

Project 4
Scamazon.com
140 Points

due by 12:30 P.M. ET on Thursday, 27 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 recommended that you begin answering this document's questions immediately thereafter. After all, the sooner you finish, the more time you'll have for your final project!

And it is also recommended that you request, per problem 11 herein, a "subscription ID" for Amazon's Web Services immediately, as it may take some time for your request to be processed.

Goals.

The goals of this project are to:

- Challenge you to build your own e-commerce engine and fulfillment partner.
- Give you hands-on experience with DTD, SOAP, web services, and XML Schema, as well as additional experience with JavaServer Pages 2.0, Java Servlet 2.4, TrAX, XSL-FO, and XSLT.
- Introduce you to Apache's Axis 1.3 and to FOP 0.20.5's API.
- Synthesize Lectures 1 through 11 and equip you for your final project and future XML endeavors!



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? Do your files contain lines of code that have been commented out? ³

Academic Honesty.

Again for this project, you will be running your own instance of a webserver on `nice.fas.harvard.edu` which 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 DTD, Java, JSP, XHTML, XML, XML Schema, XPath, XSL-FO, and XSLT.

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

³ They shouldn't.

Getting Started.

1. (0 points.) This project expects use of J2SE SDK 1.4.2, Apache's Ant 1.6.5, Apache's Axis 1.3, Apache's FOP 0.20.5, 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, all of which have been installed on `nice.fas.harvard.edu` and are available for download via the course's website.⁴
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.^{5,6}

```
cp -r ~cscie259/pub/distribution/projects/project4-6.0/ ~/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/      project4-6.0/
```

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 `project4-6.0/` 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. Be sure, of course, to download and install the requisite versions of J2SE SDK, Ant, Axis, FOP, Tomcat, Xalan-J, and Xerces-J, as well as Sun's JavaBeans Activation Framework (JAF) 1.0.2 and JIMI. Finally, be sure to examine `cscie259.cshrc` in `~cscie259/pub/etc/` on `nice.fas.harvard.edu`, so that you know how to configure your machine.^{7,8}

⁴ Provided you ran `cscie259setup` for Project 1, your FAS account is automatically configured for these versions.

⁵ These commands assume that you already created (for Projects 1, 2, and 3) a directory called `cscie259/` in your FAS account's home directory.

⁶ 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. Henceforth, however, all footnotes assume a Windows platform.

⁷ 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 5.5.16'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 5.5.16's `common/lib/` directory; copy `axis.jar`, `commons-discovery-0.2.jar`, `commons-logging-1.0.4.jar`, `jaxrpc.jar`, `log4j-1.2.8.jar`, `saaj.jar`, and `wsdl4j-1.5.1.jar` from Axis 1.3's `lib/` directory into Tomcat 5.5.16's `common/lib/` directory; copy `avalon-framework-cvs-20020806.jar`, `batik.jar`, and `fop.jar`, all of which ship with FOP 0.20.5, into Tomcat 5.5.16's `common/lib/` directory; copy `log4j.properties` from Axis 1.3's `lib/` directory and `i18n.properties` from Axis 1.3's `webapps/axis/WEB-INF/classes/` directory into Tomcat 5.5.16's `common/classes/` directory; copy `activation.jar` (which ships with JAF 1.0.2) into Tomcat 5.5.16's `common/lib/` directory; copy `JimiProClasses.zip` (which ships with JIMI) into Tomcat 5.5.16's `common/lib/` directory, renaming it `jimi-1.0.jar`; optionally copy `mailapi.jar` (which ships with JavaMail 1.3.3) and `xmlsec-1.3.0.jar` (which ships with Apache-XML-Security-J 1.3.0) into Tomcat 5.5.16'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 5.5.16; define an environment variable called

Notice, now, that your `project4-6.0/` directory is structured as follows.

```
project4-6.0/  
  conf/  
  src/  
    com/  
      amazon/  
        soap/  
          AWSECommerceService/
```

