



BEA

WebLogic Server

Programming

WebLogic HTTP Servlets

BEA WebLogic Server 6.0
Document Date: March 3, 2001

Copyright

Copyright © 2001 BEA Systems, Inc. All Rights Reserved.

Restricted Rights Legend

This software and documentation is subject to and made available only pursuant to the terms of the BEA Systems License Agreement and may be used or copied only in accordance with the terms of that agreement. It is against the law to copy the software except as specifically allowed in the agreement. This document may not, in whole or in part, be copied photocopied, reproduced, translated, or reduced to any electronic medium or machine readable form without prior consent, in writing, from BEA Systems, Inc.

Use, duplication or disclosure by the U.S. Government is subject to restrictions set forth in the BEA Systems License Agreement and in subparagraph (c)(1) of the Commercial Computer Software-Restricted Rights Clause at FAR 52.227-19; subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013, subparagraph (d) of the Commercial Computer Software--Licensing clause at NASA FAR supplement 16-52.227-86; or their equivalent.

Information in this document is subject to change without notice and does not represent a commitment on the part of BEA Systems. THE SOFTWARE AND DOCUMENTATION ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. FURTHER, BEA Systems DOES NOT WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE, OR THE RESULTS OF THE USE, OF THE SOFTWARE OR WRITTEN MATERIAL IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE.

Trademarks or Service Marks

BEA, WebLogic, Tuxedo, and Jolt are registered trademarks of BEA Systems, Inc. How Business Becomes E-Business, BEA WebLogic E-Business Platform, BEA Builder, BEA Manager, BEA eLink, BEA WebLogic Commerce Server, BEA WebLogic Personalization Server, BEA WebLogic Process Integrator, BEA WebLogic Collaborate, BEA WebLogic Enterprise, and BEA WebLogic Server are trademarks of BEA Systems, Inc.

All other product names may be trademarks of the respective companies with which they are associated.

Programming WebLogic HTTP Servlets

Part Number	Document Date	Software Version
	March 6, 2001	BEA WebLogic Server Version 6.0

Contents

About This Document

Audience.....	v
e-docs Web Site.....	vi
How to Print the Document.....	vi
Related Information.....	vi
Contact Us!.....	vii
Documentation Conventions.....	viii

1. Overview

What Is a Servlet?.....	1-1
What You Can Do with Servlets.....	1-2
Overview of Servlet Development.....	1-3
Servlets and J2EE.....	1-3
HTTP Servlet API Reference.....	1-4

2. Introduction to Programming

Writing a Basic HTTP Servlet.....	2-1
Advanced Features.....	2-4
Complete HelloWorldServlet Example.....	2-5

3. Programming Tasks

Initializing a Servlet.....	3-1
Initializing a Servlet when WebLogic Server Starts.....	3-2
Overriding the init() Method.....	3-3
Providing an HTTP Response.....	3-4
Retrieving Client Input.....	3-6
Other Methods for Using the HTTP Request.....	3-7

Example: Retrieving Input by Using Query Parameters	3-8
Using Session Tracking from a Servlet	3-10
A History of Session Tracking	3-10
Tracking a Session with an HttpSession Object.....	3-11
Lifetime of a Session	3-12
How Session Tracking Works	3-12
Detecting the Start of a Session.....	3-13
Setting and Getting Session Name/Value Attributes	3-13
Logging Out and Ending a Session	3-14
Configuring Session Tracking	3-15
Using URL Rewriting.....	3-15
URL Rewriting and Wireless Access Protocol (WAP).....	3-16
Making Sessions Persistent	3-16
Scenarios to Avoid When Using Sessions	3-17
Use Serializable Attribute Values	3-18
Configuring Session Persistence	3-18
Using Cookies in a Servlet	3-18
Setting Cookies in an HTTP servlet.....	3-18
Retrieving Cookies in an HTTP Servlet.....	3-19
Using Cookies that Are Transmitted by Both HTTP and HTTPS	3-20
Application Security and Cookies	3-21
Using WebLogic Services from an HTTP Servlet	3-21
Accessing Databases.....	3-22
Connecting to a Database Using a JDBC Connection Pool	3-22
Using a Connection Pool in a Servlet.....	3-23
Connecting to a Database Using a DataSource Object	3-24
Using a DataSource in a Servlet.....	3-24
Connecting Directly to a Database Using a JDBC Driver	3-25
Threading Issues in HTTP Servlets	3-25
SingleThreadModel	3-25
Shared Resources.....	3-26
Dispatching Requests to Another Resource	3-26
Forwarding a Request.....	3-28
Including a Request	3-28

4. Administration and Configuration

Overview	4-1
Using Deployment Descriptors to Configure and Deploy Servlets	4-2
web.xml (Web Application Deployment Descriptor)	4-2
weblogic.xml (Weblogic-Specific Deployment Descriptor)	4-3
WebLogic Server Administration Console	4-3
Directory Structure	4-5
URL	4-5
Security	4-6
Authentication	4-6
Authorization (Security Constraints)	4-7
Servlet Development Tips	4-7
Clustering Servlets.....	4-8

Index



About This Document

This document provides information on programming and deploying WebLogic HTTP Servlets.

The document is organized as follows:

- Chapter 1, “Overview,” provides an overview of Hypertext transfer protocol (HTTP) servlet programming and explains how to use HTTP servlets with WebLogic Server.
- Chapter 2, “Introduction to Programming,” introduces basic HTTP servlet programming.
- Chapter 3, “Programming Tasks,” provides information about writing HTTP servlets in a WebLogic Server environment.
- Chapter 4, “Administration and Configuration,” provides information about writing HTTP servlets in a WebLogic Server environment.

Audience

This document is written for application developers who want to build e-commerce applications using HTTP servlets and the Java 2 Platform, Enterprise Edition (J2EE) from Sun Microsystems. It is assumed that readers know Web technologies, object-oriented programming techniques, and the Java programming language.

e-docs Web Site

BEA product documentation is available on the BEA corporate Web site. From the BEA Home page, click on Product Documentation or go directly to the WebLogic Server Product Documentation page at <http://e-docs.bea.com/wls/docs60>.

How to Print the Document

You can print a copy of this document from a Web browser, one main topic at a time, by using the File→Print option on your Web browser.

A PDF version of this document is available on the WebLogic Server documentation Home page on the e-docs Web site (and also on the documentation CD). You can open the PDF in Adobe Acrobat Reader and print the entire document (or a portion of it) in book format. To access the PDFs, open the WebLogic Server documentation Home page, click Download Documentation, and select the document you want to print.

Adobe Acrobat Reader is available at no charge from the Adobe Web site at <http://www.adobe.com>.

Related Information

- [Package javax.servlet](http://java.sun.com/products/servlet/2.2/javadoc/javax/servlet/package-summary.html)
(<http://java.sun.com/products/servlet/2.2/javadoc/javax/servlet/package-summary.html>)
- [Package javax.servlet.http](http://java.sun.com/products/servlet/2.2/javadoc/javax/servlet/http/package-summary.html)
(<http://java.sun.com/products/servlet/2.2/javadoc/javax/servlet/http/package-summary.html>)
- [Servlet 2.2 specification](http://java.sun.com/products/servlet/download.html#specs)
(<http://java.sun.com/products/servlet/download.html#specs>)

-
- [Deploying and Configuring Applications](#) at
http://e-docs.bea.com/wls/docs60/adminguide/config_web_app.html
 - [Writing Web Application Deployment Descriptors](#) at
<http://e-docs.bea.com/wls/docs60/programming/webappdeployment.html>

Contact Us!

Your feedback on BEA documentation is important to us. Send us e-mail at docsupport@bea.com if you have questions or comments. Your comments will be reviewed directly by the BEA professionals who create and update the documentation.

In your e-mail message, please indicate the software name and version your are using, as well as the title and document date of your documentation. If you have any questions about this version of BEA WebLogic Server, or if you have problems installing and running BEA WebLogic Server, contact BEA Customer Support through BEA WebSupport at <http://www.bea.com>. You can also contact Customer Support by using the contact information provided on the Customer Support Card, which is included in the product package.

When contacting Customer Support, be prepared to provide the following information:

- Your name, e-mail address, phone number, and fax number
- Your company name and company address
- Your machine type and authorization codes
- The name and version of the product you are using
- A description of the problem and the content of pertinent error messages

Documentation Conventions

The following documentation conventions are used throughout this document.

