

Project 3
Wahoo!
124 Points

due by 12:30 P.M. ET on Tuesday, 6 April 2006

It is recommended that you read the entirety of this document,
well in advance of this project's deadline, before answering any of its questions.

It is further recommend that you begin answering this document's questions immediately thereafter.
Really.

Goals.

The goals of this project are to:

- Challenge you to implement your own XML-based portal.
- Give you hands-on experience with JavaServer Pages 2.0, Java Servlet 2.4, and TrAX, as well as additional experience with XSLT.
- Introduce you to an industry-standard application server, Apache's Tomcat 5.5.16.

WAAHOO!

Grading Metric.

Each question is worth the number of points specified parenthetically in line with it.

Your responses to questions requiring exposition will be graded on the basis of their clarity and correctness. Your responses to questions requiring code will be graded on the following bases.¹

Basis	Considerations
Correctness	Does your code work in accordance with the question's guidelines?
Design	Does your code make sense, given the question's framework? Is your code written logically, clearly, and succinctly? Is your code efficient? Is your code divided into logical units (<i>e.g.</i> , multiple methods and templates)?
Style	Is your code rigorously documented with inline comments and Javadoc doc comments and tags? ² Is it clear from your comments alone how your code operates? Is your code pretty-printed? Are your attributes, data members, elements, methods, templates, and variables aptly named?

Academic Honesty.

For this project, you will be running your own instance of a webserver on `nice.fas.harvard.edu` that does not require that any files or directories being served be world-executable or world-readable.

If, however, you opt to publish world-readable files on `www.people.fas.harvard.edu` for some reason, it is expected that you will ensure the privacy of your work by password-protecting any Web-accessible directories and naming files in such a (random) way as to render their viewing on `nice.fas.harvard.edu` unlikely.

Recall that, for Project 1, you configured your account with a `~/public_html/` directory and a password-protected `~/public_html/cscie259/` directory.

Needless to say, attempting to view the work of another student, even if published in a world-accessible directory, is considered academic dishonesty and will be handled accordingly.

¹ By code, we mean Java, JSP, XHTML, XML, XPath, and XSLT.

² With regard to Javadoc, we expect descriptions for all methods and appropriate use of the `@param`, `@return`, and `@throws` block tags.

Getting Started.

1. (1 point.) Okay, okay, fine, here's a point.
2. (0 points.) If you intend to do your work on `nice.fas.harvard.edu`, SSH to that machine and execute the following sequence of commands.^{3,4}

```
cp -r ~cscie259/pub/distribution/projects/project3-6.0/ ~/cscie259/  
cd ~/cscie259/  
ls
```

You should see that you have the following in your current working directory.

```
project1-6.0/      project2-6.0/      project3-6.0/
```

If, on the other hand, you do not intend to do your work on `nice.fas.harvard.edu`, you may proceed to download a gzip-compressed tarball or a ZIP file containing this `project3-6.0/` directory from the course's website to your local machine. Or, of course, can you transfer the directory itself via SFTP to your local machine. However, prior to disconnecting from `nice.fas.harvard.edu`, be sure to examine `cscie259.cshrc` in `~cscie259/pub/etc/`, so that you know how to configure your machine; note that we expect use of J2SE SDK 1.4.2, Apache's Ant 1.6.5, Apache's Tomcat 5.5.16 (and its compatibility package), Apache's Xerces-J 2.8.0, and Apache's Xalan-J 2.7.0.^{5,6}

³ These commands assume that you already created (for Projects 1 and 2) a directory called `cscie259/` in your FAS account's home directory.

⁴ Beware the distinction between `~cscie259` and `~/cscie259`. Also, if you do your work on a Windows machine, take care to substitute backslashes for any forward slashes as well as `dir` for any instances of `ls` in this document's directions.

⁵ Specifically, if your local machine runs Linux or UNIX: install each of these distributions to some directory on your hard drive; define an environment variable called `"JAVA_HOME"` that points to your installation of J2SE SDK 1.4.2; create (in accordance with Sun's "Endorsed Standards Override Mechanism") in the SDK's `jre/lib/` directory a directory called `"endorsed"`, thereafter copying into it `serializer.jar` and `xalan.jar` from your installation of Xalan-J 2.7.0 as well as `xercesImpl.jar` and `xml-apis.jar` from your installation of Xerces-J 2.8.0; also copy `serializer.jar`, `xalan.jar`, `xercesImpl.jar`, and `xml-apis.jar` into Tomcat 5.5.15's `common/endorsed/` directory after deleting `xercesImpl.jar` and `xmlParserAPIs.jar` from the same; copy `ant.jar` and `ant-launcher.jar` from Ant 1.6.5's `lib/` directory into Tomcat 5.5.15's `common/lib/` directory; define an environment variable called `"ANT_HOME"` that points to your installation of Ant 1.6.5; define an environment variable called `"CATALINA_HOME"` that points to your installation of Tomcat 5.5.15; define an environment variable called `"JAVACMD"` whose value is `"$JAVA_HOME/bin/java"`; define an environment variable called `CATALINA_BASE` whose value is `"."`; add `./build` to your `CLASSPATH`; add `"$ANT_HOME/lib/ant.jar"` and `"$ANT_HOME/lib/ant-launcher.jar"` to your `CLASSPATH`; add to your `CLASSPATH` both `"$CATALINA_HOME/bin/bootstrap.jar"` and `"$CATALINA_HOME/common/lib/servlet-api.jar"`; and prepend `"$JAVA_HOME/bin"` and `"$ANT_HOME/bin"` to your `PATH`.

