already been set. You can use either the `setStatus()` or `sendError()` method to specify the status code sent back to the server. `HttpServletResponse` defines a long list of integer constants that represent specific status codes (we'll see some of these shortly). You typically don't need to worry about setting a status code, as the default code is 200 ("OK"), meaning that the servlet sent a normal response. However, a servlet that is part of a complex application structure (such as the file servlet included in the Java Web Server that handles the dispatching of HTML pages) may need to use a variety of status codes. Finally, the `sendRedirect()` method allows you to issue a page redirect. Calling this method sets the `Location` header to the specified location and uses the appropriate status code for a redirect.

Request Dispatching

Request dispatching allows a servlet to delegate request handling to other components on the server. A servlet can either forward an entire request to another servlet or include bits of content from other components in its own output. In either case, this is done with a `RequestDispatcher` object that is obtained from the `ServletContext` via the `getRequestDispatcher()` method (also available via the `HttpServletRequest` object.) When you call this method, you specify the path to the servlet to which you are dispatching the request. The path should be relative to the servlet context. If you want to dispatch a request to the `/servlet/TargetServlet` URI within the `/app` context (which is accessed from a user's browser by `/app/servlet/TargetServlet`), request a dispatcher for `/servlet/TargetServlet`.

When you dispatch a request, you can set request attributes using the `setAttribute()` method of `ServletRequest` and read them using the `getAttribute()` method. A list of available attributes is returned by `getAttributeNames()`. All three of these methods were new in Version 2.1. Rather than taking only `String` objects (like parameters), an attribute may be any valid `Java` object.

`RequestDispatcher` provides two methods for dispatching requests: `forward()` and `include()`. To forward an entire request to another servlet, use the `forward()` method. When using `forward()`, the `ServletRequest` object is updated to include the new target URL. If a `ServletOutputStream` or `PrintWriter` has already been retrieved from the `ServletResponse` object, the `forward()` method throws an `IllegalStateException`.

The `include()` method of `RequestDispatcher` causes the content of the dispatchee to be included in the output of the main servlet, just like a server-side include. To see how this works, let's look at part of a servlet that does a keep-alive check on several different servers. The `ServerMonitorServlet` referenced in this example relies on the `serverurl` attribute to determine which server to display monitoring information for:

```
out.println("Uptime for our servers");

// Get a RequestDispatcher to the ServerMonitorServlet
RequestDispatcher d = getServletContext().
    getRequestDispatcher("/servlet/ServerMonitorServlet");

req.setAttribute("serverurl", new URL("http://www1.company.com"));
d.include(req, res);

req.setAttribute("serverurl", new URL("http://www2.company.com"));
d.include(req, res);
```

Error Handling

Sometimes things just go wrong. When that happens, it's nice to have a clean way out. The Servlet API gives you two ways of to deal with errors: you can manually send an error message back to the client or you can throw a `ServletException`. The easiest way to handle an error is simply to write an error message to the servlet's output stream. This is the appropriate technique to use when the error is part of a servlet's normal operation, such as when a user forgets to fill in a required form field.

Status codes

When an error is a standard HTTP error, you should use the `sendError()` method of `HttpServletResponse` to tell the server to send a standard error status code. `HttpServletResponse` defines integer constants for all the major HTTP status codes. Table 5-1 lists the most common status codes. For example, if a servlet can't find a file the user has requested, it can send a 404 ("File Not Found") error and

let the browser display it in its usual manner. In this case, we can replace the typical `setContentType()` and `getWriter()` calls with something like this:

```
response.sendError(HttpServletResponse.SC_NOT_FOUND);
```

If you want to specify your own error message (in addition to the web server's default message for a particular error code), you can call `sendError()` with an extra `String` parameter:

```
response.sendError(HttpServletResponse.SC_NOT_FOUND,
    "It's dark. I couldn't find anything.");
```

Table 5-1: Some Common HTTP Error Codes

Mnemonic Content	Code	Default Message	Meaning
SC_OK	200	OK	The client's request succeeded, and the server's response contains the requested data. This is the default status code.
SC_NO_CONTENT	204	No Content	The request succeeded, but there is no new response body to return. A servlet may find this code useful when it accepts data from a form, but wants the browser view to stay at the form. It avoids the "Document contains no data" error message.
SC_MOVED_PERMANENTLY	301	Moved Permanently	The requested resource has permanently moved to a new location. Any future reference should use the new location given by the <code>Location</code> header. Most browsers automatically access the new location.
SC_MOVED_TEMPORARILY	302	Moved Temporarily	The requested resource has temporarily moved to another location, but future references should still use the original URL to access the resource. The temporary new location is given by the <code>Location</code> header. Most browsers automatically access the new location.
SC_UNAUTHORIZED	401	Unauthorized	The request lacked proper authorization. Used in conjunction with the <code>WWW-Authenticate</code> and <code>Authorization</code> headers.
SC_NOT_FOUND	404	Not Found	The requested resource is not available.
SC_INTERNAL_SERVER_ERROR	500	Internal Server Error	An error occurred inside the server that prevented it from fulfilling the request.
SC_NOT_IMPLEMENTED	501	Not Implemented	The server doesn't support the functionality needed to fulfill the request.
SC_SERVICE_UNAVAILABLE	503	Service Unavailable	The server is temporarily unavailable, but service should be restored in the future. If the server knows when it will be available again, a <code>Retry-After</code> header may also be supplied.

Servlet exceptions

The Servlet API includes two `Exception` subclasses, `ServletException` and its derivative, `UnavailableException`. A servlet throws a `ServletException` to indicate a general servlet problem. When a server catches this exception, it can handle the exception however it sees fit.