Convention	Usage
Ctrl+Tab	Keys you press simultaneously.
<i>italics</i>	Emphasis and book titles.
monospace text	Code samples, commands and their options, Java classes, data types, directories, and file names and their extensions. Monospace text also indicates text that you enter from the keyboard. <i>Examples:</i> <pre>import java.util.Enumeration; chmod u+w * config/examples/applications .java config.xml float</pre>
<i>monospace italic text</i>	Variables in code. <i>Example:</i> <pre>String CustomerName;</pre>
UPPERCASE TEXT	Device names, environment variables, and logical operators. <i>Examples:</i> <pre>LPT1 BEA_HOME OR</pre>
{ }	A set of choices in a syntax line.
[]	Optional items in a syntax line. <i>Example:</i> <pre>java utils.MulticastTest -n name -a address [-p portnumber] [-t timeout] [-s send]</pre>

Convention	Usage
	<p>Separates mutually exclusive choices in a syntax line. <i>Example:</i></p> <pre>java weblogic.deploy [list deploy undeploy update] password {application} {source}</pre>
...	<p>Indicates one of the following in a command line:</p> <ul style="list-style-type: none"> ■ An argument can be repeated several times in the command line. ■ The statement omits additional optional arguments. ■ You can enter additional parameters, values, or other information
.	<p>Indicates the omission of items from a code example or from a syntax line.</p> <p>.</p> <p>.</p>



1 Overview

This section provides an overview of Hypertext transfer protocol (HTTP) servlet programming and explains how to use HTTP servlets with WebLogic Server. The following topics are discussed:

- [What Is a Servlet?](#)
- [What You Can Do with Servlets](#)
- [Overview of Servlet Development](#)
- [Servlets and J2EE](#)
- [HTTP Servlet API Reference](#)

What Is a Servlet?

A servlet is a Java class that runs in a Java-enabled server. An *HTTP* servlet is a special type of servlet that handles an HTTP request and provides an HTTP response, usually in the form of an HTML page. The most common use of WebLogic HTTP Servlets is to create interactive applications using standard Web browsers for the client-side presentation while WebLogic Server handles the business logic as a server-side process. WebLogic HTTP servlets can access databases, Enterprise JavaBeans, messaging APIs, HTTP sessions, and other facilities of WebLogic Server.

WebLogic Server fully supports HTTP servlets as defined in the Servlet 2.2 specification from Sun Microsystems. HTTP servlets form an integral part of the Java 2 Enterprise Edition (J2EE) standard.

What You Can Do with Servlets

- Create dynamic Web pages that utilize HTML forms to get end-user input and provide HTML pages that respond to that input. Examples of this include online shopping carts, financial services, and personalized content.
- Create collaborative systems such as online conferencing.
- Servlets running in WebLogic Server have access to a variety of APIs and services. For example:
 - Session tracking—allows a Web site to track a user’s progress across multiple Web pages. This functionality supports Web sites such as e-commerce sites that use shopping carts. WebLogic Server supports session persistence to a database, providing fail-over between server down time and session sharing between clustered servers. For more information see [Using Session Tracking from a Servlet](#) on page 3-10.
 - JDBC drivers (including BEA’s)—provide basic database access. With Weblogic Server’s multitier JDBC implementations, you can take advantage of connection pools, server-side data caching, and transactions. For more information see [Accessing Databases](#) on page 3-22.
 - Security—you can apply various types of security to servlets, including using ACLs for authentication and Secure Sockets Layer (SSL) to provide secure communications.
 - Enterprise JavaBeans—servlets can use Enterprise JavaBeans (EJB) to encapsulate sessions, data from databases, and other functionality.
 - Java Messaging Service (JMS)—JMS allows your servlets to exchange messages with other servlets and Java programs.
 - Java JDK APIs—servlets can use the standard Java JDK APIs.
 - Forwarding requests—servlets can forward a request to another servlet or other resource.
- Servlets written for any J2EE-compliant Servlet engine can be easily deployed on WebLogic Server.
- Servlets and Java Server Pages (JSP) can work together to create an application.

Overview of Servlet Development

- Programmers of HTTP servlets utilize a standard API from JavaSoft, `javax.servlet.http` to create interactive applications.
- HTTP servlets can read HTTP headers and write HTML coding to deliver a response to a browser client.
- Servlets are deployed on WebLogic Server as part of a Web Application. A Web Application is a grouping of application components such as servlet classes, JavaServer Pages (JSP), static HTML pages, images, and security. For more information see “[Administration and Configuration](#)” on page 4-1.

Servlets and J2EE

The [Servlet 2.2 specification](#) (available at <http://java.sun.com/products/servlet/download.html#specs>), part of the Java 2 Platform, Enterprise Edition, defines the implementation of the servlet API and the method by which servlets are deployed in enterprise applications. Deploying servlets on a J2EE-compliant server, such as WebLogic Server, is accomplished by packaging the servlets and other resources that make up an enterprise application into a single unit called a *Web Application*. A Web Application utilizes a specific directory structure to contain its resources and a deployment descriptor that defines how these resources interact and how the application is accessed by a client. A Web Application may also be deployed as an archive file called a `.war` file.

For more information on creating Web Applications, see “[Deploying and Configuring Web Applications](#)” section of the *WebLogic Server Administration Guide* at http://e-docs.bea.com/wls/docs60/adminguide/config_web_app.html. For an overview of servlet administration and deployment issues, see “[Administration and Configuration](#)” on page 4-1.

HTTP Servlet API Reference

WebLogic Server supports the `javax.servlet.http` package in the Java Servlet 2.2 API. You can find additional documentation for the package from Sun Microsystems:

- API documentation
 - [Package javax.servlet](http://java.sun.com/products/servlet/2.2/javadoc/javax/servlet/package-summary.html)
(<http://java.sun.com/products/servlet/2.2/javadoc/javax/servlet/package-summary.html>)
 - [Package javax.servlet.http](http://java.sun.com/products/servlet/2.2/javadoc/javax/servlet/http/package-summary.html)
(<http://java.sun.com/products/servlet/2.2/javadoc/javax/servlet/http/package-summary.html>)
- [Servlet 2.2 specification](http://java.sun.com/products/servlet/download.html#specs)
(<http://java.sun.com/products/servlet/download.html#specs>)

2 Introduction to Programming

This section introduces basic HTTP servlet programming. The following topics are discussed:

- [Writing a Basic HTTP Servlet](#)
- [Advanced Features](#)
- [Complete HelloWorldServlet Example](#)

Writing a Basic HTTP Servlet

The section provides a procedure for writing a basic HTTP servlet. A complete code example (the `HelloWorldServlet`) illustrating these steps is included at the end of this section. Additional information about using various J2EE and WebLogic Server services, such as JDBC, RMI, and JMS, in your servlet are discussed later in this document.

1. Import the appropriate packages and classes, including the following:

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

2. Extend `javax.servlet.http.HttpServlet`. For example:

```
public class HelloWorldServlet extends HttpServlet{
```

3. Implement a `service()` method. The main function of a servlet is to accept an HTTP request from a Web Browser, and return an HTTP response. This work is done by the `service()` method of your servlet. Service methods include *response* objects used to create output and *request* objects used to receive data from the client.

You may have seen other servlet examples implement the `doPost()` and/or `doGet()` methods. These methods reply only to POST or GET requests; if you want to handle all request types from a single method, your servlet can simply implement the `service()` method. (However, if you choose to implement the `service()` method, you will not be able to implement the `doPost()` or `doGet()` methods, unless you call `super.service()` at the beginning of the `service()` method.) The HTTP servlet specification describes other methods used to handle other request types, but all of these methods are collectively referred to as *service* methods.

All the service methods take the same parameter arguments. An `HttpServletRequest` provides information about the request, and an `HttpServletResponse`, is used by your servlet to reply to the HTTP client.

```
public void service(HttpServletRequest req,
                    HttpServletResponse res) throws IOException
{
```

4. Set the content type, as follows:

```
res.setContentType("text/html");
```
5. Get a reference to a `java.io.PrintWriter` object to use for output, as follows:

```
PrintWriter out = res.getWriter();
```

6. Create some HTML using the `println()` method on the `PrintWriter` object, as shown in the following example:

```
out.println("<html><head><title>Hello World!</title></head>");
out.println("<body><h1>Hello World!</h1></body></html>");
}
}
```

7. Compile the servlet, as follows
 - a. Set up a [development environment shell](http://e-docs.bea.com/wls/docs60/programming/environment.html) (see <http://e-docs.bea.com/wls/docs60/programming/environment.html>) with the correct classpath and path settings.

- b. From the directory containing the Java source code for your servlet, compile your servlet into the `WEB-INF/classes` directory of the Web Application that contains your servlet. For example:

```
javac -d /myWebApplication/WEB-INF/classes myServlet.java
```

8. Deploy the servlet as part of a Web Application hosted on WebLogic Server. For an overview of servlet deployment, see [“Administration and Configuration”](#) on page 4-1.
9. Call the servlet from a browser.

The URL you use to call a servlet is determined by: (a) the name of the Web Application containing the servlet and (b) the name of the servlet as mapped in the deployment descriptor of the Web Application. Request parameters can also be included in the URL used to call a servlet.

Generally the URL for a servlet conforms to the following:

```
http://host:port/webApplicationName/mappedServletName?parameter
```

The components of the URL are defined as follows:

- *host* is the name of the machine running WebLogic Server
- *port* is the port at which the above machine is listening for HTTP requests
- *webApplicationName* is the name of the Web Application containing the servlet
- *parameters* are one or more name-value pairs containing information sent from the browser that can be used in your servlet

For example, to use a Web browser to call the `HelloWorldServlet` (the example featured in this document), which is deployed in the `examplesWebApp` and served from a WebLogic Server running on your machine, enter the following URL:

```
http://localhost:7001/examplesWebApp/HelloWorldServlet
```

The `host:port` portion of the URL can be replaced by a DNS name that is mapped to a Web Server running in WebLogic Server.