⁶ Specifically, if your local machine runs some version of Windows: install each of these distributions in the root of your hard drive (or, at least, not within any directory whose name contains spaces); define an environment variable called `"JAVA_HOME"` that points to your installation of J2SE SDK 1.4.2; create (in accordance with Sun's "Endorsed Standards Override Mechanism") in the SDK's `jre\lib\` directory a directory called `"endorsed"`, thereafter copying into it `serializer.jar` and `xalan.jar` from your installation of Xalan-J 2.7.0 as well as `xercesImpl.jar` and `xml-apis.jar` from your installation of Xerces-J 2.8.0; also copy `serializer.jar`, `xalan.jar`, `xercesImpl.jar`, and `xml-apis.jar` into Tomcat 5.5.15's `common\endorsed\` directory after deleting `xercesImpl.jar` and `xmlParserAPIs.jar` from the same; copy `ant.jar` and `ant-launcher.jar` from Ant 1.6.5's `lib\` directory into Tomcat 5.5.15's `common\lib\` directory; define an environment variable called `"ANT_HOME"` that points to your installation of Ant 1.6.5; define an environment variable called `"CATALINA_HOME"` that points to your installation of Tomcat 5.5.15; define an environment

Notice, now, that your `project3-6.0/` directory is structured as follows.

```
project3-6.0/  
  conf/  
  src/  
    cscie259/  
      project3/  
        wahoo/  
temp/  
webapps/  
  ROOT/  
    docs/  
      cscie259/  
        project3/  
          wahoo/  
    dtd/  
    images/  
    WEB-INF/  
    xml/  
      cache/  
    xsl/
```

Ensure that your account or machine is in order by executing the following command from within any directory.⁷

```
EnvironmentCheck
```

Confirm that your account is configured to use Ant 1.6.5, Xerces-J 2.8.0, and Xalan-J 2.7.0.

Next, execute the following command from within your `project3-6.0/` directory.⁸

```
ant compile
```

By default, ant will compile `cscie259.project3.wahoo.*`, storing the resulting bytecodes in `project3-6.0/webapps/ROOT/WEB-INF/classes/`.

In fact, take a look at `project3-6.0/build.xml` so that you know what ant can do for you during this project.

variable called `"CATALINA_BASE"` whose value is `"."`; add `".\build"` to your `CLASSPATH`; add `"%ANT_HOME%\lib\ant.jar"` and `"%ANT_HOME%\lib\ant-launcher.jar"` to your `CLASSPATH`; add to your `CLASSPATH` both `"%CATALINA_HOME%\bin\bootstrap.jar"` and `"%CATALINA_HOME%\common\lib\servlet-api.jar"`; and prepend `"%JAVA_HOME%\bin"` and `"%ANT_HOME%\bin"` to your `PATH`.

⁷ Note that, on `nice.fas.harvard.edu`, ``EnvironmentCheck`` is an alias for ``java org.apache.xalan.xslt.EnvironmentCheck``.

⁸ Alternatively, you can execute ant without any arguments.

Next, proceed to edit `project3-6.0/conf/server.xml` with your favorite text editor. Per the file's comments, locate the `Server` element's `port` attribute and assign it an integral value between 1024 and 65535, inclusive; then, locate the `Connector` element's `port` attribute and assign it an integral value between 1024 and 65535, inclusive, as well, taking care not to reuse the value you chose for the `Server` element's `port` attribute. After saving and closing the file, execute the following command from within `project3-6.0/`.⁹

```
tomcat
```

Then, assuming you've connected via SSH to `nice.fas.harvard.edu`, take notice of the particular host to which you're connected by examining your prompt (which should be of the form `"$USER@$ICEBOX (%~):"`, per Project 1's configuration of your account) or by executing the following command.¹⁰

```
echo $ICEBOX
```