"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"; prepend "\$JAVA_HOME/bin" and "\$ANT_HOME/bin" to your PATH; define an environment variable called D_LOCAL_HOST whose value is "-Dlocal.host=localhost" or "-Dlocal.host=127.0.0.1"; define an environment variable called D_VALIDATION_OFF whose value is ' ' (*i.e.*, a single space); define an environment variable called D_VALIDATION_DTD whose value is "-Denable.validation=dtd"; define an environment variable called D_VALIDATION_XSD whose value is "-Denable.validation=xsd"; define an environment variable called CATALINA_OPTS whose value is "\$D_LOCAL_HOST \$D_VALIDATION_OFF"; add to your CLASSPATH avalon-framework-cvs-20020806.jar, batik.jar, and fop.jar, all of which ship with FOP 0.20.5; and add to your CLASSPATH axis.jar, commons-discovery-0.2.jar, commons-logging-1.0.4.jar, jaxrpc.jar, log4j-1.2.8.jar, saaj.jar, and wsdl4j-1.5.1.jar, all from Axis 1.3's lib/ directory. None of the values of these environment variables, by the way, should contain the quotes with which we've flanked them in this text.

⁸ 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 5.5.16'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 5.5.16's common\lib\ directory; copy axis.jar, commons-discovery-0.2.jar, commons-logging-1.0.4.jar, jaxrpc.jar, log4j-1.2.8.jar, saaj.jar, and wsdl4j-1.5.1.jar from Axis 1.3's lib\ directory into Tomcat 5.5.16's common\lib\ directory; copy avalon-framework-cvs-20020806.jar, batik.jar, and fop.jar, all of which ship with FOP 0.20.5, into Tomcat 5.5.16's common\lib\ directory; copy log4j.properties from Axis 1.3's lib\ directory and il8n.properties from Axis 1.3's webapps\axis\WEB-INF\classes\ directory into Tomcat 5.5.16's common\classes\ directory; copy activation.jar (which ships with JAF 1.0.2) into Tomcat 5.5.16's common\lib\ directory; copy JimiProClasses.zip (which ships with JIMI) into Tomcat 5.5.16's common\lib\ directory, renaming it jimi-1.0.jar; optionally copy mailapi.jar (which ships with JavaMail 1.3.3) and xmlsec-1.3.0.jar (which ships with Apache-XML-Security-J 1.3.0) into Tomcat 5.5.16'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 5.5.16; 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"; prepend "%JAVA_HOME%\bin" and "%ANT_HOME%\bin" to your PATH; define an environment variable called D_LOCAL_HOST whose value is "-Dlocal.host=localhost" or "-Dlocal.host=127.0.0.1"; define an environment variable called D_VALIDATION_OFF whose value is ' ' (*i.e.*, a single space); define an environment variable called D_VALIDATION_DTD whose value is "-Denable.validation=dtd"; define an environment variable called D_VALIDATION_XSD whose value is "-Denable.validation=xsd"; define an environment variable called CATALINA_OPTS whose value is "%D_LOCAL_HOST% %D_VALIDATION_OFF%"; add to your CLASSPATH avalon-framework-cvs-20020806.jar, batik.jar, and fop.jar, all of which ship with FOP 0.20.5; and add to your CLASSPATH axis.jar, commons-discovery-0.2.jar, commons-logging-1.0.4.jar, jaxrpc.jar, log4j-1.2.8.jar, saaj.jar, and wsdl4j-1.5.1.jar, all from Axis 1.3's lib\ directory. None of the values of these environment variables, by the way, should contain the quotes with which we've flanked them in this text.

```
cscie259/  
  project4/  
    scamazon/  
      warehouse/  
temp/  
webapps/  
  ROOT/  
    docs/  
      com/  
        amazon/  
          soap/  
            AWSECommerceService/  
cscie259/  
  project4/  
    scamazon/  
      warehouse/  
WEB-INF/  
scamazon/  
  dtd/  
  fo/  
  images/  
  WEB-INF/  
    lib/  
  xml/  
    cache/  
  xsd/  
  xsl/  
warehouse/  
  dtd/  
  images/  
  WEB-INF/  
  xml/  
  xsd/  
  xsl/
```

Ensure that your account or machine is in order by executing the following command from within any directory.⁹

```
EnvironmentCheck
```

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

Next, execute the following command from within your `project4-6.0/` directory.¹⁰

```
ant compile
```

⁹ Note that, on `nice.fas.harvard.edu`, ``EnvironmentCheck`` is an alias for ``java org.apache.xalan.xslt.EnvironmentCheck``.

¹⁰ Alternatively, you can execute ant without any arguments.

By default, ant will compile (among other things) `cscie259.project4.scamazon.*`, storing the resulting bytecodes in `project4-6.0/webapps/scamazon/WEB-INF/classes/`, and `cscie259.project4.warehouse.*`, storing the resulting bytecodes in `project4-6.0/webapps/warehouse/WEB-INF/classes/`.