Advanced Features

The preceding steps create a basic servlet. You will probably also want to use some of the more advanced features of servlets, which are described briefly below:

- Handling HTML form data—HTTP servlets can receive and process data received from a browser client in HTML forms.
 - [“Retrieving Client Input”](#) on page 3-6.
- Application design—HTTP servlets offer many ways to design your application. The following sections provide detailed information about writing servlets:
 - [“Providing an HTTP Response”](#) on page 3-4.
 - [“Threading Issues in HTTP Servlets”](#) on page 3-25.
 - [“Dispatching Requests to Another Resource”](#) on page 3-26.
- Initializing a servlet—if your servlet needs to initialize data, accept initialization arguments, or perform other actions when the servlet is initialized, you can override the `init()` method.
 - [“Initializing a Servlet”](#) on page 3-1.
- Use of *sessions* and *persistence* in your servlet—sessions and persistence allow you to track your users within and between HTTP sessions. Session management includes the use of *cookies*. For more information, see the following sections:
 - [“Using Session Tracking from a Servlet”](#) on page 3-10.
 - [“Using Cookies in a Servlet”](#) on page 3-18.
 - [“Configuring Session Persistence”](#) on page 3-18
- Use of WebLogic services in your servlet—WebLogic Server provides a variety of services and APIs that you can use in your Web applications. These services include Java Database Connectivity (JDBC) drivers, JDBC database connection pools, Java Messaging Service (JMS), Enterprise JavaBeans (EJB), and Remote Method Invocation (RMI). For more information, see the following sections:
 - [“Using WebLogic Services from an HTTP Servlet”](#) on page 3-21.
 - [“Security”](#) on page 4-6.
 - [“Accessing Databases”](#) on page 3-22.

Complete HelloWorldServlet Example

This section provides the complete Java source code for the example used in the preceding procedure. The example is a simple servlet that provides a response to an HTTP request. Later in this document, this example is expanded to illustrate how to use HTTP parameters, cookies, and session tracking.

Listing 2-1 HelloWorldServlet.java

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

public class HelloWorldServlet extends HttpServlet {
    public void service(HttpServletRequest req,
                       HttpServletResponse res)
        throws IOException
    {
        // Must set the content type first
        res.setContentType("text/html");
        // Now obtain a PrintWriter to insert HTML into
        PrintWriter out = res.getWriter();

        out.println("<html><head><title>" +
                   "Hello World!</title></head>");
        out.println("<body><h1>Hello World!</h1></body></html>");
    }
}
```

You can find the source code and instructions for compiling and running all the examples used in this document in the `samples/examples/servlets` directory of your WebLogic Server distribution.

3 Programming Tasks

This section provides information about writing HTTP servlets in a WebLogic Server environment. The following topics are discussed:

- [Initializing a Servlet](#)
- [Providing an HTTP Response](#)
- [Retrieving Client Input](#)
- [Using Session Tracking from a Servlet](#)
- [Using Cookies in a Servlet](#)
- [Using WebLogic Services from an HTTP Servlet](#)
- [Accessing Databases](#)
- [Threading Issues in HTTP Servlets](#)
- [Dispatching Requests to Another Resource](#)

Initializing a Servlet

When a servlet is initialized, WebLogic Server executes the `init()` method of the servlet. Once the servlet has been initialized, it is not initialized again until you restart WebLogic Server or the servlet code is recompiled and reloaded. By overriding the `init()` method, your servlet can perform certain tasks, such as establishing database connections, when the servlet is initialized.

Normally, WebLogic Server initializes a servlet when the first request is made for the servlet. Subsequently, if the servlet is modified, the `destroy()` method is called on the existing version of the servlet. Then, after a request is made for the modified servlet, the `init()` method of the modified servlet is executed. For more information, see “[Servlet Development Tips](#)” on page 4-7.

Initializing a Servlet when WebLogic Server Starts

You can also configure WebLogic Server to initialize a servlet when the server starts by specifying the servlet class in the `<load-on-startup>` element in the Web Application deployment descriptor. For more information see “[Servlet element](#)” at http://e-docs.bea.com/wls/docs60/programming/web_xml.html#web_xml_servlet.

You can pass parameters to an HTTP servlet during initialization by defining these parameters in the Web Application containing the servlet. You can use these parameters to pass values to your servlet every time the servlet is initialized without having to rewrite the servlet. For more information see “[Writing Web Application Deployment Descriptors](#)” at <http://e-docs.bea.com/wls/docs60/programming/webappdeployment.html>.

For example, the following entries in the Web Application deployment descriptor define two initialization parameters: `greeting`, which has a value of `Welcome` and `person`, which has a value of `WebLogic Developer`.

```
<servlet>
  ...
  <init-param>
    <param-name>greeting</param-name>
    <param-value>Welcome</param-value>
    <description>The salutation</description>
  </init-param>
  <init-param>
    <param-name>person</param-name>
    <param-value>WebLogic Developer</param-value>
    <description>name</description>
  </init-param>
</servlet>
```

To retrieve initialization parameters by call the `getInitParameter(String name)` method from the parent `javax.servlet.GenericServlet` class. When passed the name of the parameter, this method returns the parameter's value as a `String`.

Overriding the `init()` Method

You can have your servlet execute tasks at initialization time by overriding the `init()` method. The following code fragment reads the `<init-param>` tags that define a greeting and a name in the Web Application deployment descriptor:

```
String defaultGreeting;
String defaultName;

public void init(ServletConfig config)
    throws ServletException {
    if ((defaultGreeting = getInitParameter("greeting")) == null)
        defaultGreeting = "Hello";

    if ((defaultName = getInitParameter("person")) == null)
        defaultName = "World";
}
```

The values of each parameter are stored in the class instance variables `defaultGreeting` and `defaultName`. Note that the first code tests to see if the parameters have null values and, if null values are returned, provides appropriate default values.

You can then use the `service()` method to use these variables in the response. For example:

```
out.print("<body><h1>");
out.println(defaultGreeting + " " + defaultName + "!");
out.println("</h1></body></html>");
```

The full source code and instructions for compiling, installing, and trying out an example called `HelloWorld2.java`, which illustrates the use of the `init()` method, can be found in the `samples/examples/servlets` directory of your WebLogic Server distribution.

Providing an HTTP Response

This section describes how to provide a response to the client in your HTTP servlet. All responses should be delivered using the `HttpServletResponse` object that is passed as a parameter to the `service()` method of your servlet.

1. Configure the `HttpServletResponse`.

Using the `HttpServletResponse` object, you can set several servlet properties that are translated into HTTP header information:

- *At a minimum*, set the content type using the `setContentType()` method before you obtain the output stream to which you write the page contents. For HTML pages, set the content type to `text/html`. For example:

```
res.setContentType("text/html");
```

- (optional) You can also use the `setContentType()` method to set the character encoding. For example:

```
res.setContentType("text/html;ISO-88859-4");
```

- Set header attributes using the `setHeader()` method. For dynamic responses, it is useful to set the “pragma” attribute to `no-cache`, which causes the browser to always reload the page and ensures the data is current. For example:

```
res.setHeader("Pragma", "no-cache");
```

2. Compose the HTML page.

The response that your servlet sends back to the client must look like regular HTTP content. For the most part, this means you must send back a response that is formatted as an HTML page. Your servlet sends back an HTTP response via an output stream that you obtain using the response parameter of the `service()` method. To send an HTTP response:

- a. Obtain an output stream using the `HttpServletResponse` object, using one of the methods shown in the following two examples:

- `PrintWriter out = res.getWriter();`
- `ServletOutputStream out = res.getOutputStream();`

You can use both `PrintWriter` and `ServletOutputStream` in the same servlet (or in another servlet that is included in a servlet). The output of both is written to the same buffer.

- b. Write the contents of the response to the output stream using the `print()` method. You can use HTML tags in these statements. For example:

```
out.print("<html><head><title>My Servlet</title>");
out.print("</head><body><h1>");
out.print("Welcome");
out.print("</h1></body></html>");
```

Do not close the output stream using the `close()` method and avoid flushing the contents of the stream. If you do not close or flush the output stream, WebLogic Server can take advantage of persistent HTTP connections, as described in the next step.

3. Optimize the response.

By default, WebLogic Server attempts to use HTTP persistent connections whenever possible. A persistent connection attempts to reuse the same HTTP TCP/IP connection for a series of communications between client and server. Application performance improves because a new connection need not be opened for each request. Persistent connections are useful for HTML pages containing many in-line images, where each requested image would otherwise require a new TCP/IP connection.

You can configure the amount of time that WebLogic Server keeps an HTTP connection open using the WebLogic Server Administration Console. For more information, see [“KeepAliveSecs”](#) (see <http://e-docs.bea.com/wls/docs60/ConsoleHelp/virtualhost.html#keepliveenabled>).

WebLogic Server must know the length of the HTTP response in order to establish a persistent connection and automatically adds a `Content-Length` property to the HTTP response header. In order to determine the content length, WebLogic Server must buffer the response. However, if your servlet explicitly flushes the `ServletOutputStream`, WebLogic Server is unable to determine the length of the response and therefore cannot use persistent connections. For this reason, you should avoid explicitly flushing the HTTP response in your servlets.

You may decide that, in some cases, it is better to flush the response early to display information in the client before the page has completed, for example to display a banner advertisement while some time-consuming page content is calculated. Conversely, you may want to increase the size of the buffer used by the servlet engine to accommodate a larger response before flushing the response. You can manipulate the size of the response buffer by using the related methods of the [javax.servlet.ServletResponse](#) interface (at