You should see output resembling the below, where `$ICEBOX` is the host to which you're connected, `m` is your choice of ports for the `Server` element's `port` attribute, and `n` is your choice of ports for the `Connector` element's `port` attribute.

```
Using CATALINA_BASE:  .
Using CATALINA_HOME:  /home/c/s/cscie259/pub/local/jvm/apache-tomcat-5.5.16
Using CATALINA_TMPDIR: ./temp
Using JRE_HOME:       /home/c/s/cscie259/pub/local/i386/j2sdk1.4.2_10
Created MBeanServer with ID: 161d36b:109ce4e5ce3:-m:$ICEBOX.fas.harvard.edu:1
Mar 6, 2006 1:46:16 AM org.apache.coyote.http11.Http11BaseProtocol init
INFO: Initializing Coyote HTTP/1.1 on http-n
Mar 6, 2006 1:46:16 AM org.apache.catalina.startup.Catalina load
INFO: Initialization processed in 1149 ms
Mar 6, 2006 1:46:16 AM org.apache.catalina.core.StandardService start
INFO: Starting service Catalina
Mar 6, 2006 1:46:17 AM org.apache.catalina.core.StandardEngine start
INFO: Starting Servlet Engine: Apache Tomcat/5.5.16
Mar 6, 2006 1:46:17 AM org.apache.catalina.core.StandardHost start
INFO: XML validation enabled
Mar 6, 2006 1:46:18 AM org.apache.coyote.http11.Http11BaseProtocol start
INFO: Starting Coyote HTTP/1.1 on http-n
Mar 6, 2006 1:46:18 AM org.apache.catalina.startup.Catalina start
INFO: Server startup in 1830 ms
```

At this point, an instance of Tomcat is running under your `userid` on port `n` of the host to which you're connected. (Tomcat should now have control of your terminal; you should not

⁹ Note that, on `nice.fas.harvard.edu`, ``tomcat`` is an alias for ``catalina.sh run``. If you opt to develop on a Windows machine, you should instead execute ``catalina.bat run`` (assuming `"%CATALINA_HOME%\bin\"` is in your `PATH`); alternatively, you can achieve this same aliasing effect by creating in any directory in your `PATH` a file called `tomcat.bat` containing the following two lines.

```
@echo off
%CATALINA_HOME%\bin\catalina.bat run
```

¹⁰ Though we've always referred to it as one system, `nice.fas.harvard.edu` actually refers to a cluster of machines, all of which mount via NFS your home directory; you are connected to one essentially at random.

be returned to your shell's prompt.) Confirm as much by visiting the URL below, taking care to substitute `$ICEBOX` and `n` with their appropriate values.^{11,12}

`http://$ICEBOX.fas.harvard.edu:n/happyenv.jsp`

Confirm that Tomcat, like your account, is configured to use Ant 1.6.5, Xerces-J 2.8.0, and Xalan-J 2.7.0.

Next, visit the URL below, again taking care to substitute `$ICEBOX` and `n` with their appropriate values.

`http://$ICEBOX.fas.harvard.edu:n/`

You should be automatically directed to `/servlet/login`, which should display Wahoo's login page. Meanwhile, output resembling the below should have been outputted to Tomcat's console (*i.e.*, your terminal).

```
Adding "jharvard" to main memory with password "crimson"... Done.
Associating "Boston news" with user "jharvard"... Done.
Associating "Top stories" with user "jharvard"... Done.
Associating "Top technology stories" with user "jharvard"... Done.
```

Go ahead and log in with username "jharvard" and password "crimson". Though, seeing as you haven't yet implemented your portal, the results won't be all that interesting.

Proceed to terminate Tomcat by hitting `Ctrl-C` at its console (*i.e.*, your terminal).

If your account or system does not behave as described above, simply contact `cscie259@lists.dce.harvard.edu` or the staff for assistance.

My Wahoo!

3. (0 points.) "Moreover Technologies is a provider of real-time information management solutions, delivering essential online information in time to impact business decisions.

"Unlike traditional news services that resell archived information, Moreover's sophisticated technology continually scours the Internet to capture breaking news and business information from more than 4,500 qualified, handpicked sources. Headline links to relevant information are filtered according to clients' specific business needs, and delivered in real time to any platform or business application through one of Moreover's customizable Connected Intelligence™ solutions.

¹¹ If you opt to develop on your own machine, you should instead visit `http://localhost:n/` or `http://127.0.0.1:n/`.

¹² Although you needn't type "`http://`" to visit most websites with most browsers these days, be sure to do so in this case.

“Corporations use Moreover’s technology solutions to stay competitive and increase their margins. Today’s quickly evolving business environment requires that companies have the right information, all the time, to spot market opportunities, track industry developments, and react quickly to announcements. Ultimately, Moreover helps every employee make better and faster decisions to stay ahead of the competition.