UnavailableException is a bit more useful, however. When a servlet throws this exception, it is notifying the server that it is unavailable to service requests. You can throw an UnavailableException when some factor beyond your servlet's control prevents it from dealing with requests. To throw an exception that indicates permanent unavailability, use something like this:

```
throw new UnavailableException(this, "This is why you can't use the  
servlet.");
```

UnavailableException has a second constructor to use if the servlet is going to be temporarily unavailable. With this constructor, you specify how many seconds the servlet is going to be unavailable, as follows:

```
throw new UnavailableException(120, this, "Try back in two minutes");
```

One caveat: the servlet specification doesn't mandate that servers actually try again after the specified interval. If you choose to rely on this capability, you should test it first.

A file serving servlet

Example 5-2 demonstrates both of these error-handling techniques, along with another method for reading data from the server. FileServlet reads a pathname from a form parameter and returns the associated file. Note that this servlet is designed only to return HTML files. If the file can't be found, the servlet sends the browser a 404 error. If the servlet lacks sufficient access privileges to load the file, it sends an UnavailableException instead. Keep in mind that this servlet exists as a teaching exercise: you should not deploy it on your web server. (For one thing, any security exception renders the servlet permanently unavailable, and for another, it can serve files from the root of your hard drive.)

Example 5-2: Serving Files

```
import javax.servlet.*;  
import javax.servlet.http.*;  
import java.io.*;  
  
public class FileServlet extends HttpServlet {  
  
    public void doGet(HttpServletRequest req, HttpServletResponse resp)  
        throws ServletException, IOException {  
  
        File r;  
        FileReader fr;  
        BufferedReader br;  
        try {  
            r = new File(req.getParameter("filename"));  
            fr = new FileReader(r);  
            br = new BufferedReader(fr);  
            if(!r.isFile()) { // Must be a directory or something else  
                resp.sendError(resp.SC_NOT_FOUND);  
                return;  
            }  
        }  
    }  
}
```


Example 5-2: Serving Files (continued)

```
catch (FileNotFoundException e) {
    resp.sendError(resp.SC_NOT_FOUND);
    return;
}
catch (SecurityException se) { // Be unavailable permanently
    throw(new UnavailableException(this,
        "Servlet lacks appropriate privileges."));
}

resp.setContentType("text/html");
PrintWriter out = resp.getWriter();
String text;
while( (text = br.readLine()) != null)
    out.println(text);

br.close();
}
```

Custom Servlet Initialization

At the beginning of this chapter, we talked about how a servlet's persistence can be used to build more efficient web applications. This is accomplished via class variables and the `init()` method. When a server loads a servlet for the first time, it calls the servlet's `init()` method and doesn't make any service calls until `init()` has finished. In the default implementation, `init()` simply handles some basic housekeeping, but a servlet can override the method to perform whatever one-time tasks are required. This often means doing some sort of I/O-intensive resource creation, such as opening a database connection. You can also use the `init()` method to create threads that perform various ongoing tasks. For instance, a servlet that monitors the status of machines on a network might create a separate thread to periodically ping each machine. When an actual request occurs, the service methods in the servlet can use the resources created in `init()`. Thus, the status monitor servlet might display an HTML table with the status of the various machines. The default `init()` implementation is not a do-nothing method, so you should remember to always call the `super.init()` method as the first action in your own `init()` routines.*

The server passes the `init()` method a `ServletConfig` object, which can include specific servlet configuration parameters (for instance, the list of machines to monitor). `ServletConfig` encapsulates the servlet initialization parameters, which are accessed via the `getInitParameter()` and `getInitParameterNames()` methods. `GenericServlet` and `HttpServlet` both implement the `ServletConfig` interface, so these methods are always available in a servlet. (One task the default `init()` implementation does is store the `ServletConfig` object for these methods, which is

* Note that you no longer have to do this with Version 2.1 of the Servlet API. The specification has been changed so that you can simply override a no-argument `init()` method, which is called by the `GenericServlet` `init(ServletConfig)` implementation.

why it is important you always call `super.init()`.) Different web servers have different ways of setting initialization parameters, so we aren't going to discuss how to set them. Consult your server documentation for details.

Every servlet also has a `destroy()` method that can be overwritten. This method is called when, for whatever reason, a server unloads a servlet. You can use this method to ensure that important resources are freed, or that threads are allowed to finish executing unmolested. Unlike `init()`, the default implementation of `destroy()` is a do-nothing method, so you don't have to worry about invoking the superclass' `destroy()` method.

Example 5-3 shows a counter servlet that saves its state between server shut-downs. It uses the `init()` method to first try to load a default value from a servlet initialization parameter. Next the `init()` method tries to open a file named `/data/counter.dat` and read an integer from it. When the servlet is shut down, the `destroy()` method creates a new `counter.dat` file with the current hit-count for the servlet.