```
http://java.sun.com/products/servlet/2.2/javadoc/  
javax/servlet/ServletResponse.html).
```

Retrieving Client Input

The HTTP servlet API provides a clean interface for retrieving user input from Web pages.

An HTTP request from a Web browser can be accompanied by information besides the URL, such as information about the client, the browser, cookies, and user query parameters. Query parameters are used to carry user input from the browser, and are either appended to the URL address (the `GET` method) or included in the HTTP request body (the `POST` method).

HTTP servlets need not deal with these details; all the information in a request is made available through the `HttpServletRequest` object and can be accessed using the `request.getParameters()` method, regardless of the send method.

You can send query parameters from the client in a number of ways:

- Encode the parameters directly into the URL of a link on a page. This approach uses the `GET` method for sending parameters. The parameters are appended to the URL after a `?` character. Multiple parameters are separated by a `&` character. Parameters are always specified in *name=value* pairs so the order in which they are listed is not important. For example, you might include the following link in a Web page, which sends the parameter `color` with the value `purple` to an HTTP servlet called `ColorServlet`:

```
<a href=  
  "http://localhost:7001/myWebApp/ColorServlet?color=purple">  
  Click Here For Purple!</a>
```

- Manually enter the URL, with query parameters, into the browser location field. This is equivalent to clicking the link shown in the previous example.
- Query the user for input with an HTML form. The contents of each user input field on the form are sent as query parameters when the user clicks the form's submit button. Specify the method used by the form to send the query parameters (`POST` or `GET`) in the `<FORM>` tag using the `METHOD="GET|POST"` attribute.

Query parameters are always sent in *name=value* pairs, and are accessed through the `HttpServletRequest` object. You can obtain an `Enumeration` of all parameter names in a query, and fetch each parameter value using its parameter name. A parameter usually has only one value, but it can also hold an array of values. Parameter values are always interpreted as `Strings`, so you may need to cast them to a more appropriate type.

The following sample from a `service()` method examines query parameter names and their values from a form. Note that `request` is the `HttpServletRequest` object.

```
Enumeration params = request.getParameterNames();
String paramName = null;
String[] paramValues = null;

while (params.hasMoreElements()) {
    paramName = (String) params.nextElement();
    paramValues = request.getParameterValues(paramName);
    System.out.println("\nParameter name is " + paramName);
    for (int i = 0; i < paramValues.length; i++) {
        System.out.println("    value " + i + " is " +
            paramValues[i].toString());
    }
}
```

Other Methods for Using the HTTP Request

This section defines the methods of the `javax.servlet.HttpServletRequest` interface that you can use to get data from the request object. You should keep the following limitations in mind:

- You cannot read request parameters using any of the `getParameter()` methods described in this section and then attempt to read the request with the `getInputStream()` method.
- You cannot read the request with `getInputStream()` and then attempt to read request parameters with one of the `getParameter()` methods.

If you attempt either of these procedures, an `illegalStateException` is thrown.

`HttpServletRequest.getMethod()`

Allows you to determine the request method, such as GET or POST.

`HttpServletRequest.getQueryString()`
Allows you to access the query string. (The remainder of the requested URL, following the ? character.)

`HttpServletRequest.getParameter()`
Returns the value of a parameter.

`HttpServletRequest.getParameterNames()`
Returns an array of parameter names.

`HttpServletRequest.getParameterValues()`
Returns an array of values for a parameter.

`HttpServletRequest.getInputStream()`
Reads the body of the request as binary data. If you call this method after reading the request parameters with `getParameter()`, `getParameterNames()`, or `getParameterValues()`, an `IllegalStateException` is thrown.

Example: Retrieving Input by Using Query Parameters

In this example, the `HelloWorld2.java` servlet example is modified to accept a username as a query parameter, in order to display a more personal greeting. (For the complete code, see the `HelloWorld3.java` servlet example, located in the `samples/examples/servlets` directory of your WebLogic Server distribution.) The `service()` method is shown here:

Listing 3-1

```
public void service(HttpServletRequest req,
                    HttpServletResponse res)
    throws IOException
{
    String name, paramName[];
    if ((paramName = req.getParameterValues("name"))
        != null) {
        name = paramName[0];
    }
    else {
        name = defaultName;
    }
}
```

```
// Set the content type first
res.setContentType("text/html");
// Obtain a PrintWriter as an output stream
PrintWriter out = res.getWriter();

out.print("<html><head><title>" +
        "Hello World!" + </title></head>");
out.print("<body><h1>");
out.print(defaultGreeting + " " + name + "!");
out.print("</h1></body></html>");
}
```

The `getParameterValues()` method retrieves the value of the `name` parameter from the HTTP query parameters. You retrieve these values in an array of type `String`. A single value for this parameter is returned and is assigned to the first element in the name array. If the parameter is not present in the query data, `null` is returned; in this case, `name` is assigned to the default name that was read from the `<init-param>` by the `init()` method.

Your servlet code should not presume that parameters are included in an HTTP request. Because the `getParameter()` method has been deprecated, you might be tempted to shorthand the `getParameterValues()` method by tagging an array subscript to the end. However, this method can return `null` if the specified parameter is not available, resulting in a `NullPointerException`.

For example, the following code will trigger a `NullPointerException`:

```
String myStr = req.getParameterValues("paramName")[0];
```

Instead, use the following code:

```
if ((String myStr[] =
    req.getParameterValues("paramName"))!=null) {
    // Now you can use the myStr[0];
}
else {
    // paramName was not in the query parameters!
}
```

Using Session Tracking from a Servlet

Session tracking enables you to track a user's progress over multiple servlets or HTML pages, which, by nature, are stateless. A *session* is defined as a series of related browser requests that come from the same client during a certain time period. Session tracking ties together a series of browser requests—think of these requests as pages—that may have some meaning as a whole, such as a shopping cart application.

A History of Session Tracking

Before session tracking matured conceptually, developers tried to build state into their pages by stuffing information into hidden fields on a page or embedding user choices into URLs used in links with a long string of appended characters. You can see good examples of this at most search engine sites, many of which still depend on CGI. These sites track user choices with URL parameter *name=value* pairs that are appended to the URL, after the reserved HTTP character `?`. This practice can result in a very long URL that the CGI script must carefully parse and manage. The problem with this approach is that you cannot pass this information from session to session. Once you lose control over the URL—that is, once the user leaves one of your pages—the user information is lost forever.

Later, Netscape introduced browser *cookies*, which enable you to store user-related information about the client for each server. However, some browsers still do not fully support cookies, and some users prefer to turn off the cookie option in their browsers. Another factor that should be considered is that most browsers limit the amount of data that can be stored with a cookie.

The HTTP servlet specification defines a solution that allows the server to store user details on the server, and protects your code from the complexities of tracking sessions. Your servlets can use an `HttpSession` object to track a user's input over the span of a single session and to share session details among multiple servlets.

Tracking a Session with an HttpSession Object

According to the Java Servlet API, which WebLogic Server implements and supports, each servlet can access a server-side session by using its `HttpSession` object. You can access an `HttpSession` object in the `service()` method of the servlet using the `HttpServletRequest` object, as shown in the following example using the variable `request`:

```
HttpSession session = request.getSession(true);
```

An `HttpSession` object is created if one does not already exist for that client when the `request.getSession(true)` method is called with the argument `true`. The session object lives on WebLogic Server for the lifetime of the session, during which the session object accumulates information related to that client. Your servlet adds or removes information from the session object as necessary. A session is associated with a particular client. Each time the client visits your servlet, the same associated `HttpSession` object is retrieved when the `getSession()` method is called.

For more details on the methods supported by the `HttpSession`, refer to the [HttpServlet API](http://java.sun.com/j2ee/j2sdkee/techdocs/api/javax/servlet/http/HttpSession.html) at <http://java.sun.com/j2ee/j2sdkee/techdocs/api/javax/servlet/http/HttpSession.html>.

In the following example, the `service()` method counts the number of times a user requests the servlet during a session.