“Moreover Technologies has operations in San Francisco and London, and is backed by prominent venture firms including Reuters Venture Capital and Atlas Ventures.”¹³

How’s that for marketing hype?

In other words, Moreover is a syndicator of news. The firm aggregates headlines from around the world and produces “over 330 feeds that are free to non-for-profit webmasters,” who can then incorporate those headlines and the associated articles into their own sites.¹⁴

Sound too good to be true? Take a look for yourself at Moreover’s offerings by visiting the URL below!

http://w.moreover.com/categories/category_list.html

Clearly, Moreover has incorporated their headlines into a portal of their own. But it’s no Wahoo! That’s where you come in. Read on.

4. (0 points.) Not only does Moreover offer its newsfeeds as HTML, it also offers them as XML. Needless to say, we’ll be using XML, Moreover’s “preferred format [for] developers’ use for integration.”¹⁵

Before you can access a newsfeed, however, you must know its name. Fortunately, Moreover provides an XML-based list of “channels” and categories into which it has placed its many newsfeeds, the DTD for which appears below.¹⁶

```
<!ELEMENT nested_category_list (channel*)>
<!ELEMENT channel (channel_name, category*)>
<!ELEMENT channel_name (#PCDATA)>
<!ELEMENT category (category_name, feed_name)>
<!ELEMENT category_name (#PCDATA)>
<!ELEMENT feed_name (#PCDATA)>
```

The list of channels and their categories itself is available (as, of course, XML) at the URL below.¹⁷

¹³ Excerpted from http://w.moreover.com/main_site/aboutus/company.html.

¹⁴ Quotation excerpted from <http://w.moreover.com/dev/xml/> prior to the page’s removal.

¹⁵ Quotation excerpted from <http://w.moreover.com/dev/> prior to the page’s alteration.

¹⁶ This DTD is also available at http://w.moreover.com/xml/xml_nestedcatlist.dtd as well as in `project3-6.0/webapps/ROOT/xml/cache/`.

¹⁷ A sample list, generated automatically at 3:13 P.M. ET on Friday, 3 March 2006, is available in `project3-6.0/webapps/ROOT/xml/cache/nested_category_list.xml`.

http://w.moreover.com/categories/nested_category_list.xml

Note that each newsfeed mentioned in the list of channels and categories can be accessed by way of its category_name. For instance, to access the “Biotech news” feed, you can visit

<http://p.moreover.com/cgi-local/page?c=Biotech%20news&o=xml>

where the HTTP GET string’s parameters specify the category’s name (c) and the newsfeed’s output format (o), respectively.¹⁸

Each feed is then structured according to the DTD below.¹⁹

```
<!ELEMENT moreovernews (article*)>
<!ELEMENT article (url, headline_text, source, media_type, cluster,
tagline, document_url, harvest_time, access_registration, access_status)>
<!ATTLIST article id ID #IMPLIED>
<!ELEMENT url (#PCDATA)>
<!ELEMENT headline_text (#PCDATA)>
<!ELEMENT source (#PCDATA)>
<!ELEMENT media_type (#PCDATA)>
<!ELEMENT cluster (#PCDATA)>
<!ELEMENT tagline (#PCDATA)>
<!ELEMENT document_url (#PCDATA)>
<!ELEMENT harvest_time (#PCDATA)>
<!ELEMENT access_registration (#PCDATA)>
<!ELEMENT access_status (#PCDATA)>
```

For instance, below is an excerpt from the “Biotech news” feed at 1:54 A.M. ET on Monday, 6 March 2006.²⁰

```
<moreovernews>
[... ]
<article id=" 479411492">
  <url>http://c.moreover.com/click/here.pl?x479411492</url>
  <headline_text>Biometrics look ready for prime time </headline_text>
  <source>Government Computer News</source>
  <media_type>text</media_type>
  <cluster>moreover...</cluster>
  <tagline></tagline>
  <document_url>http://www-gcn3.iproduction.com</document_url>
  <harvest_time>Mar 4 2006 10:54PM</harvest_time>
  <access_registration></access_registration>
  <access_status></access_status>
</article>
[... ]
</moreovernews>
```