Example 5-3: A Persistent Counter Servlet

```
import javax.servlet.*;
import javax.servlet.http.*;
import java.io.*;

public class LifeCycleServlet extends HttpServlet {

    int timesAccessed;

    public void init(ServletConfig conf) throws ServletException {

        super.init(conf);

        // Get initial value
        try {
            timesAccessed = Integer.parseInt(getInitParameter("defaultStart"));
        }
        catch(NullPointerException e) {
            timesAccessed = 0;
        }
        catch(NumberFormatException e) {
            timesAccessed = 0;
        }

        // Try loading from the disk
        try {
            File r = new File("./data/counter.dat");
            DataInputStream ds = new DataInputStream(new FileInputStream(r));
            timesAccessed = ds.readInt();
        }
        catch (FileNotFoundException e) {
            // Handle error
        }
        catch (IOException e) {
            // This should be logged
        }
    }
}
```

Example 5-3: A Persistent Counter Servlet (continued)

```

    }
    finally {
        ds.close();
    }
}

public void doGet(HttpServletRequest req, HttpServletResponse resp)
    throws ServletException, IOException {

    resp.setContentType("text/html");
    PrintWriter out = resp.getWriter();

    timesAccessed++;

    out.println("<HTML>");
    out.println("<HEAD>");
    out.println("<TITLE>Life Cycle Servlet</TITLE>");
    out.println("</HEAD><BODY>");

    out.println("I have been accessed " + timesAccessed + " time[s]");
    out.println("</BODY></HTML>");
}

public void destroy() {

    // Write the Integer to a file
    File f = new File("../data/counter.dat");
    try {
        DataOutputStream dout = new DataOutputStream(new FileOutputStream(f));
        dout.writeInt(timesAccessed);
    }
    catch(IOException e) {
        // This should be logged
    }
    finally {
        dout.close();
    }
}
}

```

Servlet Context Initialization

Version 2.3 of the Servlet API adds support for application-level events via a listener-style interface. Classes that implement the `ServletContextListener` interface can be associated with a servlet context, and will be notified when the context is initialized or destroyed. This provides programmers with the opportunity to create application-level resources, such as database connection pools, before any servlets are initialized, and to share single resources among multiple servlets using the `ServletContext` attribute functionality.

`ServletContextListener` contains two methods, `contextInitialized()` and `contextDestroyed()`, which take a `ServletContextEvent`. Context listeners are

associated with their context in the *web.xml* file for the web application. Example 5-4 defines a listener that creates a hashtable of usernames and unencrypted passwords and associates it as a context attribute. We use it in a later example:

Example 5-4: A Servlet Context Listener

```
import javax.servlet.ServletContextListener;
import javax.servlet.ServletContextEvent;

public class ContextResourceLoader implements ServletContextListener {

    public void contextInitialized(ServletContextEvent sce) {
        java.util.Hashtable users = new Hashtable();
        users.put("test", "test");
        users.put("admin", "bob3jk");
        sce.getServletContext().setAttribute("enterprise.users", users);
    }

    public void contextDestroyed(ServletContextEvent sce) {
        // This is where we clean up resources on server shutdown/restart
    }
}
```

Obviously, a real application would retrieve the usernames and passwords in a more efficient manner. In this case, we can count on the JVM to properly garbage-collect the Hashtable object. If we do something more complex (such as maintaining a pool of connections to a relational database), we would use the `contextDestroyed()` method to make sure those resources were properly freed.

Security

Servlets don't have to handle their own security arrangements. Instead, they can rely on the capabilities of the web server to limit access where required. The security capabilities of most web servers are limited to basic on-or-off access to specific resources, controlled by username and password (or digital certificate), with possible encryption-in-transmission using SSL. Most servers are limited to basic authentication, which transmits passwords more or less in the clear, while some support the more advanced digest authentication protocol, which works by transmitting a hash of the user's password and a server-generated value, rather than the password itself. Both of these approaches look the same to the user; the familiar "Enter username and password" window pops up in the web browser.

Recent versions of the Servlet API take a much less hands-off approach to security. The *web.xml* file can be used to define which servlets and resources are protected, and which users have access. The user access model is the J2EE User-Role model, in which users can be assigned one or more Roles. Users with a particular role are granted access to protected resources. A user named Admin might have both the Administrator role and the User role, while users Bob and Ted might only have the User role.

In addition to basic, digest and SSL authentication, the web application framework allows for HTML form-based logins. This approach allows the developer to specify an HTML or JSP page containing a form like the following:

```
<form method="POST" action="j_security_check">  
<input type="text" name="j_username">  
<input type="password" name="j_password">  
<input type="submit" value="Log In">  
</form>
```

Note that forms-based authentication is insecure, and will only work if the client session is being tracked via Cookies or SSL signatures.

In Servlets 2.0, the `HttpServletRequest` interface included a pair of basic methods for retrieving standard HTTP user authentication information from the web server. If your web server is equipped to limit access, a servlet can retrieve the username with `getRemoteUser()` and the authentication method (basic, digest, or SSL) with `getAuthType()`. Version 2.2 of the Servlet API added the `isUserInRole()` and `getUserPrincipal()` methods to `HttpServletRequest`. `isUserInRole()` allows the program to query whether the current user is member of a particular role (useful for dynamic content decisions that can not be made at the container level). The `getUserPrincipal()` method returns a `java.security.Principal` object identifying the current user.

The process used to authenticate users (by validating their usernames and passwords) is up the developer of the servlet container.

Servlet Chains and Filters

So far, we have looked at servlets that take requests directly from the server and return their results directly to the client. Servlets were designed as a generic server extension technology, however, rather than one devoted solely to performing CGI-like functions. A servlet can just as easily take its input from another servlet, and a servlet really doesn't care very much about where its output goes.

Most web servers that implement servlets have also implemented a feature called *servlet chaining*, where the server routes a request through an administrator-defined chain of servlets. At the end of the sequence, the server sends the output to the client. Alternately, some servers can be configured to route certain MIME types through certain servlets. If a filtering servlet is configured to take all of the output with the MIME type "servlet/filterme," another servlet can produce data with that MIME type, and that data will be passed to the filtering servlet. The filtering servlet, after doing its work, can output HTML for the browser. MIME-based filtering also allows servlets to filter objects that don't come from a servlet in the first place, such as HTML files served by the web server.