```
public void service(HttpServletRequest request,
                   HttpServletResponse response)
    throws IOException
{
    // Get the session and the counter param attribute
    HttpSession session = request.getSession (true);
    Integer ival = (Integer)
        session.getAttribute("simplesession.counter");
    if (ival == null) // Initialize the counter
        ival = new Integer (1);
    else // Increment the counter
        ival = new Integer (ival.intValue () + 1);
    // Set the new attribute value in the session
    session.setAttribute("simplesession.counter", ival);
    // Output the HTML page
    out.print("<HTML><body>");
    out.print("<center> You have hit this page ");
    out.print(ival + " times!");
    out.print("</body></html>");
}
```

Lifetime of a Session

A session tracks the selections of a user over a series of pages in a single transaction. A single transaction may consist of several tasks, such as searching for an item, adding it to a shopping cart, and then processing a payment. A session is transient, and its lifetime ends when one of the following occurs:

- A user leaves your site and the user's browser does not accept cookies.
- A user quits the browser.
- The session is timed out due to inactivity.
- The session is completed and invalidated by the servlet.
- The user logs out and is invalidated by the servlet.

For more persistent, long-term storage of data your servlet should write details to a database using JDBC or EJB and associate the client with this data using a long-lived cookie and/or username and password. Although this document describes that sessions use cookies and persistence internally, *you should not use sessions as a general mechanism for storing data about a user.*

How Session Tracking Works

How does WebLogic Server know which session is associated with each client? When an `HttpSession` is created in a servlet, it is associated with a unique ID. The browser must provide this session ID with its request in order for the server to find the session data again. The server attempts to store this ID by setting a cookie on the client. Once the cookie is set, each time the browser sends a request to the server it includes the cookie containing the ID. The server automatically parses the cookie and supplies the session data when your servlet calls the `getSession()` method.

If the client does not accept cookies, the only alternative is to encode the ID into the URL links in the pages sent back to the client. For this reason, you should always use the `encodeURL()` method when you include URLs in your servlet response. WebLogic Server knows whether the browser accepts cookies and does not unnecessarily encode URLs. WebLogic automatically parses the session ID from an encoded URL and retrieves the correct session data when you call the `getSession()`

method. Using the `encodeURL()` method ensures no disruption to your servlet code, regardless of the procedure used to track sessions. For more information, see [“Using URL Rewriting”](#) on page 3-15.

Detecting the Start of a Session

After you obtain a session using the `getSession(true)` method, you can tell whether the session has just been created by calling the `HttpSession.isNew()` method. If this method returns `true`, then the client does not already have a valid session, and at this point it is unaware of the new session. The client does not become aware of the new session until a reply is posted back from the server.

Design your application to accommodate new or existing sessions in a way that suits your business logic. For example, your application might redirect the client’s URL to a login/password page if you determine that the session has not yet started, as shown in the following code example:

```
HttpSession session = request.getSession(true);
if (session.isNew()) {
    response.sendRedirect(welcomeURL);
}
```

On the login page, provide an option to log in to the system or create a new account. You can also specify a login page in your Web Application. For more information, see [login-config](#) at http://e-docs.bea.com/wls/docs60/programming/web_xml.html#login-config.

Setting and Getting Session Name/Value Attributes

You can store data in an `HttpSession` object using `name=value` pairs. Data stored in a session is available through the session. To store data in a session, use these methods from the `HttpSession` interface:

```
getAttribute()
getAttributeNames()
setAttribute()
removeAttribute()
```

The following code fragment shows how to get all the existing `name=value` pairs:

```
Enumeration sessionNames = session.getAttributeNames();
String sessionName = null;
Object sessionValue = null;

while (sessionNames.hasMoreElements()) {
    sessionName = (String)sessionNames.nextElement();
    sessionValue = session.getAttribute(sessionName);
    System.out.println("Session name is " + sessionName +
        ", value is " + sessionValue);
}
```

To add or overwrite a named attribute, use the `setAttribute()` method. To remove a named attribute altogether, use the `removeAttribute()` method.

Note: You can add any Java descendant of `Object` as a session attribute and associate it with a name. However, if you are using session persistence, your attribute *value* objects must implement `java.io.Serializable`.

Logging Out and Ending a Session

If your application deals with sensitive information you might consider offering the ability to log out of the session. This is a common feature when using shopping carts and Internet email accounts. When the same browser returns to the service, the user must log back in to the system. To log a user out of a session, invalidate the current session by calling the following method:

```
session.invalidate()
```

Do not reference an invalidated session after calling this method. If you do, an `IllegalStateException` is thrown. The next time a user visits your servlet from the same browser, the session data will be missing, and a new session will be created when you call the `getSession(true)` method. At that time you can send the user to the login page again.

Configuring Session Tracking

WebLogic Server provides many configurable attributes that determine how WebLogic Server handles session tracking. For details about configuring these session tracking attributes, see “[Session descriptor](#)” at http://e-docs.bea.com/wls/docs60/programming/weblogic_xml.html#session-descriptor.

Using URL Rewriting

In some situations, a browser may not accept cookies, which means that session tracking using cookies is not possible. URL rewriting is a workaround to this scenario that can be substituted automatically when WebLogic Server detects that the browser does not accept cookies. URL rewriting involves encoding the session ID into the hyperlinks on the Web pages that your servlet sends back to the browser. When the user subsequently clicks these links, WebLogic Server extracts the ID from the URL and finds the appropriate `HttpSession`. Then you use the `getSession()` method to access session data.

To enable URL rewriting in WebLogic Server, set the `UrlRewritingEnabled` attribute to true in the [Session descriptor](#) element of the WebLogic-specific deployment descriptor (at http://e-docs.bea.com/wls/docs60/programming/weblogic_xml.html#session-descriptor).

To make sure your code correctly handles URLs in order to support URL rewriting, consider the following guidelines:

- You should avoid writing a URL straight to the output stream, as shown here:

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

Instead, use the `HttpServletResponse.encodeURL()` method. For example:

```
out.println("<a href=\"" +  
    + response.encodeURL("myshop/catalog.jsp")  
    + "\">catalog</a>");
```

- Calling the `encodeURL()` method determines if the URL needs to be rewritten and, if necessary, rewrites the URL by including the session ID in the URL.

- Encode URLs that send redirects, as well as URLs that are returned as a response to WebLogic Server. For example:

```
if (session.isNew())  
    response.sendRedirect(response.encodeRedirectUrl(welcomeURL));
```

WebLogic Server uses URL rewriting when a session is new, even if the browser accepts cookies, because the server cannot determine, during the first visit of a session, whether the browser accepts cookies.

Your servlet may determine whether a given session was returned from a cookie by checking the Boolean returned from the

`HttpServletRequest.isRequestedSessionIdFromCookie()` method. Your application may respond appropriately, or it may simply rely on URL rewriting by WebLogic Server.

URL Rewriting and Wireless Access Protocol (WAP)

If you are writing a WAP application, you must use URL rewriting because the WAP protocol does not support cookies. In addition, some WAP devices impose a 128-character limit (including parameters) on the length of a URL, which limits the amount of data that can be transmitted using URL rewriting. To allow more space for parameters, you can limit the size of the session ID that is randomly generated by WebLogic Server by specifying the number of bytes with the `IDLength` attribute in the [session-descriptor](#) element of the WebLogic-specific deployment descriptor, `weblogic.xml` (see http://e-docs.bea.com/wls/docs60/programming/weblogic_xml.html#session-descriptor).

The minimum value is 8 bytes; the default value is 52 bytes; the maximum value is `Integer.MAX_VALUE`.

Making Sessions Persistent

You can set up WebLogic Server to record session data in a persistent store. If you are using session persistence, you can expect the following characteristics:

- Good failover, because sessions are saved when servers fail.
- Better load balancing, because any server can handle requests for any number of sessions, and use caching to optimize performance. For more information, see

the `cacheEntries` property, under “[Configuring session persistence](#)” at http://e-docs.bea.com/wls/docs60/adminguide/config_web_app.html#session-persistence.

- Sessions can be shared across clustered WebLogic Servers. Note that session persistence is no longer a requirement in a WebLogic Cluster. Instead, you can use in-memory replication of state. For more information, see [Using WebLogic Server Clusters](#) at <http://e-docs.bea.com/wls/docs60/cluster/index.html>.
- For customers who want the highest in servlet session persistence, JDBC-based persistence is the best choice. For customers who want to sacrifice some amount of session persistence in favor of drastically better performance, in-memory replication is the appropriate choice. JDBC-based persistence is noticeably slower than in-memory replication. In some cases, in-memory replication has outperformed JDBC-based persistence for servlet sessions by a factor of eight.
- It is legal to put any kind of Java object into a session, but for file, JDBC, and in-memory replication, only objects that are `java.io.Serializable` can be stored in a session. For more information, see “[Configuring session persistence](#)” at http://e-docs.bea.com/wls/docs60/adminguide/config_web_app.html#session-persistence.

Scenarios to Avoid When Using Sessions

Do not use session persistence for storing long-term data between sessions. In other words, do not rely on a session still being active when a client returns to a site at some later date. Instead, your application should record long-term or important information in a database.

Sessions are not a convenience wrapper around cookies. Do not attempt to store long-term or limited-term client data in a session. Instead, your application should create and set its own cookies on the browser. Examples include an auto-login feature that allows a cookie to live for a long period, or an auto-logout feature that allows a cookie to expire after a short period of time. Here, you should not attempt to use HTTP sessions. Instead, you should write your own application-specific logic.

Use Serializable Attribute Values

When you use persistent sessions, all attribute value objects that you add to the session must implement `java.io.Serializable`. For more details on writing serializable classes, refer to the online java tutorial about [serializable objects](http://java.sun.com/docs/books/tutorial/essential/io/providing.html) at <http://java.sun.com/docs/books/tutorial/essential/io/providing.html>. If you add your own serializable classes to a persistent session, make sure that each instance variable of your class is also serializable. Otherwise, you can declare it as `transient`, and WebLogic Server does not attempt to save that variable to persistent storage. One common example of an instance variable that must be made `transient` is the `HttpSession` object. (See the notes on using serialized objects in sessions in the section “[Making Sessions Persistent](#)” on page 3-16.)

Configuring Session Persistence

For details about setting up persistent sessions, see “[Configuring session persistence](#)” at http://e-docs.bea.com/wls/docs60/adminguide/config_web_app.html#session-persistence.

Using Cookies in a Servlet

A cookie is a piece of information that the server asks the client browser to save locally on the user’s disk. Each time the browser visits the same server, it sends all cookies relevant to that server with the HTTP request. Cookies are useful for identifying clients as they return to the server.

Each cookie has a name and a value. A browser that supports cookies generally allows each server domain to store up to 20 cookies with a size of up to 4k per cookie.

Setting Cookies in an HTTP servlet

To set a cookie on a browser, create the cookie, give it a value, and add it to the `HttpServletResponse` object that is the second parameter in your servlet’s service method. For example:

```
Cookie myCookie = new Cookie("ChocolateChip", "100");
myCookie.setMaxAge(Integer.MAX_VALUE);
response.addCookie(myCookie);
```

This examples shows how to add a cookie called `ChocolateChip` with a value of `100` to the browser client when the response is sent. The expiration of the cookie is set to the largest possible value, which effectively makes the cookie last forever. Because cookies accept only string-type values, you should cast to and from the desired type that you want to store in the cookie. When using EJBs, a common practice is to use the *home handle* of an EJB instance for the cookie value and to store the user's details in the EJB for later reference.

Retrieving Cookies in an HTTP Servlet

You can retrieve a cookie object from the `HttpServletRequest` that is passed to your servlet as an argument to the `service()` method. The cookie itself is presented as a `javax.servlet.http.Cookie` object.

In your servlet code, you can retrieve all the cookies sent from the browser by calling the `getCookies()` method. For example:

```
Cookie[] cookies = request.getCookies();
```

This method returns an array of all cookies sent from the browser, or `null` if no cookies were sent by the browser. Your servlet must process the array in order to find the correct named cookie. You can get the name of a cookie using the `Cookie.getName()` method. It is possible to have more than one cookie with the same name, but different path attributes. If your servlets set multiple cookies with the same names, but different path attributes, you also need to compare the cookies using the `Cookie.getPath()` method. The following code illustrates how to access the details of a cookie sent from the browser. It assumes that all cookies sent to this server have unique names, and that you are looking for a cookie called `ChocolateChip` that may have been set previously in a browser client. For example:

```
Cookie[] cookies = request.getCookies();
boolean cookieFound = false;

for(int i=0; i < cookies.length; i++) {
    thisCookie = cookies[i];
    if (thisCookie.getName().equals("ChocolateChip")) {
        cookieFound = true;
        break;
    }
}
```

```
    }  
  }  
  
  if (cookieFound) {  
    // We found the cookie! Now get its value  
    int cookieOrder = String.parseInt(thisCookie.getValue());  
  }  
}
```

For more details on cookies, see:

- [The Cookie API](http://java.sun.com/j2ee/j2sdkee/techdocs/api/javax/servlet/http/Cookie.html) at <http://java.sun.com/j2ee/j2sdkee/techdocs/api/javax/servlet/http/Cookie.html>
- [The Java Tutorial: Using Cookies](http://java.sun.com/docs/books/tutorial/servlets/client-state/cookies.html) at <http://java.sun.com/docs/books/tutorial/servlets/client-state/cookies.html>

Using Cookies that Are Transmitted by Both HTTP and HTTPS

Because HTTP and HTTPS requests are sent to different ports, some browsers may not include the cookie sent in an HTTP request with a subsequent HTTPS request (or vice-versa). This may cause new sessions to be created when servlet requests alternate between HTTP and HTTPS. To ensure that all cookies set by a specific domain are sent to the server every time a request in a session is made, set the `CookieDomain` attribute to the name of the domain. Set the `CookieDomain` attribute with the `<session-descriptor>` element of the WebLogic-specific deployment descriptor (`weblogic.xml`) for the Web Application that contains your servlet. For example:

```
<session-descriptor>  
  <session-param>  
    <param-name>CookieDomain</param-name>  
    <param-value>mydomain.com</param-value>  
  </session-param>  
</session-descriptor>
```

The `CookieDomain` attribute instructs the browser to include the proper cookie(s) for all requests to hosts in the domain specified by `mydomain.com`. For more information about this property or configuring session cookies, see “[Setting Up Session Management](#)” at http://e-docs.bea.com/wls/docs60/adminguide/config_web_app.html#session-management.

Application Security and Cookies

Using cookies that enable automatic account access on a machine is convenient, but can be undesirable from a security perspective. When designing an application that uses cookies, follow these guidelines:

- Do not assume that a cookie is always correct for a user. Sometimes machines are shared or the same user may want to access a different account.
- Allow your users to make a choice about leaving cookies on the server. On shared machines, users may not want to leave automatic logins for their account. Do not assume that users know what a cookie is; instead, ask a question like:
`Automatically login from this computer?`
- Always ask for passwords from users logging on to obtain sensitive data. Unless a user requests otherwise, you can store this preference and the password in the user's session data. Configure the session cookie to expire when the user quits the browser.

Using WebLogic Services from an HTTP Servlet

When you write an HTTP servlet, you have access to many of the rich features of WebLogic Server, such as JNDI, EJB, JDBC, and JMS.

The following documents provide additional information about these features:

- [Programming WebLogic EJB](http://e-docs.bea.com/wls/docs60/ejb/index.html) at `http://e-docs.bea.com/wls/docs60/ejb/index.html`
- [Programming WebLogic JDBC](http://e-docs.bea.com/wls/docs60/jdbc/index.html) at `http://e-docs.bea.com/wls/docs60/jdbc/index.html`
- [Programming WebLogic JNDI](http://e-docs.bea.com/wls/docs60/jndi/index.html) at `http://e-docs.bea.com/wls/docs60/jndi/index.html`
- [Programming WebLogic JMS](http://e-docs.bea.com/wls/docs60/jms/index.html) at `http://e-docs.bea.com/wls/docs60/jms/index.html`

Accessing Databases

WebLogic Server supports the use of Java Database Connectivity (JDBC) from server-side Java classes, including servlets. JDBC allows you to execute SQL queries from a Java class and to process the results of those queries. For more information on JDBC and WebLogic Server, see *Using WebLogic JDBC* at <http://e-docs.bea.com/wls/docs60/jdbc/index.html>.

There three ways you can use JDBC in servlets, which are discussed in detail in the following sections:

- “Connecting to a Database Using a JDBC Connection Pool” on page 3-22.
- “Connecting to a Database Using a DataSource Object” on page 3-24.
- “Connecting Directly to a Database Using a JDBC Driver” on page 3-25

Connecting to a Database Using a JDBC Connection Pool

A connection pool is a named group of identical JDBC connections to a database that are created when the connection pool is registered, usually when starting WebLogic Server. Your servlets “borrow” a connection from the pool, use it, and then return it to the pool by closing it. This process is far more efficient than creating a new connection for every client each time the client needs to access the database. Another advantage is that you do not need to include details about the database in your servlet code.

When connecting to a JDBC connection pool, use one of the following multitier JDBC drivers:

- Pool driver, used for most server-side operations:
 - Driver URL: `jdbc:weblogic:pool`
 - Driver package name: `weblogic.jdbc.pool.Driver`
- JTS pool driver, used when database operations require transactional support.
 - Driver URL: `jdbc:weblogic:jts`
 - Driver package name: `weblogic.jdbc.jts.Driver`

Using a Connection Pool in a Servlet

The following example demonstrates how to use a database connection pool from a servlet.

1. Load the pool driver and cast it to `java.sql.Driver`. The full pathname of the driver is `weblogic.jdbc.pool.Driver`. For example:

```
Driver myDriver = (Driver)
    Class.forName("weblogic.jdbc.pool.Driver").newInstance();
```

2. Create a connection using the URL for the driver, plus (optionally) the name of the registered connection pool. The URL of the pool driver is `jdbc:weblogic:pool`.

You can identify the pool in either of two ways:

- Specify the name of the connection pool in a `java.util.Properties` object using the key `connectionPoolID`. For example:

```
Properties props = new Properties();
props.put("connectionPoolID", "myConnectionPool");
Connection conn =
    myDriver.connect("jdbc:weblogic:pool", props);
```

- Add the name of the pool to the end of the URL. In this case you do not need a `Properties` object unless you are setting a username and password for using a connection from the pool. For example:

```
Connection conn =
    myDriver.connect("jdbc:weblogic:pool:myConnectionPool",
    null);
```

Note that the `Driver.connect()` method is used in these examples instead of the `DriverManger.getConnection()` method. Although you may use `DriverManger.getConnection()` to obtain a database connection, we recommend that you use `Driver.connect()` because this method is not synchronized and provides better performance.

Note that the `Connection` returned by `connect()` is an instance of `weblogic.jdbc.pool.Connection`.

3. Call the `close()` method on the `Connection` object when you have finished with your JDBC calls, so that the connection is properly returned to the pool. A good coding practice is to create the connection in a `try` block and then close the connection in a `finally` block, to make sure the connection is closed in all cases.