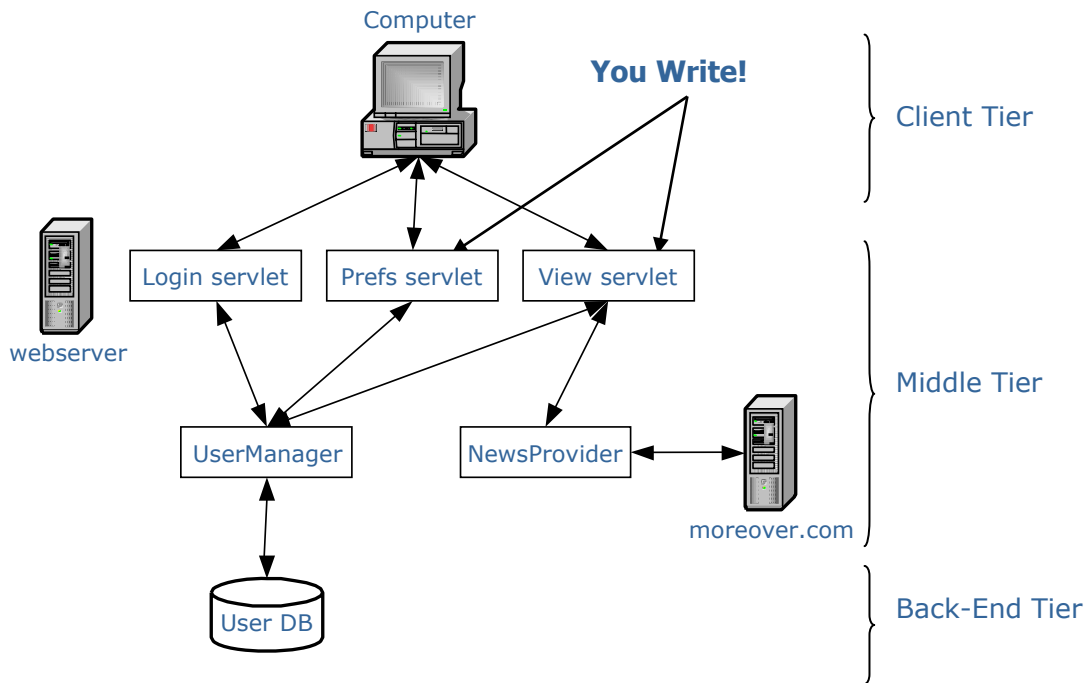
¹⁸ Note that any space in a category’s name must be encoded in the application/x-www-form-urlencoded MIME format (*i.e.*, as '+' or '%20').

¹⁹ This DTD is also available at http://p.moreover.com/xml_dtds/moreovernews.dtd.

²⁰ A copy of this excerpt’s source is available in project3-6.0/webapps/ROOT/xml/cache/Biotech%20news.xml.

Clearly, the root element of the newsfeed is `moreovernews`. Meanwhile, an `article` element represents an individual article, uniquely identified by way of an `id` attribute. Each article has a number of children: `url` specifies a proxy address for the actual article; `headline_text` specifies the article's headline; `source` specifies the article's source; `media_type` specifies the article's format; `cluster` specifies (so far as I can tell) some sort of grouping for the article; `tagline` specifies the article's tagline, if any; `document_url` specifies the homepage of the article's source; `harvest_time` specifies the moment at which the headline was discovered; and `access_registration` and `access_status` specify registration requirements for the article's source, if any.

5. (0 points.) Lest you be overwhelmed by the challenge of implementing a web-based portal, rest assured that we've provided you with a framework. Per Lecture 7, Wahoo's architecture resembles the below.



In order to better understand this architecture, spend some time looking through the code in `project3-6.0/src/cscie259/project3/wahoo/`.

Start by examining the `Login servlet` by way of `Login.java`, which handles the task of authenticating and creating users; notice how it uses `login.xml` in `project3-6.0/webapps/ROOT/xml/` to generate the login screen. Notice how HTTP GETs and HTTP POSTs are both handled by `doWork(·, ·, ·)`, which essentially drives the entire servlet. Since Wahoo's login screen is rather minimalist, `displayLoginForm(·, ·, ·)` does little more than apply (by way of TrAX) `login.xml` to a dummy XML file (*i.e.*, `project3-6.0/webapps/ROOT/xml/dummy.xml`); certainly a more interesting login screen could make use of actual XML content. Note, by the way, that the XHTML generated

by `login.xml` includes a form that will be submitted via an HTTP POST operation back to the Login servlet. Finally, notice how the mere presence of a username in the session implies that a user is currently authenticated (useful knowledge if you wish to detect whether a user's attempted to visit some servlet prior to authenticating).

Take a quick peek, by the way, at `project3-6.0/webapps/ROOT/index.jsp`, a minimalist JSP that redirects users to the Login servlet. In other words, if a user visits `/` on your instance of Tomcat, the user will be redirected automatically to `/servlet/login`.

Next, take a look at the View servlet by way of `View.java`; recall from your perusal of `Login.java` that the Login server redirects a user to this View servlet upon successful authentication. Eventually, this View servlet will generate the user's personalized portal; at present, it does very little.

Now, stroll on over to the Prefs servlet by way of `Prefs.java`. Whereas the View servlet will eventually display a user's favorite headlines, the Prefs servlet will eventually provide the user with a XHTML form-based interface with which to add newsfeeds to and remove newsfeeds from his favorites. At present, the Prefs servlet does very little.

