

# CSC309 Tutorial – Servlets

**TA: Lei Jiang**

March 17, 2003




# What Are Servlets?

---

- n A servlet can be thought of as a server-side applet.
  - n Is a Java class which is loaded and executed by a web server.
  - n Like applets are loaded and executed by a web browser.
- n A servlet accepts requests from a client (via the web server), performs some task, and returns the results.

# What Are Servlets?



1. The client (most likely a web browser) makes a request via HTTP.
2. The web server receives the request and forwards it to the servlet.
  - n If the servlet has not yet been loaded, the web server loads it into the Java Virtual Machine (JVM) and executes it.
3. The servlet receives the HTTP request and performs some type of process.
4. The servlet returns a response to the web server.
5. The web server forwards the response to the client.



# Servlets v.s. CGI

---

## n Performance

### n CGI:

- n Need to create new processes
- n Web servers spend as much time on operating system overhead as they did serving pages.
- n very inefficient

### n Servlets:

- n Run in the same process and address space as its host Web server.
- n Little inter-process overhead
- n Much more efficient



# Servlets v.s. CGI

---

## n Persistence

### n CGI:

#### n Transient:

- n each time a request is made to a CGI script, it must be loaded and executed by the web server.
- n All program initialization (such as connecting to a database) must be repeated each time a CGI script is used.

### n Servlets:

#### n Persistent:

- n loaded only once by the web server and can maintain services (such as a database connection) between requests.



# Tomcat architecture

---

1. A web browser makes a request to the server.  
e.g. [http://some.server.com/servlet/my\\_servlet?param=value](http://some.server.com/servlet/my_servlet?param=value)
2. Tomcat receives this request,
  - n through its built-in web server on port 8080 (standalone), or
  - n from a *web server connector*, a piece of native code attached to the web server forwarding requests to the servlet engine (embedded)
3. If a static resource, such as a file, is requested, the **DefaultServlet** will open the file, read it, and send it back to the client.
4. If a servlet is requested, the **InvokerServlet** will execute the servlet.



# Tomcat architecture

---

- n When Tomcat starts, it reads a web application descriptor file, web.xml, which sets the default values for all web applications

A sample  
descriptor file:

```
<!-- The mapping for the default servlet -->
<servlet-mapping>
  <servlet-name>default</servlet-name>
  <url-pattern>/</url-pattern>
</servlet-mapping>

<!-- The mapping for the invoker servlet -->
<servlet-mapping>
  <servlet-name>invoker</servlet-name>
  <url-pattern>/servlet/*</url-pattern>
</servlet-mapping>
```



# Tomcat architecture

```
<!-- The mapping for the default servlet -->
<servlet-mapping>
  <servlet-name>default</servlet-name>
  <url-pattern>/</url-pattern>
</servlet-mapping>

<!-- The mapping for the invoker servlet -->
<servlet-mapping>
  <servlet-name>invoker</servlet-name>
  <url-pattern>/servlet/*</url-pattern>
</servlet-mapping>
```

- n The servlet-name elements refer to the class name of the servlet
- n Tomcat attempts to match the URL:  
    “http://some.server.com/servlet/my\_servlet?param=value”  
with the list of mappings (longest mapping first):  
    /servlet/\*  
    /  
  
n The /servlet/\* pattern matches, and the servlet **InvokerServlet** is invoked
  - n It ensures that the requested servlet **my\_servlet** is loaded and then pass control off.



# Tomcat architecture

---

n When **InvokerServlet** gets control, what happens next?  
The basic servlet lifecycle is as follows:

1. The servlet is loaded.
  - n The servlet is loaded only once;
  - n multiple threads of the same servlet will handle multiple client requests.
2. The servlet is initialized.
  - n The **init()** method on the servlet is called to allow for the servlet to perform some type of initialization
3. The servlet **service()** method is invoked.
  - n The servlet performs some type of processing and returns a response.
4. If the servlet engine is terminated, or the servlet engine decides to unload the servlet for any reason, the **destroy()** method is invoked.



# Tomcat architecture

---

- n The Tomcat architecture has made it very easy:
  - n To focus on writing the servlet pieces needed to perform work
- n you don't need to be concerned about:
  - n loading and unloading a servlet,
  - n handling the HTTP protocol,
  - n performing chaining, invoking filters, and so on.