```
conn.close();
```

Connecting to a Database Using a DataSource Object

A `DataSource` is a server-side object that references a connection pool. The connection pool registration defines the JDBC driver, database, login, and other parameters associated with a database connection. `DataSource` objects and connection pools are created using the Administration Console. *Using a `DataSource` object is recommended when creating J2EE-compliant applications.*

Using a DataSource in a Servlet

1. Register a connection pool using the Administration Console. For more information, see “[Create a Connection Pool](http://e-docs.bea.com/wls/docs60/ConsoleHelp/jdbcconnectionpool.html)” at <http://e-docs.bea.com/wls/docs60/ConsoleHelp/jdbcconnectionpool.html>.
2. Register a `DataSource` object that points to the connection pool. For more information, see “[JDBC DataSources](http://e-docs.bea.com/wls/docs60/ConsoleHelp/jdbcdatasource.html)” at <http://e-docs.bea.com/wls/docs60/ConsoleHelp/jdbcdatasource.html>.

3. Look up the `DataSource` object in the JNDI tree. For example:

```
Context ctx = null;

// Get a context for the JNDI look up
ctx = new InitialContext(ht);

// Look up the DataSource object
javax.sql.DataSource ds
    = (javax.sql.DataSource) ctx.lookup ("myDataSource");
```

4. Use the `DataSource` to create a JDBC connection. For example:

```
java.sql.Connection conn = ds.getConnection();
```

5. Use the connection to execute SQL statements. For example:

```
Statement stmt = conn.createStatement();
stmt.execute("select * from emp");
. . .
```

Connecting Directly to a Database Using a JDBC Driver

Connecting directly to a database is the least efficient way of making a database connection because a new database connection must be established for each request. You can use any JDBC driver to connect to your database. BEA provides JDBC drivers for Oracle, Microsoft SQL Server, and Informix. For more information, see [Programming WebLogic JDBC](#) at <http://e-docs.bea.com/wls/docs60/jdbc/index.html>.

Threading Issues in HTTP Servlets

When you design a servlet, you should consider how the servlet is invoked by WebLogic Server under high load. It is inevitable that more than one client will hit your servlet simultaneously. Therefore, write your servlet code to guard against sharing violations on shared resources or instance variables. The following tips can help you to design around this issue.

SingleThreadModel

An instance of a class that implements the `SingleThreadModel` is guaranteed not to be invoked by multiple threads simultaneously. Multiple instances of a `SingleThreadModel` servlet are used to service simultaneous requests, each running in a single thread.

To use the `SingleThreadModel` efficiently, WebLogic Server creates a pool of servlet instances for each servlet that implements `SingleThreadModel`. WebLogic Server creates the pool of servlet instances when the first request is made to the servlet and increments the number of servlet instances in the pool as needed.

The attribute “[SingleThreadedModelPoolSize](#)” (see <http://e-docs.bea.com/wls/docs60/ConsoleHelp/webappcomponent.html#SingleThreadedModelPoolSize>), set using the WebLogic Server Administration Console, specifies the initial number of servlet instances that are created when the servlet is first requested. Set this attribute to the average number of concurrent requests that you expect your `SingleThreadModel` servlets to handle.

When designing your servlet, consider how you use shared resources outside of the servlet class such as file and database access. Because there are multiple instances of servlets that are identical, and may use exactly the same resources, there are still synchronization and sharing issues that must be resolved, even if you do implement the `SingleThreadModel`.

Shared Resources

We recommend that shared-resource issues be handled on an individual servlet basis. Consider the following guidelines:

- Wherever possible, avoid synchronization, because it causes subsequent servlet requests to bottleneck until the current thread completes.
- Define variables that are specific to each servlet request within the scope of the service methods. Local scope variables are stored on the stack and, therefore, are not shared by multiple threads running within the same method, which avoids the need to be synchronized.
- Access to external resources should be synchronized on a Class level, or encapsulated in a transaction.

Dispatching Requests to Another Resource

This section provides an overview of commonly used methods for dispatching requests.

A servlet can pass on a request to another resource, such as a servlet, JSP, or HTML page. This process is referred to as *request dispatching*. When you dispatch requests, you use either the `include()` or `forward()` method of the `RequestDispatcher` interface. There are limitations regarding when output can be written to the response object using the `forward()` or `include()` methods. These limitations are also discussed in this section.

For a complete discussion of request dispatching, see section 8.1 of the [Servlet 2.2 specification](http://java.sun.com/products/servlet/download.html#specs) (see <http://java.sun.com/products/servlet/download.html#specs>) from Sun Microsystems.

By using the `RequestDispatcher`, you can avoid sending an HTTP-redirect response back to the client. The `RequestDispatcher` passes the HTTP request to the requested resource.

To dispatch a request to a particular resource:

1. Get a reference to a `ServletContext`:

```
ServletContext sc = getServletConfig().getServletContext();
```

2. Look up the `RequestDispatcher` object using one of the following methods:

- `RequestDispatcher rd = sc.getRequestDispatcher(String path);`
path should be relative to the root of the Web Application.

- `RequestDispatcher rd = sc.getNamedDispatcher(String name);`
Replace *name* with the name assigned to the servlet in a Web Application deployment descriptor with the `<servlet-name>` element. For details, see “[Servlet element](#)” at

```
http://e-docs.bea.com/wls/docs60/programming/  
web_xml.html#web_xml_servlet
```

- `RequestDispatcher rd =
 ServletRequest.getRequestDispatcher(String path);`

This method returns a `RequestDispatcher` object and is similar to the `ServletContext.getRequestDispatcher(String path)` method except that it allows the *path* specified to be relative to the current servlet. If the path begins with a `/` character it is interpreted to be relative to the Web Application.

You can obtain a `RequestDispatcher` for any HTTP resource within a Web Application, including HTTP Servlets, JSP pages, or plain HTML pages by requesting the appropriate URL for the resource in the `getRequestDispatcher()` method. Use the returned `RequestDispatcher` object to forward the request to another servlet.

3. Forward or include the request using the appropriate method:

- `rd.forward(request, response);`
- `rd.include(request, response);`

These methods are discussed in the next two sections.

Forwarding a Request

Once you have the correct `RequestDispatcher`, your servlet forwards a request using the `RequestDispatcher.forward()` method, passing `HttpServletRequest` and `HttpServletResponse` as arguments. If you call this method when output has already been sent to the client an `IllegalStateException` is thrown. If the response buffer contains pending output that has not been committed, the buffer is reset.

The servlet must not attempt to write any previous output to the response. If the servlet retrieves the `ServletOutputStream` or the `PrintWriter` for the response before forwarding the request, an `IllegalStateException` is thrown.

All other output from the original servlet is ignored after the request has been forwarded.

Including a Request

Your servlet can include the output from another resource by using the `RequestDispatcher.include()` method, and passing `HttpServletRequest` and `HttpServletResponse` as arguments. When you include output from another resource, the included resource has access to the request object.

The included resource can write data back to the `ServletOutputStream` or `Writer` objects of the response object and then can either add data to the response buffer or call the `flush()` method on the response object. Any attempt to set the response status code or to set any HTTP header information from the included servlet response is ignored.

In effect, you can use the `include()` method to mimic a “server-side-include” of another HTTP resource from your servlet code.

4 Administration and Configuration

This section provides a brief overview of administration and configuration tasks for WebLogic HTTP servlets. For a complete discussion of servlet administration and configuration see “Using WebLogic Server HTTP/Web Services” in the *WebLogic Server Administration Guide* at

http://e-docs.bea.com/wls/docs60/adminguide/config_web_app.html.

This section discusses the following topics:

- [Overview](#)
- [URL](#)
- [Directory Structure](#)
- [Security](#)
- [Servlet Development Tips](#)
- [Clustering Servlets](#)

Overview

Consistent with the Java 2 Enterprise Edition standard, HTTP servlets are deployed as part of a *Web Application*. A Web Application is a grouping of application components, such as servlet classes, JavaServer Pages (JSP), static HTML pages, images, and utility classes.

In a Web Application the components are deployed using a standard directory structure. This directory structure can be archived into a file called a `.war` file and then deployed on WebLogic Server. Information about the resources and operating parameters of a Web Application are defined using two *deployment descriptors*, which are packaged with the Web Application.

Using Deployment Descriptors to Configure and Deploy Servlets

The first deployment descriptor, `web.xml`, is defined in the Servlet 2.2 specification from Sun Microsystems and provides a standardized format that describes the Web Application. The second deployment descriptor, `weblogic.xml`, is a WebLogic-specific deployment descriptor that maps resources defined in the `web.xml` file to resources available in WebLogic Server, defines JSP behavior, and defines HTTP session parameters.

web.xml (Web Application Deployment Descriptor)

In the Web Application deployment descriptor you define the following attributes for HTTP servlets:

- Servlet name
- Java class of the servlet
- servlet initialization parameters
- Whether or not the `init()` method of the servlet is executed when WebLogic Server starts
- URL pattern which, if matched, will call this servlet
- Security
- MIME type
- Error pages
- References to EJBs
- References to other resources

For a complete discussion of creating `web.xml` file see [“Writing Web Application Deployment Descriptors”](#) at

<http://e-docs.bea.com/wls/docs60/programming/webappdeployment.html>.

weblogic.xml (Weblogic-Specific Deployment Descriptor)

In the WebLogic-specific deployment descriptor you define the following attributes for HTTP servlets:

- HTTP session configuration
- Cookie configuration
- EJB resource mapping
- JSP Configuration

For a complete discussion of creating the `weblogic.xml` file see [“Writing Web Application Deployment Descriptors”](#) at

<http://e-docs.bea.com/wls/docs60/programming/webappdeployment.html>.

WebLogic Server Administration Console

Use the WebLogic Server Administration Console to set the following parameters:

- HTTP parameters
- Log files
- URL rewriting
- Keep alive
- Default MIME types
- Clustering parameters
- URL mapping for virtual hosting

For more information see the following resources:

- Administration Console: [“WebApp Component”](#) at <http://e-docs.bea.com/wls/docs60/ConsoleHelp/webappcomponent.html>

- Administration Console: “[Virtual Hosts](http://e-docs.bea.com/wls/docs60/ConsoleHelp/virtualhost.html)” at <http://e-docs.bea.com/wls/docs60/ConsoleHelp/virtualhost.html>

Directory Structure

Use the following directory structure for all Web Applications:

```

WebApplicationRoot/(Publicly available files, such as
|   .jsp, .html, .jpg, .gif)
|
+WEB-INF/+
|
+ classes/(directory containing
|           Java classes including
|           servlets used by the
|           Web Application)
|
+ lib/(directory containing
|      jar files used by the
|      Web Application)
|
+ web.xml
|
+ weblogic.xml

```

URL

The URL used to reference a servlet in a Web Application is constructed as follows:

`http://myHostName:port/myContextPath/myRequest/?myRequestParameters`

The components of this URL are defined as follows:

myHostName

the DNS name mapped to the Web Server defined in the WebLogic Server Administration Console

this portion of the URL can be replaced with *host:port*, where *host* is the name of the machine running WebLogic Server and *port* is the port at which WebLogic Server is listening for requests

port

the port at which WebLogic Server is listening for requests

myContextPath

the name of the Web Application defined in the WebLogic Server Administration Console

myRequest

the name of the servlet as defined in the `web.xml` file

myRequestParameters

optional HTTP request parameters encoded in the URL, which can be read by an HTTP servlet.

Security

Security for servlets is defined in the context of the Web Application containing the servlet. Security can be handled by WebLogic Server, or it can be incorporated programmatically into your servlet classes.

For more information see “[Configuring Security in Web Applications](http://e-docs.bea.com/wls/docs60/adminguide/config_web_app.html#configure-security)” at http://e-docs.bea.com/wls/docs60/adminguide/config_web_app.html#configure-security.

Authentication

You can incorporate user authentication into your servlets using any of the following three techniques:

- BASIC—uses the browser to collect a username and password.
- FORM—uses HTML forms to collect a username and password
- Client Certificate—uses digital certificates to authenticate the user. For more information, see “[Digital Certificates](http://e-docs.bea.com/wls/docs60/security/concepts.html#concepts008)” at <http://e-docs.bea.com/wls/docs60/security/concepts.html#concepts008>.

The BASIC and FORM techniques call into a security *realm* that contains user and password information. You can use a default realm provided with WebLogic Server, or a variety of existing realms, including realms for Windows NT, UNIX, RDBMS, and user-defined realms. For more information about security realms, see “[Security Fundamentals](#)” at

<http://e-docs.bea.com/wls/docs60/security/concepts.html#concepts002>.

Authorization (Security Constraints)

You can restrict access to servlets and other resources in a Web Application by using *security constraints*. Security constraints are defined in the Web Application deployment descriptor (`web.xml`). Resources may be constrained by roles and/or by resource. Roles can be mapped to a principal in a realm. Specific resources can be constrained by matching a URL pattern to a resource in a Web Application. Finally, the use of Secure Sockets Layer (SSL) can be specified for a security constraint.

You can also perform authorization programmatically, using one of the following methods of the `HttpServletRequest` interface:

- `getRemoteUser()`
- `isUserInRole()`
- `getUserPrincipal()`

For more information see the [javax.servlet API](#) at

<http://java.sun.com/products/servlet/2.2/javadoc/index.html>.

Servlet Development Tips

Consider the following tips when writing HTTP servlets:

- Make sure you compile your servlet classes into the `WEB-INF/classes` directory of your Web Application.
- Make sure your servlet is registered in the Web Applications deployment descriptor (`web.xml`).

- When responding to a request for a servlet, WebLogic Server checks the time stamp of the servlet class file and compares it to the servlet instance in memory. If a newer version of the servlet class is found, WebLogic Server re-loads the servlet class. When a servlet is re-loaded, the `init()` method of the servlet is called.

You can set the interval (in seconds) at which WebLogic Server checks the time stamp with the `Servlet Reload` attribute. This attribute is set on the `Files` tab of your Web Application, in the Administration Console. If you set this attribute to zero, WebLogic Server checks the time stamp on every request, which can be useful while developing and testing servlets but is needlessly time consuming in a production environment. If this attribute is set to `-1`, WebLogic Server does not check for modified servlets.

Clustering Servlets

To deploy an HTTP servlet in a cluster of WebLogic Servers, deploy the Web Application containing the servlet on all of the servers in the cluster. Depending on the cluster configuration, requests are routed to various servers in the cluster, each of which has access to a copy of the Web Application resources.

For more information on clustering, see [Using WebLogic Server Clusters](http://e-docs.bea.com/wls/docs60/cluster/index.html) at <http://e-docs.bea.com/wls/docs60/cluster/index.html>.

Index

A

- addCookie() 3-19
- administration
 - console 4-3
- administration console 4-3
- API 1-4
- authentication 4-6

C

- classpath 2-2
- clustering 3-17, 4-8
- compiling 2-2
- connection pools 3-22
 - DataSource 3-24
 - driver 3-23
 - JDBC 3-22
 - using 3-23

- contentType 2-2

- cookies 3-18

- and EJB 3-19
- and logging in 3-21
- and passwords 3-21
- domain 3-20
- HTTP and HTTPS 3-20
- retrieving 3-19
- using in servlets 3-18

- customer support contact information
- vii

D

- databases 3-22
- DataSource 3-22, 3-24
- deployment 2-3
- deployment descriptor 4-2
- Developing 1-3
- development
 - classpath 2-2
 - compiling 4-7
 - tips 4-7
- development environment 2-2
- dispatching 3-26
- documentation, where to find it vi

E

- EJB 3-21
- encodeURL() 3-15
- environment, development
 - environment 2-2

F

- forward() 3-26
- forwarding 3-27, 3-28

G

- getAttribute() 3-13
- getAttributeNames() 3-13
- getCookies() 3-19
- getParameters() 3-6
- getParameterValues() 3-9

getSession() 3-11, 3-13
H
HelloWorldServlet 2-5
HTTP
 response 3-4
HttpServletRequest 2-2
 methods 3-7
HttpServletResponse 2-2, 3-4
HttpSession object 3-11
I
IDLength 3-16
IllegalStateException 3-14
import 2-1
include() 3-26
including 3-27
including a request 3-28
init parameters 3-2
init() method 3-1, 3-3
initialization
 init() method 3-1
 on server start up 3-2
 parameters 3-2
init-param 3-3
in-memory replication 3-17
input
 query paramters 3-8
J
J2EE 1-3
javax.servlet 1-4
JDBC 3-21, 3-22, 3-25
JDBC session persistence 3-17
JMS 3-21
JNDI 3-21
JTS pool driver 3-22
K
keep alive 3-5
L
load-on-startup 3-2
logging out 3-14
N
name/value pairs 3-13
P
packages 2-1
Pool driver 3-22
printing product documentation vi
PrintWriter object 2-2
Q
query parameters 3-6, 3-7, 3-8
R
removeAttribute() 3-13
RequestDispatcher() 3-27
requests
 dispatching 3-26
 forwarding 3-27, 3-28
 including 3-27, 3-28
response 3-4
 buffer 3-5
 optimizing 3-5
retrieving input 3-6
S
security 4-6
 applying programatically 4-7
 authentication 4-6
 authorization 4-7
 constraints 4-7
 realms 4-7
security constraints 4-7
service method 2-2
Servlet 2.2 Specification 1-4

servlets
 and clustering 4-8

session persistence
 JDBC 3-17

sessions
 and clusters 3-17
 and persistence 3-16
 and serialization 3-17
 cookies 3-12, 3-15
 detecting start of 3-13
 encodeURL() method 3-15
 ending 3-14
 history of tracking 3-10
 lifetime 3-12
 logging out 3-14
 name/value attributes 3-13
 tracking 3-10, 3-12
 tracking with HttpSession object
 3-11
 tracking, configuration 3-15
 URL rewriting 3-15

setAttribute() 3-13

SingleThreadModel 3-25

SingleThreadModelPoolSize 3-25

support
 technical vii

T

threading 3-25
 SingleThreadModel 3-25

U

URL rewriting 3-15
 and WAP 3-16
 and Wireless Access Protocol
 3-16

URLs 4-5

W

WAP 3-16

Web Applications
 and security 4-6
 deployment descriptor 4-2
 directory structure 4-5
 URLs 4-5

web.xml 4-2

weblogic.xml 4-2

Wireless Access Protocol 3-16