That's it for Wahoo's servlets. To understand how they're accessed by way of `/servlet/`, take a quick look at `project3-6.0/webapps/ROOT/WEB-INF/web.xml`.

Okay, let's now examine the remaining classes in `project3-6.0/src/cscie259/project3/wahoo/`, all of which provide support for Wahoo's middle tier, but don't interact with the client tier directly.

First, take a gander at `WahooServlet.java` for `WahooServlet`, the abstract class from which the Login, Prefs, and View servlets descend. Besides the signature for each of those classes' `doWork` methods, `WahooServlet` also offers a pair of helper methods, one of which (*i.e.*, `redirect(·,·)`) the Login servlet already makes use of, the other of which (*i.e.*, `forward(·,·,·)`) you may or may not find useful during development.

Next, steal a glance at `User.java`, which implements the concept of a user. Then, read over `UserManager.java`, which handles the storage and retrieval of user data, including usernames, passwords, and favorite newsfeeds. In fact, recall the output on Tomcat's console when you tried to log in earlier: it's this `UserManager` class whose methods loaded John Harvard and his favorites into memory. Not sure how such came to pass? Well, upon the user's first visit to the Login servlet, the servlet container automatically executes that servlet's `init()` method, inherited from `WahooServlet`. (Similarly does the servlet container handle the user's first visit to the Prefs servlet.) That `init()` method, meanwhile, invokes the `UserManager` class's own `init(·)` method, which proceeds to instantiate, if it hasn't already, the class's one and only instance, the result of which is execution of the class's private, explicit constructor, which parses `project3-6.0/webapps/ROOT/xml/users.xml` and loads

John Harvard and his favorites into memory. Make sense?²¹ Take a look at `users.xml`, so that you see exactly what data is loaded into memory.

Some additional discussion of the `User` and `UserManager` classes is perhaps in order. First, it's worth noting that calling, for example,

```
UserManager.getUser("jharvard")
```

will return John Harvard's `User` object, provided John is registered in `users.xml` (which he is, by default); otherwise, this method returns `null`.

Second, calling, for instance,

```
UserManager.addUser("jharvard", "crimson")
```

will add to Wahoo's database a user with username `"jharvard"` and password `"crimson"`, thereafter returning a reference to the newly instantiated `User` object, unless username `"jharvard"` is already in the database, in which case `null` is returned.

Finally, calling

```
UserManager.save()
```

will save information about users to `users.xml`. Be sure to call this method every time you change user information (*e.g.*, add a user, add or remove a favorite newsfeed for a user, *etc.*), else the information will be lost when Tomcat terminates.

Note that each of these three methods is static; you will never manipulate a reference to `UserManager` object directly. Additionally, you should never instantiate a `User` object directly; allow the `UserManager` class's `addUser(·, ·)` method to handle such details.²²

Make sense? Okay, last but not least, look over `NewsProvider.java`, which handles the retrieval of headlines from Moreover.

If you learn better by way of Javadoc, feel free to view `project3-6.0/webapps/ROOT/docs/` with your favorite browser. Note that we've placed the project's initial Javadoc in `project3-6.0/webapps/ROOT/` rather than in `project3-6.0/` so that you can view our (and, soon, your) Javadoc with your own instance of Tomcat!²³ In fact, if Tomcat is running,

²¹ Feel free to insert temporarily, for example,

```
(new Exception()).printStackTrace()
```

anywhere you'd like to force a printing of the call stack; such should make this sequence of executions all the more clear.

²² In the interests of testing, you're welcome, of course, to add `user` elements to `users.xml` with your favorite text editor.

²³ Know that, during the development of this project, you can update your Javadoc by executing ``ant javadoc`` from within your `project3-6.0/` directory.

you can visit the URL below, again taking care to substitute `$ICEBOX` and `n` with their appropriate values.

`http://$ICEBOX.fas.harvard.edu:n/docs/`

Got it? Lots of code, to be sure. But, trust us: spend a good amount of time now getting to know the distribution code. Only if you understand the framework will building upon it be straightforward.