Example 5-5 demonstrates a basic servlet, derived from `HttpServletRequest`, that examines incoming text for a `<DATE>` tag and replaces the tag with the current date. This servlet is never called on its own, but instead after another servlet (such as an HTML generator) has produced the actual content.

Example 5-5: Date Filtering Servlet

```

import javax.servlet.*;
import javax.servlet.http.*;
import java.io.*;
import java.util.*;

public class DateFilter extends HttpServlet {

    public void doGet(HttpServletRequest req, HttpServletResponse resp)
        throws ServletException, IOException {

        PrintWriter out = resp.getWriter();

        String contentType = req.getContentType();
        if (contentType == null)
            return; // No incoming data

        // Note that if we were using MIME filtering we would have to set this to
        // something different to avoid an infinite loop
        resp.setContentType(contentType);

        BufferedReader br = new BufferedReader(req.getReader());

        String line = null;
        Date d = new Date();
        while ((line = br.readLine()) != null) {
            int index;
            while ((index=line.indexOf("<DATE>")) >= 0)
                line = line.substring(0, index) + d + line.substring(index + 6);
            out.println(line);
        }

        br.close();
    }
}

```

The `DateFilter` servlet works by reading each line of input, scanning for the text `<DATE>`, and replacing it with the current date. This example introduces the `getReader()` method of `HttpServletRequest`, which returns a `PrintReader` that points to the original request body. When you call `getReader()` in an `HttpServlet`, you can read the original HTTP form variables, if any. When this method is used within a filtering servlet, it provides access to the output of the previous servlet in the chain.

Filters

Version 2.3 of the Servlet API introduced a new method of handling requests, via the `javax.servlet.Filter` class. When filters are used, the servlet container creates a *filter chain*. This consists of zero or more `Filter` objects and a destination resource, which can be either a servlet or another resources available on the web server (such as an HTML or JSP file).

Filters are installed in the server and associated with particular request paths (just like servlets). When a filtered resource is requested, the servlet constructs a filter chain and calls the `doFilter()` method of the first filter in the filter chain, passing a `ServletRequest`, a `ServletResponse`, and the `FilterChain` object. The filter can then perform processing on the request. The processing is sometimes noninterventionary (such as logging characteristics of the request or tracking a clickstream). However, the filter can also wrap the `ServletRequest` and `ServletResponse` classes with its own versions, overriding particular methods. For instance, one of the example filters included with the Tomcat server adds support for returning compressed output to browsers that support it.

After the filter has processed the response, it can call the `doFilter()` method of the `FilterChain` to invoke the next filter in the sequence. If there are no more filters, the request will be passed on to its ultimate destination. After calling `doFilter()`, the filter can perform additional processing on the response received from farther down the chain.

In the event of an error, the filter can stop processing, returning to the client whatever response has already been created, or forwarding the request on to a different resource.

Example 5-6 duplicates the form-based authentication feature that already exists in the servlet API, but could be customized to provide additional functionality not available directly from the server (for instance, authenticating users against systems other than those supported by the servlet container). It works by intercepting each request and checking the `HttpSession` for an attribute called "enterprise.login." If that attribute contains a `Boolean.TRUE`, access is permitted. If not, the filter checks for request parameters named "login_name" and "login_pass," and searches for a match in a hashtable containing valid username/password pairs. If valid login credentials are found, processing the filter chain is allowed to continue. If not, the user is served a login page located at `/login.jsp`, retrieved via a `RequestDispatcher`.*

Astute readers will note that we try to retrieve the users' hashtable from a servlet context attribute. We showed how to set this attribute at server startup in the section "Custom Servlet Initialization." In case you don't have that set up, the Filter's `init()` method will create its own if it can't find one in the context.

Example 5-6: AuthenticationFilter

```
import javax.servlet.*;
import javax.servlet.http.*;
```

* This isn't a highly secure system. Unless the client has connected via SSL, the username/password combination is transmitted unencrypted over the Internet. Also, successful logins leave the `login_name` and `login_pass` parameters in the request when processing it, potentially making them available to a malicious JSP file or servlet. This can be an issue when designing a shared security scheme for dynamic content created by a group of different users (such as at an ISP). One way to get around this is to create a custom `HttpServletRequest` wrapper that filters out the `login_name` and `login_pass` parameters for filters and resources further down the chain.

Example 5-6: AuthenticationFilter (continued)

```
import java.util.Hashtable;

public class AuthenticationFilter implements Filter {

    private Hashtable users = null;

    public void init(FilterConfig config)
        throws javax.servlet.ServletException {

        users = (Hashtable)config.getServletContext().getAttribute(
            "enterprise.users");

        if(users == null) {
            users = new Hashtable(5);
            users.put("test", "test");
        }
    }

    public void doFilter(
        ServletRequest req, ServletResponse res, FilterChain chain)
        throws java.io.IOException, javax.servlet.ServletException {

        HttpServletRequest request = (HttpServletRequest)req;
        HttpSession sess = request.getSession(true);

        if(sess != null) {
            Boolean loggedIn = (Boolean)sess.getAttribute("enterprise.login");
            if (loggedIn != Boolean.TRUE) {
                String login_name = request.getParameter("login_name");
                String login_pass = request.getParameter("login_pass");
                if((login_name != null) && (login_pass != null))
                    if(users.get(login_name).toString().equals(login_pass)) {
                        loggedIn = Boolean.TRUE;
                        sess.setAttribute("enterprise.login", Boolean.TRUE);
                        sess.setAttribute("enterprise.loginname", login_name);
                    }
            }

            if (loggedIn == Boolean.TRUE) {
                chain.doFilter(req, res);
            } else {
                request.setAttribute("originaluri", request.getRequestURI());
                request.getRequestDispatcher("/login.jsp").forward(req, res);
            }
        }
    }

    public void destroy() {
        // Code cleanup would be here
    }
}
```