# What Makes a Servlet a Servlet?

---

- n To create a servlet, write a program that implements the **javax.servlet.Servlet** interface, which includes following methods
  - n **init()** Called by the servlet engine (servlet container) after the servlet is loaded.
    - n The `init()` method is passed a configuration object, which the servlet can use to read initialization parameters.
  - n **service()** Allows the servlet to process a request.
  - n **destroy()** Called when the servlet is being removed from the servlet engine.
  - n **getServletConfig()** Returns the `javax.servlet.ServletConfig` object,
    - n which holds initialization and startup information for the servlet.
  - n **getServletInfo()** Returns information about the servlet,
    - n such as the name, version, and author.
- n Every servlet must implement these methods.
  - n You could implement each of these methods every time you write a servlet, or,
  - n You can extend two abstract base classes instead, saving your time:
    - n **javax.servlet.GenericServlet**: extend the class when writing a protocol-independent servlet.
    - n **javax.servlet.http.HttpServlet** extend this class when creating a servlet responding to HTTP requests.



# Servlet Example – Basic Steps

---

- n Writing a servlet for HTTP request generally requires only two basic steps:
  1. Create a new servlet class extending **javax.servlet.http.HttpServlet**
  2. Override one or both of the **doGet()** and **doPost()** methods.
    - n **HttpServlet** calls these methods from the **service()** method depending upon the type of HTTP request.
- n A servlet can also optionally override the **init()** and **destroy()** methods.
  - n initialize a database connection in the **init()** method and close the connection in the **destroy()** method
- n An example: **Properties** servlet
  - n simply return an HTML page to the client
  - n The response includes
    - n information about the client,
    - n any parameters that were passed,
    - n and all the system properties of the server.

# The Properties servlet

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

/**
 * First sample servlet. Returns any parameters and
 * lists server properties.
 */
public class Properties
    extends HttpServlet
{
    public void init(ServletConfig config)
        throws ServletException
    {
        // Initialize the servlet here.
        super.init(config);
    }

    public void doGet(HttpServletRequest req,
        HttpServletResponse resp)
        throws ServletException, IOException
    {
        // Get an output stream
        PrintWriter out = resp.getWriter();
        // Print the HTML header to the output stream
        out.println("<html>");
```

```
        out.println("<head>");
        out.println("<title>My First Servlet</title>");
        out.println("</head>");
        out.println("<h2><center>");
        out.println("Information About You</center></h2>");
        out.println("<br>");

        // Create a table with information about the client
        out.println("<center><table border>");
        out.println("<tr>");
        out.println("<td>Method</td>");
        out.println("<td>" + req.getMethod() + "</td>");
        out.println("</tr>");

        out.println("<tr>");
        out.println("<td>User</td>");
        out.println("<td>" + req.getRemoteUser() + "</td>");
        out.println("</tr>");

        out.println("<tr>");
        out.println("<td>Client</td>");
        out.println("<td>" + req.getRemoteHost() + "</td>");
        out.println("</tr>");

        out.println("<tr>");
        out.println("<td>Protocol</td>");
        out.println("<td>" + req.getProtocol() + "</td>");
        out.println("</tr>");
```

# The Properties servlet

```
// Get and dump all of the request parameters
java.util.Enumeration enum =
    req.getParameterNames();
while (enum.hasMoreElements()) {
    String name = (String) enum.nextElement();
    out.println("<tr>");
    out.println("<td>Parameter '" + name +
        "'</td>");
    out.println("<td>" + req.getParameter(name) +
        "</td>");
    out.println("</tr>");
}

out.println("</table></center><br><hr><br>");

// Create a table with information about the server
out.println("<h2><center>");
out.println("Server Properties</center></h2>");
out.println("<br>");
out.println("<center><table border width=80%>");
java.util.Properties props = System.getProperties();
enum = props.propertyNames();
```

```
while (enum.hasMoreElements()) {
    String name = (String) enum.nextElement();
    out.println("<tr>");
    out.println("<td>" + name + "</td>");
    out.println("<td>" + props.getProperty(name) +
        "</td>");
    out.println("</tr>");
}
out.println("</table></center>");

// Wrap up
out.println("</html>");
out.flush();
}

public void destroy()
{
    // Clean up here
}
}
```



# Servlet Example – Packages

---

- n When writing servlets, almost always need to import:
  - n **javax.servlet.\***
    - n Contains the *generic* servlet interfaces and classes,
    - n Contains special servlet exceptions.
  - n **javax.servlet.http.\***
    - n contains the *special* interfaces and classes used when writing servlets responding to HTTP requests.



## Servlet Example – `init()`

---

- n The first method in the program is an override of the **`init()`** method
  - n It override its superclass's `init()` method by calling **`super.init(config)`**
  - n Can use the **`ServletConfig`** object to get initialization parameters
- n required in every HTTP servlet.



# Servlet Example – doXXX

---

- n doGet is invoked by the Web server each time a client browser requests via an HTTP GET request.
- n There are other doXXX methods corresponding to other HTTP requests
  - n doDelete, doOptions, doPost and doPut
- n doXXX methods take two arguments:
  - n **ServletRequest** and **ServletResponse**
- n doXXX methods must throw two exceptions:
  - n **ServletException** and **IOException**
- n **A typical doGet method signature:**