As always, take a look at `project4-6.0/build.xml` so that you know what ant can do for you during this project.

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.¹¹

```
echo $ICEBOX
```

Next, proceed to edit `project4-6.0/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 `project4-6.0/`.¹²

```
tomcat
```

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.¹³

```
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:109f5f7ca02:-m:$ICEBOX.fas.harvard.edu:1
- Initializing Coyote HTTP/1.1 on http-n
- Initialization processed in 1133 ms
- Starting service Catalina
- Starting Servlet Engine: Apache Tomcat/5.5.16
- XML validation enabled
- Starting Coyote HTTP/1.1 on http-n
- Server startup in 2739 ms
```

¹¹ Though we referred to it as one system prior to Project 3, `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.

¹² 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 `\Windows\System32\`, `\Winnt\System32\`, or any other directory in your `PATH` a file called `tomcat.bat` containing the following two lines.

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

¹³ This output's report that XML validation is enabled refers only to Tomcat's validation of its own configuration files (e.g., this project's several `web.xml` files) not to your parsers' validation of XML inputs (e.g., `catalog.xml`). Validation of the latter is to be enabled manually, per this specification's discussion of `CATALINA_OPTS`.

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 be returned to your shell's prompt.) Confirm as much by visiting the URL below, taking care to substitute `$ICEBOX` and `n` with their appropriate values, per your `$ICEBOX` and Connector's port.^{14,15}

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

You should see a directory listing for your container's root directory (namely, `project4-6.0/webapps/ROOT/`), the contents of which include `docs/` (the project's initial Javadoc), `happyaxis.jsp` (a JSP that ships with Axis 1.3), and `happyenv.jsp` (a JSP that the staff whipped up). Confirm that the listing's footer reports version 5.5.16 of Tomcat.

Confirm also that Axis is functioning properly by clicking on `happyaxis.jsp`: all Needed Components and Optional Components should be present. That is, just above a horizontal rule, the page should report: *"The core axis libraries are present. The optional components are present."*¹⁶ Then, scroll down a bit and confirm that the application server (*i.e.*, Tomcat) implements version 2.4 of Java Servlet, that the JVM in use is 1.4.2, that the `local.host` system property is present and its value is your host's fully qualified domain name or IP address,¹⁷ and that the remaining properties appear as expected.

Next, go back to your root directory's listing and click on `happyenv.jsp`. Confirm that Tomcat is configured to use J2SE SDK 1.4.2, Ant 1.6.5, Xalan-J 2.7.0, and Xerces-J 2.8.0 and that the remaining properties appear as expected.

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

If your account or system ain't so happy, feel free to contact `cscie259@lists.dce.harvard.edu` or the staff for assistance.

Earth's Smallest Selection.

3. (0 points.) Your mission for this project, if you choose to accept it,¹⁸ is to build an e-commerce engine whose orders will be transmitted to and (not actually) fulfilled by a local warehouse. In fact, upon completion, your project will be the place (not) to shop! Visitors will be able to browse your project's dynamically generated aisle(s), add items to your project's

¹⁴ If you opt to develop on your own machine, you should instead visit `http://localhost:n/` or `http://127.0.0.1:n/` in lieu of this and all future references to `http://$ICEBOX.fas.harvard.edu:n/`.

¹⁵ Although you needn't type `"http://"` to visit most websites with most browsers these days, be sure to do so in this case.