Here's the JSP page used to display the login form. The important thing to note is that the form submits back to the original URI. The filter uses the `setAttribute()` method of `HttpServletRequest` to specify the URI to post the form back to; the filter is then reapplied, and if the user has provided appropriate credentials access to the resource is granted. For more on JSP, see Chapter 6.

```
<html><body bgcolor="white">

  <% out.print ("<FORM METHOD=POST ACTION=\""+request.
    getAttribute("originaluri").toString()+"\">"); %>
  Login Name: <INPUT TYPE=TEXT NAME="login_name"><br>
  Password: <INPUT TYPE=PASSWORD NAME="login_pass">
  <INPUT TYPE=SUBMIT VALUE="Log In">
</FORM>

</body></html>
```

When configuring the filter, map it to the paths you wish to protect. Mapping it to `/*` will not work, as that would also protect the `/login.jsp` file (which will be run through its own filter chain by the `RequestDispatcher` object). If you did want to protect your whole application, you could build the login form internally to the filter; but this is generally considered bad practice.

Thread Safety

In a typical scenario, only one copy of any particular servlet or filter is loaded at any given time. Each servlet might, however, be called upon to deal with multiple requests at the same time. This means that a servlet needs to be thread-safe. If a servlet doesn't use any class variables (that is, any variables with a scope broader than the service method itself), it is generally already thread-safe. If you are using any third-party libraries or extensions, make sure that those components are also thread-safe. However, a servlet that maintains persistent resources needs to make sure that nothing untoward happens to those resources. Imagine, for example, a servlet that maintains a bank balance using an `int` in memory.* If two servlets try to access the balance at the same time, you might get this sequence of events:

1. User 1 connects to the servlet to make a \$100 withdrawal.
2. The servlet checks the balance for User 1, finding \$120.
3. User 2 connects to the servlet to make a \$50 withdrawal.
4. The servlet checks the balance for User 2, finding \$120.
5. The servlet debits \$100 for User 1, leaving \$20.
6. The servlet debits \$50 for User 2, leaving -\$30.
7. The programmer is fired.

Obviously, this is incorrect behavior, particularly that last bit. We want the servlet to perform the necessary action for User 1, and then deal with User 2 (in this case, by giving him an insufficient funds message). We can do this by surrounding

* Hey, bear with us on this one. This is an example.

sections of code with synchronized blocks. While a particular synchronized block is executing, no other sections of code that are synchronized on the same object (usually the servlet or the resource being protected) can execute. For more information on thread safety and synchronization, see *Java Threads* by Scott Oaks and Henry Wong (O'Reilly).

Example 5-7 implements the ATM display for the First Bank of Java. The `doGet()` method displays the current account balance and provides a small ATM control panel for making deposits and withdrawals, as shown in Figure 5-3.*

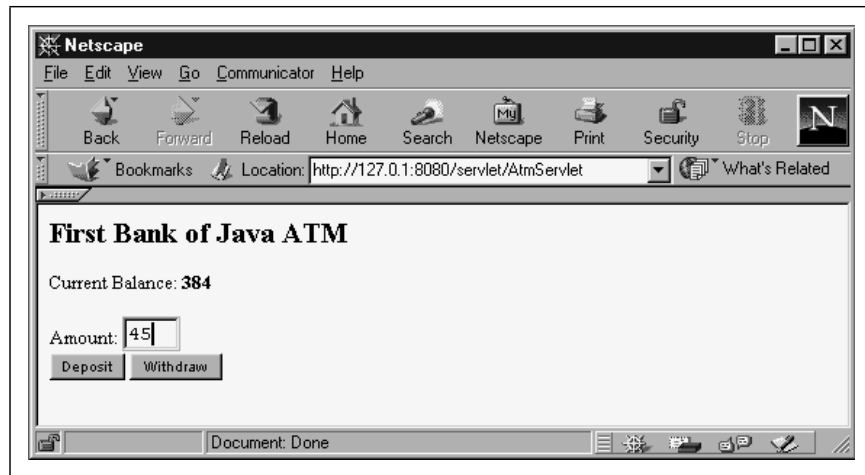


Figure 5-3: The First Bank of Java ATM display

The control panel uses a POST request to send the transaction back to the servlet, which performs the appropriate action and calls `doGet()` to redisplay the ATM screen with the updated balance.

Example 5-7: An ATM Servlet

```
import javax.servlet.*;
import javax.servlet.http.*;
import java.util.*;
import java.io.*;

public class AtmServlet extends HttpServlet {

    Account act;

    public void init(ServletConfig conf) throws ServletException {
        super.init(conf);
        act = new Account();
        act.balance = 0;
    }
}
```

* Despite the fact that Java is a very large island, there's still only one account.