```
public void doPost(HttpServletRequest request, HttpServletResponse response)
    throws ServletException, IOException
```



# Servlet Example – ServletRequest

---

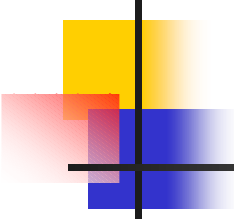
- n **ServletRequest** class contains all the information the servlet needs regarding the current request.
- n Its methods include:
  - n `getLength()`
    - n Returns the length, in bytes, of the request body
    - n For HTTP servlets, this is the same as the value of the CGI variable `CONTENT_LENGTH`.
  - n `getContentType()`
    - n returns the MIME type of the body of the request
    - n for HTTP servlets, this is the same as the value of the CGI variable `CONTENT_TYPE`.
  - n `getParameter(String)`
    - n returns the value of a request parameter
    - n use it when the parameter has only one value



# Servlet Example – ServletRequest

---

- n `getParameterNames()`
  - n returns an Enumeration the names of the parameters
- n `getParameterValues(String)`
  - n returns an array of the values for the given parameter
- n `getRemoteHost()`
  - n returns the fully qualified name of the client that sent the request
- n `getServerPort()`
  - n Returns the port number on which this request was received
- n `getServerName()`
  - n Returns the hostname of the server that received the request
- n `getMethod()`
  - n Returns the name of the HTTP method



# Servlet Example – ServletResponse

---

- n **ServletResponse** class allows the servlet to return its results to the Web server.
- n Its methods include:
  - n **setContentType(String)**
    - n Sets the content type of the response being sent.
    - n Allowable types include text/html, text/plain, and other MIME types.
  - n **getWriter()**
    - n Returns a **PrintWriter** object
    - n Suitable for writing HTML pages or plain text in the response
  - n **getOutputStream()**
    - n Returns a **ServletOutputStream** object
    - n Suitable for writing binary data (e.g. images) in the response
    - n Can be used to return a part of a page at a time



# Web Applications – Overview

---

- n According to the Servlet specification:
  - n A web application is a collection of servlets, HTML pages, classes, and other resources
    - n That make up a complete application on a web server.
  - n The web application can be bundled and run on multiple containers from multiple vendors
- n The notion of a web application addresses issues like:
  - n Where to place servlet classes
  - n How to supply initialization parameters for servlets
  - n Where to place other resources, such as HTML files and images
  - n How to configure the servlet container
  - n How to install supporting libraries (JAR files)



# Web Applications – Directory Structure

- n A web application contains a predefined directory structure
  - n that defines where to put classes, libraries, and general resources.

/	<ul style="list-style-type: none"><li>• The web application root directory.</li><li>• General resources, such as HTML files, can be placed here.</li><li>• You are encouraged to create subdirectories for better organization.</li></ul>
/WEB-INF	<ul style="list-style-type: none"><li>• Contains the web application descriptor file named <b>web.xml</b>.</li><li>• This directory is not visible outside of the web application.</li><li>• In other words, a client cannot request to view the contents of the /WEB-INF directory (or any of its subdirectories).</li></ul>
/WEB-INF/classes	<ul style="list-style-type: none"><li>• The location of any classes used by the web application.</li><li>• This is the root directory of any class packages.</li></ul>
/WEB-INF/lib	<ul style="list-style-type: none"><li>• The location of any supporting libraries (JAR files).</li><li>• All JAR files found in this location will automatically be placed on the classpath.</li><li>• Servlets can be packaged in JAR files and placed here as well.</li></ul>

# Web Applications – Directory Structure

A sample directory structure

/webapps/sample/index.html

/webapps/sample/images/banner.gif

/webapps/sample/images/logo.gif

**/webapps/sample/WEB-INF/web.xml** ←

**/webapps/sample/WEB-INF/classes/Properties.class**

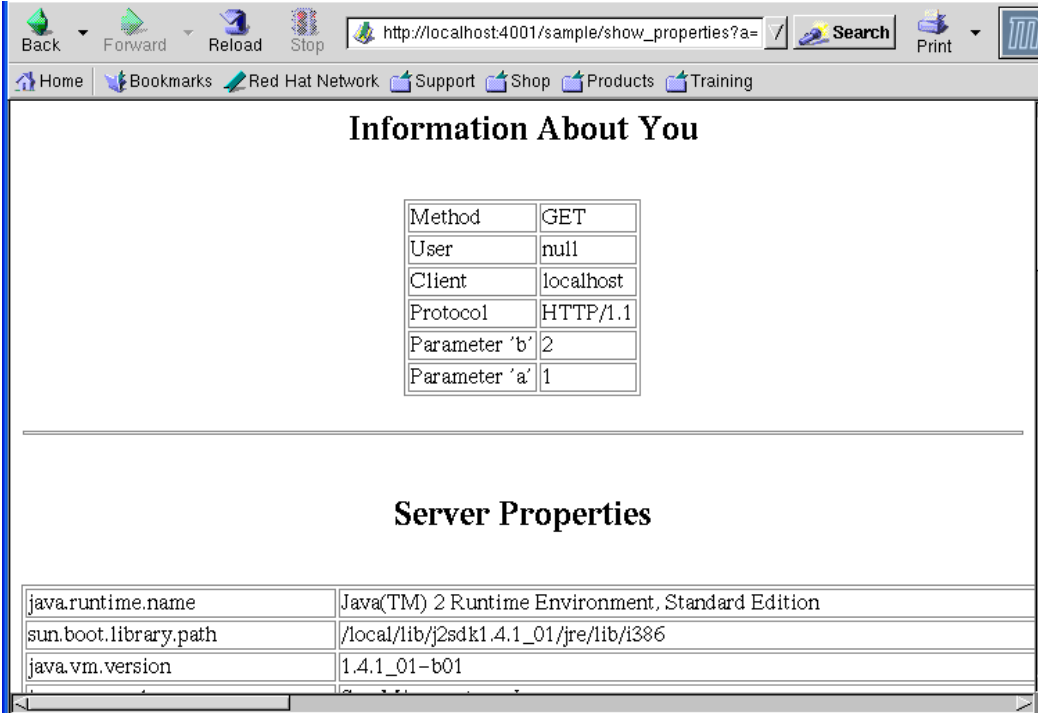
/webapps/sample/WEB-INF/lib/jaxp.jar

/webapps/sample/WEB-INF/lib/xalan.jar