6. (1 point.) You're doing well, so far. Here's a point for hanging in there!
7. (40 points.) Okay, here we go! Your first challenge is to implement the essence of the portal: the `View` servlet, which, once a user has authenticated, displays, at least, the user's favorite headlines and links to the corresponding articles. The design of this page (or, rather, set of pages, all of which will be generated dynamically) is largely up to you, but we do ask that you do adhere to a few guidelines.
 - i. The servlet should allow a user to view headlines from any of his or her favorite newsfeeds. For instance, by clicking a link labeled "Boston news", John Harvard should be able to view Boston's latest headlines (along with links to the corresponding articles); that link should pass back to the servlet the user's choice of newsfeeds (*i.e.*, categories) by way of an HTTP parameter. In other words, the servlet should *not* output the headlines from all of John Harvard's favorites newsfeeds in one, large page. We suggest that you include at the bottom of any page generated by the servlet a list of the user's favorite newsfeeds, so that navigation among their headlines is simple.
 - ii. The servlet should utilize `project3-6.0/webapps/ROOT/xsl/view.xsl` in order to generate these pages of headlines dynamically, making use, of course, of the XML returned by the `NewsProvider` class's `getHeadlines()` method.
 - iii. If a user attempts to visit the servlet prior to authenticating, you should redirect the user to the `Login` servlet.
 - iv. Any page generated by the servlet should contain a link to the `Prefs` servlet, which, of course, will be implemented by you shortly!
 - v. Any page generated by the servlet should contain a link via which the user can log out (*e.g.*, the link might pass some parameter back to the servlet, instructing it to invalidate the current session and redirect the user to the `Login` servlet).

Most likely, you will want to pass to `view.xsl` not only the XML for some newsfeed but also the names of the user's other favorite categories (in fulfillment of i, above). Clearly, you can obtain the former by way of the `NewsProvider` class's `getHeadlines()` method. The latter you can obtain by way of a `User` object's `getNewsCategories()` method. However, those categories are returned as an instance of `java.util.List`, not XML. Hence, you may want your `View` servlet to iterate through this `List` of categories, wrapping each category with some

XML start tag and corresponding end tag, concatenating the elements together to form a `String`. You can then wrap those elements with a parent element (by concatenating the `String` with some other start tag and end tag on either side). You can then concatenate that result with the XML returned by `getHeadlines()` (the root of which is a `moreovernews` element), thereafter encapsulating the data in a single, root element. You can then apply `view.xml` to that dynamically created `String`.²⁴

However, since the XML is in a `String`, you cannot utilize TrAX in the same manner that the `Login` servlet does. Recall that the `Login` servlet applies `login.xml` to `dummy.xml` by way of the following code.

```
StreamSource xmlSource =  
new StreamSource(getServletContext().getRealPath("/xml/dummy.xml"));  
transformer.transform(xmlSource, new StreamResult(out));
```

To apply `view.xml` to a `String` called, say, `foo` (as opposed to a file), you'll need to employ code like the below (provided your `Transformer` object has already been configured with `view.xml`, just as `Login.java`'s `Transformer` object is configured with `login.xml`).

```
StreamSource xmlSource =  
new StreamSource(new java.io.ByteArrayInputStream(foo.getBytes()));  
transformer.transform(xmlSource, new StreamResult(out));
```

Incidentally, although we've configured Tomcat to check the modification dates and times of class files, you may want to restart Tomcat anytime you recompile your source.

Okay, get to it!

8. (40 points.) Now that your `View` servlet is working, it's time to enable the `Prefs` servlet, so that the user can customize his or her favorite categories. While the design of this interface is also largely up to you, we do ask that you do adhere to a few guidelines.
 - i. The servlet should enable the user to add newsfeeds to his or her favorites as well as remove newsfeeds from the same. The underlying mechanism for doing so should be an XHTML form, which should be submitted back to your `Prefs` servlet by way of HTTP GET or POST. You may find that checkboxes are the most aesthetically pleasing choice of input mechanisms. However, the servlet should *not* output a list of all available newsfeeds in one, large page. Instead, we suggest that you display only one channel of categories at a time, allowing the user to navigate among the available channels by way of hyperlinks or forms, which submit some parameter that instructs your `Prefs` servlet to display a different channel.
 - ii. The servlet should utilize `project3-6.0/webapps/ROOT/xml/prefs.xml` in order to generate its pages dynamically.

²⁴ Realize that this `String`'s dynamic generation is inelegant (if not an abuse of XML); we recommend it here only so that you can focus on other aspects of the project.

- iii. If a user attempts to visit the servlet prior to authenticating, you should redirect the user to the `Login` servlet.
- iv. Any page generated by the servlet should contain a link to the `View` servlet.
- v. Any page generated by the servlet should contain a link via which the user can log out (*i.e.*, the link will pass some parameter back to the servlet, instructing it to invalidate the current session and redirect the user to the `Login` servlet).

Again, be sure to call the `UserManager` class's `save()` method anytime a user updates his or her preferences.

As in the `View` servlet, you may find it useful (if not necessary) to concatenate an XML representation of the user's current favorites with an XML representation of all available categories so that the servlet is able to report simultaneously (by way of `prefs.xml`) the user's current favorites as well as other, available categories (in, say, some channel).

Again, although we've configured Tomcat to check the modification dates and times of class files, you may want to restart Tomcat anytime you recompile your source.