Example 5-7: An ATM Servlet (continued)

```

public void doGet(HttpServletRequest req, HttpServletResponse resp)
    throws ServletException, IOException {

    resp.setContentType("text/html");
    PrintWriter out = resp.getWriter();

    out.println("<HTML><<BODY>");
    out.println("<H2>First Bank of Java ATM</H2>");
    out.println("Current Balance: <B>" + act.balance + "</B><<BR>");
    out.println("<FORM METHOD=POST ACTION=/servlet/AtmServlet>");
    out.println("Amount: <INPUT TYPE=TEXT NAME=AMOUNT SIZE=3><<BR>");
    out.println("<INPUT TYPE=SUBMIT NAME=DEPOSIT VALUE=\"Deposit\">");
    out.println("<INPUT TYPE=SUBMIT NAME=WITHDRAW VALUE=\"Withdraw\">");
    out.println("</FORM>");
    out.println("</BODY><</HTML>");
}

public void doPost(HttpServletRequest req, HttpServletResponse resp)
    throws ServletException, IOException {

    int amt=0;

    try {
        amt = Integer.parseInt(req.getParameter("AMOUNT"));
    }
    catch (NullPointerException e) {
        // No Amount Parameter passed
    }
    catch (NumberFormatException e) {
        // Amount Parameter was not a number
    }

    synchronized(act) {
        if(req.getParameter("WITHDRAW") != null) && (amt < act.balance)
            act.balance = act.balance - amt;
        if(req.getParameter("DEPOSIT") != null) && (amt > 0)
            act.balance = act.balance + amt;
    } // end synchronized block

    doGet(req, resp);           // Show ATM screen
}

public void destroy() {
    // This is where we would save the balance to a file
}

class Account {
    public int balance;
}
}

```

The `doPost()` method alters the account balance contained within an `Account` object `act` (since `Account` is so simple, we've defined it as an inner class). In order to prevent multiple requests from accessing the same account at once, any code that alters `act` is synchronized on `act`. This ensures that no other code can alter `act` while a synchronized section is running.

The `destroy()` method is defined in the `AtmServlet`, but it contains no actual code. A real banking servlet would obviously want to write the account balance to disk before being unloaded. And if the servlet were using JDBC to store the balance in a database, it would also want to destroy all its database-related objects.

A more complex servlet than `AtmServlet` might need to synchronize its entire service method, limiting the servlet to one request at a time. In these situations, it sometimes makes sense to modify the standard servlet lifecycle a little bit. We can do this by implementing the `SingleThreadModel` interface. This is a tag interface that has no methods; it simply tells the server to create a pool of servlet instances, instead of a single instance of the servlet. To handle an incoming request, the server uses a servlet from the pool and only allows each copy of the servlet to serve one request at a time. Implementing this interface effectively makes a servlet thread-safe, while allowing the server to deal with more than one connection at a time. Of course, using `SingleThreadModel` does increase resource requirements and make it difficult to share data objects within a servlet.

Another use for `SingleThreadModel` is to implement simple database connection sharing. Having multiple database connections can improve performance and avoid connection overloading. Of course, for more advanced or high-traffic applications, you generally want to manage connection pooling explicitly, rather than trusting the web server to do it for you.



Cookies

Cookies spent a year or two as a little-known feature of Netscape Navigator before becoming the focus of a raging debate on electronic privacy. Ethical and moral considerations aside, cookies allow a web server to store small amounts of data on client systems. Cookies are generally used to store basic user identification or configuration information. Because a cookie's value can uniquely identify a client, cookies are often used for session tracking (although, as we'll see shortly, the Servlet API provides higher-level support for this).

To create a cookie, the server (or, more precisely, a web application running on the server) includes a `Cookie` header with a specific value in an HTTP response. The browser then transmits a similar header with that value back to the server with subsequent requests, which are subject to certain rules. The web application can use the cookie value to keep track of a particular user, handle session tracking, etc. Because cookies use a single `Cookie` header, the syntax for a cookie allows for multiple name/value pairs in the overall cookie value.

More information about the cookies is available from the original Netscape specification document at http://home.netscape.com/newsref/std/cookie_spec.html. The Internet Engineering Task Force is currently working on a standard cookie specification, defined in RFC-2109, available at <http://www.internic.net/rfc/rfc2109.txt>.

The Servlet API includes a class, `javax.servlet.http.Cookie`, that abstracts cookie syntax and makes it easy to work with cookies. In addition, `HttpServletResponse` provides an `addCookie()` method and `HttpServletRequest` provides a `getCookies()` method to aid in writing cookies to and reading cookies from the HTTP headers, respectively. To find a particular cookie, a servlet needs to read the entire collection of values and look through it:

```
Cookie[] cookies;
cookies = req.getCookies();
String userid = null;

for (int i = 0; i < cookies.length; i++)
    if (cookies[i].getName().equals("userid"))
        userid = cookies[i].getValue();
```

A cookie can be read at any time, but can be created only before any content is sent to the client. This is because cookies are sent using HTTP headers. These headers can be sent to the client before the regular content. Once any data has been written to the client, the server can flush the output and send the headers at any time, so you can't create any new cookies safely. You must create new cookies before sending any output. Here's an example of creating a cookie:

```
String userid = createUserID(); // Create a unique ID
Cookie c = new Cookie("userid", userid);
resp.addCookie(c); // Add the cookie to the HTTP headers
```

Note that a web browser is only required to accept 20 cookies per site and 300 total per user, and the browser can limit each cookie's size to 4096 bytes.