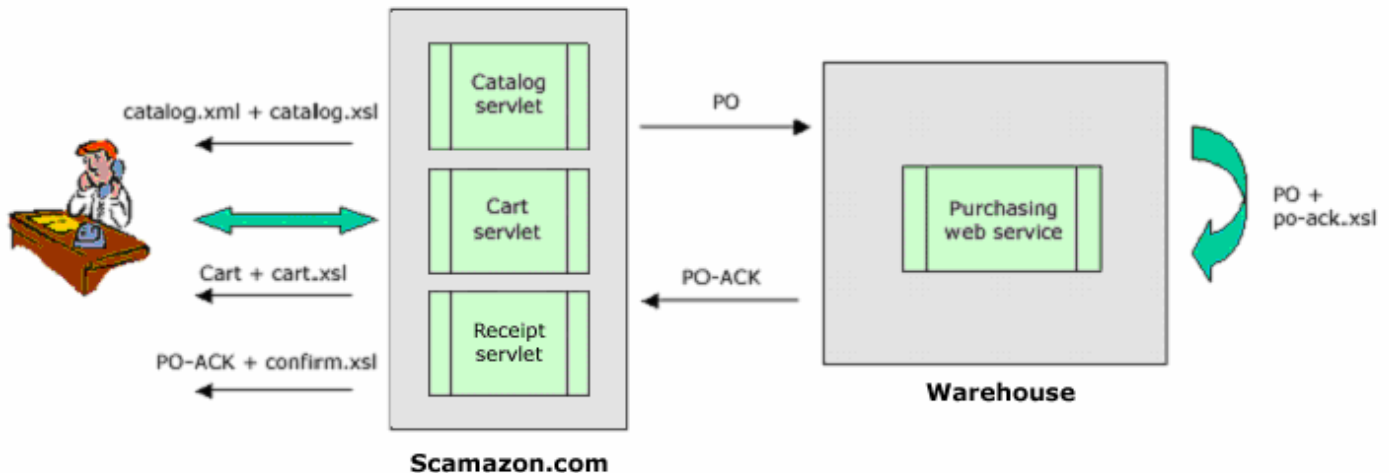
¹⁶ If you've opted to develop on your local machine, know that it is not necessary that the optional components be present. However, all "Needed Components" must be present.

¹⁷ Recall that this property should have been set by way of your `CATALINA_OPTS` environment variable.

¹⁸ On second thought, you must accept it.

shopping cart, and then check out, their order (never to be) fulfilled by your project's warehouse, a web service!
Quite the scam, eh?

Lest you be overwhelmed by thoughts of supply-chain management, inventory management, bottom dollars, and the like, rest assured that we've again provided you with a framework. Per Lecture 9, this project's architecture resembles the below.¹⁹



Before trying to wrap your mind around this architecture, take a look at `project4-6.0/webapps/scamazon/xml/catalog.xml`; the file details Scamazon.com's offerings. Next, take a look at `project4-6.0/webapps/warehouse/xml/inventory.xml`; the file specifies how many of each item is currently in stock at the warehouse.

Now, as the figure above suggests, Scamazon.com currently embodies three servlets: a Catalog servlet which will allow a visitor to browse the site's catalog, a Cart servlet which will allow that visitor to add items to his or her shopping cart and submit an order, and a Receipt servlet which will allow the user to obtain a receipt. Upon submission of an order, Scamazon.com will pass that order (in the form of an XML element called `PO`) along to the local warehouse, which is to be implemented as a web service called `Purchasing`. The warehouse will check its inventory and respond to the order (in the form of an XML element called `PO-ACK`), acknowledging which items are en route (er, would be, if this weren't a scam) and which are backordered. (For simplicity, the warehouse will operate out of the same instance of Tomcat. However, Scamazon.com will still communicate with the web service via SOAP; in fact, the former won't really know the latter is running on the same host.) The Cart servlet will then inform the user of the order's completion and allow the user to obtain a receipt in the form of a PDF.

Make sense? Okay, let's take a look at the code in `project4-6.0/src/cscie259/project4/scamazon/`.

¹⁹ The man on the phone at the desk, however, may not resemble you.

Start by examining Scamazon.com's Catalog servlet by way of `Catalog.java`; notice how it aspires to apply `project4-6.0/webapps/scamazon/xsl/catalog.xsl` to `project4-6.0/webapps/scamazon/xml/catalog.xml` in order to generate your project's aisle(s). Then take a look at the Cart servlet by way of `Cart.java`; eventually, this servlet will handle all operations involving the user's shopping cart. Next, examine the Receipt servlet by way of `Receipt.java`. Notice how, at present, this servlet allows you to submit via HTTP either the name of an XSL-FO file in `project4-6.0/webapps/scamazon/fo/` or the name of an XML file in `project4-6.0/webapps/scamazon/xml/` and the name of an XSLT file in `project4-6.0/webapps/scamazon/xsl/`, the response to which will be a dynamically generated PDF! Finally, give `Proxy.java` a once-over: it contains a proxy for (*i.e.*, a method with which to invoke) the warehouse's Purchasing service.

Let's now examine the code in `project4-6.0/src/cscie259/project4/warehouse/`. Wow, not all that much there, eh? Although `Purchasing.java` appears only to define a class called `Purchasing`, that class will be "magically" transformed into a web service by Axis, a SOAP engine we've integrated into Tomcat.