Alright, make it happen!

9. (40 points.) Pretty neat stuff, eh? By now, you're hopefully feeling pretty good about your application. And you should; you've pulled together a lot of material, a lot of skills.

It's time, then, to make that portal truly yours. Although we claim to have left the design of your portal largely up to you, it was certainly our idea to have a login page, a viewing page, and a preferences page. What functionality did we leave out, though? What would make you want to keep that instance of Tomcat running indefinitely, your portal available 24/7?

Go ahead and add a personal touch to your portal. It should go without saying that you're welcome to beautify your portal with colors and images. But we'd like you to extend your portal's functionality. The extension could be in the client tier. What's lacking in your portal's user interface? The extension could also be in the middle tier. What feature are your servlets missing? What servlet is missing? Or, the extension could be in the back-end tier. Right now, your code (presumably) relies upon the copy of `nested_category_list.html` in `project3-6.0/webapps/ROOT/xml/cache/`, which, again, was generated by Moreover at 3:13 P.M. ET on Friday, 3 March 2006. What if Moreover's offerings have been updated? Perhaps your portal should be retrieving a fresh listing on occasion, caching the results. In fact, for the sake of performance, perhaps the `NewsProvider` class should be caching headlines, at least for a short while, so that it needn't open up a TCP socket for every request! Also, what about those DTDs? Should you be validating your XML feeds? Why not spend time on all three tiers? Isn't RSS all the rage these days? Why not support feeds from other sources?

In short, make your portal better. Feel free to discuss (but not collaborate on the implementation of) possible extensions with other students by way of `cscie259@lists.dce.harvard.edu`. You may certainly incorporate additional technologies

(*e.g.*, JSP, XSL-FO, *etc.*) into your extension. We ask only that the scope of your extension be comparable to that required by question 7 or 8; in other words, the extension should demand of you about the same amount of time as did your `Login` servlet or `Prefs` servlet. You are welcome (and encouraged) to run a possible extension by the staff if you're worried that it may fall short of our expectations.

Once you've designed and implemented your extension, create in `project3-6.0/` a file called `personaltouch.{html,pdf,ps,rtf,txt}` containing two or more paragraphs detailing your extension. Be sure to explain what your extension is, how and in what file(s) you implemented it, and what challenges you encountered during its implementation.

10. (2 points.) Wahoo! You're done. Two points!

Submitting Project 3.

11. (0 points.) If you have not done your work on `nice.fas.harvard.edu`, transfer via SFTP your `project3-6.0/` directory from your local machine to the `cscie259/` directory in your FAS account's home directory, taking care to upload any text files in ASCII mode.

Next, ensure that the structure of your `project3-6.0/` directory on `nice.fas.harvard.edu` is the following.

```
project3-6.0/  
  conf/  
  src/  
    cscie259/  
      project3/  
        wahoo/  
temp/  
webapps/  
  ROOT/  
    docs/  
      cscie259/  
        project3/  
          wahoo/  
    dtd/  
    images/  
    WEB-INF/  
    xml/  
      cache/  
    xsl/
```

Your `project3-6.0/` directory should contain `build.xml` (updated, if necessary, to handle any files of your own creation) and `personaltouch.{html,pdf,ps,rtf,txt}`; `project3-6.0/conf/` should contain your application's configuration files; `project3-6.0/webapps/ROOT/docs/` should contain updated Javadoc; `project3-6.0/webapps/ROOT/dtd/` should contain those DTDs it originally contained, along with any others you might have created; `project3-6.0/webapps/ROOT/images/`

should contain the image it originally contained, along with any others you might have created; `project3-6.0/webapps/ROOT/WEB-INF/` should contain `web.xml`; `project3-6.0/src/` should contain the Java source it originally contained, along with any additional Java source you wrote; `project3-6.0/webapps/ROOT/xml/` should contain those XML files it originally contained, along with any others you might have created; and `project3-6.0/webapps/ROOT/xsl/` should contain those XSLT files it originally contained, along with any others you might have created.

Particularly if you developed on another machine, ensure that your application builds and executes properly on `nice.fas.harvard.edu`.

Once everything's in order, clean up your workspace by executing

```
ant clean
```

from within your `project3-6.0/` directory.²⁵

Finally, submit your work electronically by executing the following command from within your `project3-6.0/` directory.

```
cscie259submit project3
```

Thereafter, follow any on-screen instructions until you receive visual confirmation of your project's successful submission. You may re-submit as many times as you'd like; each re-submission will overwrite any previous submission. But take care not to re-submit after the project's deadline, as only your latest submission's timestamp is retained.

²⁵ Know that this command deletes not only your bytecodes but also Tomcat's `logs`, runtime directories, and temporary files.