Cookies can be customized to return information only in specific circumstances. In particular, a cookie can specify a particular domain, a particular path, an age after which the cookie should be destroyed, and whether the cookie requires a secure (HTTPS) connection. A cookie is normally returned only to the host that specified it. For example, if a cookie is set by `server1.company.com`, it isn't returned to `server2.company.com`. You can get around this limitation by setting the domain to `.company.com` with the `setDomain()` method of `Cookie`. By the same token, a cookie is generally returned for pages only in the same directory as the servlet that created the cookie, or it's returned under that directory. We can get around this limitation using `setPath()`. Here's a cookie that is returned to all pages on all top-level servers at `company.com`:

```
String userid = createUserID(); // Create a unique ID
Cookie c = new Cookie("userid", userid);
c.setDomain(".company.com"); // *.company.com, but not *.web.company.com
c.setPath("/"); // All pages
resp.addCookie(c); // Add the cookie to the HTTP headers
```

Session Tracking

Very few web applications are confined to a single page, so having a mechanism for tracking users through a site can often simplify application development. The Web, however, is an inherently stateless environment. A client makes a request, the server fulfills it, and both promptly forget about each other. In the past, applications that needed to deal with a user through multiple pages (for instance, a

shopping cart) had to resort to complicated dodges to hold onto state information, such as hidden fields in forms, setting and reading cookies, or rewriting URLs to contain state information.

The Servlet API provides classes and methods specifically designed to handle session tracking. A servlet can use the session-tracking API to delegate most of the user-tracking functions to the server. The first time a user connects to a session-enabled servlet, the servlet simply creates a `javax.servlet.http.HttpSession` object. The servlet can then bind data to this object, so subsequent requests can read the data. After a certain amount of inactive time, the session object is destroyed.

A servlet uses the `getSession()` method of `HttpServletRequest` to retrieve the current session object. This method takes a single `boolean` argument. If you pass `true`, and there is no current session object, the method creates and returns a new `HttpSession` object. If you pass `false`, the method returns `null` if there is no current session object. For example:

```
HttpSession thisUser = req.getSession(true);
```

When a new `HttpSession` is created, the server assigns a unique session ID that must somehow be associated with the client. Since clients differ in what they support, the server has a few options that vary slightly depending on the server implementation. In general, the server's first choice is to try to set a cookie on the client (which means that `getSession()` must be called before you write any other data back to the client). If cookie support is lacking, the API allows servlets to rewrite internal links to include the session ID, using the `encodeURL()` method of `HttpServletResponse`. This is optional, but recommended, particularly if your servlets share a system with other, unknown servlets that may rely on uninterrupted session tracking. However, this on-the-fly URL encoding can become a performance bottleneck because the server needs to perform additional parsing on each incoming request to determine the correct session key from the URL. (The performance hit is so significant that the Java Web Server disables URL encoding by default.)

To use URL encoding run all your internal links through `encodeURL()`. If you have a line of code like this:

```
out.println("<A HREF=\"/servlet/CheckoutServlet\">Check Out</A>");
```

you should replace it with:

```
out.print("<A HREF=\"");  
out.print(resp.encodeURL("/servlet/CheckoutServlet");  
out.println("\">Check Out</A>");
```

JWS, in this case, adds an identifier beginning with `$` to the end of the URL. Other servers have their own methods. Thus, with JWS, the final output looks like this:

```
<A HREF="/servlet/CheckoutServlet$FASED4W23798ASD978">Check Out</A>
```

In addition to encoding your internal links, you need to use `encodeRedirectURL()` to handle redirects properly. This method works in the same manner as `encodeURL()`.*

You can access the unique session ID via the `getSessionID()` method of `HttpSession`. This is enough for most applications, since a servlet can use some other storage mechanism (i.e., a flat file, memory, or a database) to store the unique information (e.g., hit count or shopping cart contents) associated with each session. However, the API makes it even easier to hold onto session-specific information by allowing servlets to bind objects to a session using the `putValue()` method of `HttpSession`. Once an object is bound to a session, you can use the `getValue()` method.†

Objects bound using `putValue()` are available to all servlets running on the server. The system works by assigning a user-defined name to each object (the `String` argument); this name is used to identify objects at retrieval time. In order to avoid conflicts, the general practice is to name bound objects with names of the form *applicationname.objectname*. For example:

```
session.putValue("myservlet.hitcount", new Integer(34));
```

Now that object can be retrieved with:

```
Integer hits = (Integer)session.getValue("myservlet.hitcount")
```

Example 5-8 demonstrates a basic session-tracking application that keeps track of the number of visits to the site by a particular user. It works by storing a counter value in an `HttpSession` object and incrementing it as necessary. When a new session is created (as indicated by `isNew()`, which returns `true` if the session ID has not yet passed through the client and back to the server), or the counter object is not found, a new counter object is created.

Example 5-8: Counting Visits with Sessions

```
import javax.servlet.*;
import javax.servlet.http.*;
import java.io.*;

public class VisitCounterServlet extends HttpServlet {

    public void doGet(HttpServletRequest req, HttpServletResponse resp)
        throws ServletException, IOException {

        PrintWriter out = resp.getWriter();
        resp.setContentType("text/html");
```

* These methods were introduced in Version 2.1 of the Servlet API, replacing two earlier methods named `encodeUrl()` and `encodeRedirectUrl()`. This was done to bring the capitalization scheme in line with other Java APIs.

† The `putValue()` and `getValue()` methods should be used only with Servlet 2.0 containers because they have been deprecated in favor of the work-alive `setAttribute()` and `getAttribute()` methods. The naming change was done to create consistency across various attribute-capable elements of the Servlet API.