```
<!DOCTYPE web-app
  PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.2//EN"
  "http://java.sun.com/j2ee/dtds/web-app_2_2.dtd">
<web-app>
  <!-- this is where we specify our servlet's java class -->
  <servlet>
    <servlet-name>Prop309</servlet-name>
    <servlet-class>Properties</servlet-class>
  </servlet>
  <!-- this is where we specify the path to access our servlet -->
  <servlet-mapping>
    <servlet-name>Prop309</servlet-name>
    <url-pattern>/prop</url-pattern>
  </servlet-mapping>
</web-app>
```

# Web Applications – Directory Structure

- n Given the above directory structure (under the Tomcat home), we have in a web application named “sample”
  - n It receives requests that begin with **/sample**
- n `http://localhost:4001/sample/show_properties?a=1&b=2`



The screenshot shows a web browser window with the address bar containing `http://localhost:4001/sample/show_properties?a=`. The browser's navigation bar includes Back, Forward, Reload, and Stop buttons. The page content is as follows:

### Information About You

Method	GET
User	null
Client	localhost
Protocol	HTTP/1.1
Parameter 'b'	2
Parameter 'a'	1

---

### Server Properties

java.runtime.name	Java(TM) 2 Runtime Environment, Standard Edition
sun.boot.library.path	/local/lib/j2sdk1.4.1_01/jre/lib/i386
java.vm.version	1.4.1_01-b01



# Web Applications – WAR Files

---

## n **Web Application Archive**

- n Bundle all the files together in a portable collection.
- n Is simply a JAR file that has .war extension.
- n Use the jar tool to create a WAR file and place it into a running servlet container.
  - n `jar -cf sample.war *`
  - n invoke the jar command from the root directory of the web application.
- n The servlet container recognizes the WAR file and automatically create a new mapping to handle incoming requests.



# Most Common Error

---

- n When a servlet is loaded, only one instance will be created.
  - n Each thread will handle a single request
  - n This means multiple threads may be calling a method of the same instance at the same time
- n Most common error: declare a class field and use it within `service()`, `doGet()` or `doPost()` methods.
  - n Without synchronizing access, it will cause problems
- n In general, always avoid declaring class fields within servlets
  - n unless they are to be used as read-only constants.