Example 5-8: Counting Visits with Sessions (continued)

```

HttpSession thisUser = req.getSession(true);
Integer visits;

if(!thisUser.isNew()) {           //Don't check newly created sessions
    visits = (Integer)thisUser.getValue("visitcounter.visits");
    if(visits == null)
        visits = new Integer(1);
    else
        visits = new Integer(visits.intValue() + 1);
}
else
    visits = new Integer(1);

// Put the new count in the session
thisUser.putValue("visitcounter.visits", visits);

// Finally, display the results and give them the session ID too
out.println("<HTML><HEAD><TITLE>Visit Counter</TITLE></HEAD>");
out.println("<BODY>You have visited this page " + visits + " time[s]");
out.println("since your last session expired.");
out.println("Your Session ID is " + thisUser.getId());
out.println("</BODY></HTML>");
}
}

```

HttpSessionBindingListener

Sometimes it is useful to know when an object is getting bound or unbound from a session object. For instance, in an application that binds a JDBC `java.sql.Connection` object to a session (something that, by the way, is ill-advised in all but very low traffic sites), it is important that the `Connection` be explicitly closed when the session is destroyed.

The `javax.servlet.http.HttpSessionBindingListener` interface handles this task. It includes two methods, `valueBound()` and `valueUnbound()`, that are called whenever the object that implements the interface is bound or unbound from a session, respectively. Each of these methods receives an `HttpSessionBindingEvent` object that provides the name of the object being bound or unbound and the session involved in the action. Here is an object that implements the `HttpSessionBindingListener` interface in order to make sure that a database connection is closed properly:

```

class ConnectionHolder implements HttpSessionBindingListener {

    java.sql.Connection dbCon;

    public ConnectionHolder(java.sql.Connection con) {
        dbCon = con;
    }
}

```



```
public void valueBound(HttpSessionBindingEvent event) {
    // Do nothing
}

public void valueUnbound(HttpSessionBindingEvent event) {
    dbCon.close();
}
}
```

Session Contexts

Version 2.0 of the Servlet API included the `getContext()` method of `HttpSession`, coupled with an interface named `HttpSessionContext`. Together, these allowed servlets to access other sessions running in the same context. Unfortunately, this functionality also allowed a servlet to accidentally expose all the session IDs in use on the server, meaning that an outsider with knowledge could spoof a session. To eliminate this minor security risk, the session-context functionality was deprecated in Version 2.1 of the Servlet API. Instead, web applications can use the `getAttribute()` and `setAttribute()` methods of `ServletContext` to share information across sessions.

Databases and Non-HTML Content

Most web applications need to communicate with a database, either to generate dynamic content or collect and store data from users, or both. With servlets, this communication is easily handled using the JDBC API described in Chapter 2. Thanks to JDBC and the generally sensible design of the servlet lifecycle, servlets are an excellent intermediary between a database and web clients.

Most of the general JDBC principles discussed in Chapter 2 apply to servlets. However, servlet developers should keep a few things in mind for optimal performance. First, JDBC Connection objects can be created in the servlet's `init()` method. This allows the servlet to avoid reconnecting to the database (a la CGI) with each request, saving up to a second or more on every single page request. If you anticipate high volume, you may want to create several connections and rotate between them. An excellent freeware connection-pooling system is available at <http://www.javaexchange.com>. Or, if you're using JDBC 2.0, the `javax.sql` package provides a connection-pooling mechanism. Finally, if you plan on using JDBC's transaction support, you need to create individual connections for each request or obtain exclusive use of a pooled connection.

So far, all our servlets have produced standard HTML content. Of course, this is all most servlets ever do, but it's not all that they can do. Say, for instance, that your company stores a large database of PDF documents within an Oracle database, where they can be easily accessed. Now say you want to distribute these documents on the Web. Luckily, servlets can dish out any form of content that can be defined with a MIME header. All you have to do is set the appropriate content type and use a `ServletOutputStream` if you need to transmit binary data. Example 5-9 shows how to pull an Adobe Acrobat document from an Oracle database.

Example 5-9: A Servlet That Serves PDF Files from a Database

```
import java.io.*;
import java.sql.*;
import javax.servlet.*;
import javax.servlet.http.*;

public class DBPDFReader extends HttpServlet {

    Connection con;


    public void init(ServletConfig config) throws ServletException {
        super.init(config);
        try {
            Class.forName("oracle.jdbc.driver.OracleDriver");
            con = DriverManager.getConnection("jdbc:oracle:oci8:@DBHOST",
                "user", "passwd");
        }
        catch (ClassNotFoundException e) {
            throw new UnavailableException(this, "Couldn't load OracleDriver");
        }
        catch (SQLException e) {
            throw new UnavailableException(this, "Couldn't get db connection");
        }
    }

    public void doGet(HttpServletRequest req, HttpServletResponse res)
        throws ServletException, IOException {

        try {
            res.setContentType("application/pdf");
            ServletOutputStream out = res.getOutputStream();

            Statement stmt = con.createStatement();
            ResultSet rs = stmt.executeQuery(
                "SELECT PDF FROM PDF WHERE PDFID = " + req.getParameter("PDFID"));

            if (rs.next()) {
                BufferedInputStream pdfData =
                    new BufferedInputStream(rs.getBinaryStream("PDF"));
                byte[] buf = new byte[4 * 1024]; // 4K buffer
                int len;
                while ((len = pdfData.read(buf, 0, buf.length)) != -1) {
                    out.write(buf, 0, len);
                }
            }
            else {
                res.sendError(res.SC.NOT_FOUND);
            }
        }
    }
}
```



Example 5-9: A Servlet That Serves PDF Files from a Database (continued)

```
        rs.close();
        stmt.close ();
    }
    catch(SQLException e) {
        // Report it
    }
}
```