Did you notice, by the way, that `Catalog.java` and `Purchasing.java` are already capable of validating XML inputs, either by way of DTD or XML Schema?

If you prefer Javadoc over this stroll down directory lane, feel free to view `project4-6.0/webapps/ROOT/docs/` with your favorite browser. We've placed the project's initial Javadoc in `project4-6.0/webapps/ROOT/` so that you can view our (and, soon, your) Javadoc with your own instance of Tomcat!²⁰

Now take a look at this project's configuration. Look over `web.xml` in `project4-6.0/webapps/scamazon/WEB-INF/` to see how Scamazon.com is configured. Then look over the same in `project4-6.0/webapps/warehouse/WEB-INF/` to see how the warehouse is configured; also take a look at that directory's `server-config.wsdd`. You certainly needn't understand the entirety of those files' syntax; but try to get a sense of what each file is doing for you.

Alright, on your mark, get set—

Go!

4. (5 points.) Go ahead and spawn Tomcat again from within your `project4-6.0/` directory. This time, though, point your browser at the URL below.

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

You should be immediately redirected to the URL below.

²⁰ During the development of this project, you can update your Javadoc by executing ``ant javadoc`` from within your `project4-6.0/` directory.

```
http://$ICEBOX.fas.harvard.edu:n/scamazon/servlet/catalog
```

Not the sexiest of online catalogs yet, is it?

Go ahead and kill Tomcat. Then execute the following command (at the same terminal in which you've been running Tomcat), the effect of which will be to enable validation by DTD.²¹

```
setenv CATALINA_OPTS "$D_LOCAL_HOST $D_VALIDATION_DTD"
```

Spawn Tomcat again from within your project4-6.0/ directory. Then re-visit the URL below.²²

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

Notice that a blank page is returned this time (generated automatically by Tomcat). Moreover, notice Tomcat's console: it should have reported that no grammar was found. That is, catalog.xml failed validation because it hasn't even a DTD associated with it (yet)!

Well this is no good. Go ahead and define a DTD for catalog.xml in project4-6.0/webapps/scamazon/dtd/catalog.dtd. Rest assured that more than one solution may be possible. But, in the interests of defending your application against bad data, try to make your DTD as restrictive as possible. When ready to test your design, make catalog.dtd an external subset for catalog.xml by adding the following just below the latter's XML declaration.

```
<!DOCTYPE items SYSTEM "webapps/scamazon/dtd/catalog.dtd">
```

Perfect your DTD by refreshing your view of Scamazon.com's (minimalist) catalog and watching Tomcat's console, tweaking the DTD as necessary.

You may not, however, alter the structure of catalog.xml beyond adding the aforementioned DOCTYPE element.

5. (5 points.) Given the shortcomings of DTD, your catalog.dtd probably doesn't embody the ideal definition of Scamazon.com's catalog.²³ Let's take another stab at the problem, this time with the more powerful, the more flexible, the more, er, complicated XML Schema!

²¹ Windows users should instead execute the below.

```
set CATALINA_OPTS=%D_LOCAL_HOST% %D_VALIDATION_DTD%
```

²² Be sure to reload the page; beware caching browsers.

²³ No matter how hard you tried!

First, kill Tomcat yet again and remove from `catalog.xml` the `DOCTYPE` declaration you just inserted. Then execute the following command (at the same terminal in which you've been running Tomcat), the effect of which will be to enable validation by XML Schema.²⁴

```
setenv CATALINA_OPTS "$D_LOCAL_HOST $D_VALIDATION_XSD"
```

Spawn Tomcat yet again from within your `project4-6.0/` directory. Then re-visit the URL below.

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

Again a blank page should be returned. And, again, Tomcat's console has something to say: it should have reported an inability to find the declaration of `items`. Well, of course: `catalog.xml` failed validation because it hasn't an XML schema associated with it!

In `project4-6.0/webapps/scamazon/xsd/catalog.xsd`, define an XML schema for `catalog.xml`. Be less concerned about imposing order and more concerned about enforcing data types and constraining values. Who cares, after all, if an item's name element precedes its category element?²⁵ When ready to test your design, change `catalog.xml`'s root element to the below.

```
<items xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:noNamespaceSchemaLocation="webapps/scamazon/xsd/catalog.xsd">
```

As before, perfect your schema by refreshing your view of Scamazon.com's Catalog servlet and watching Tomcat's console, tweaking the schema as necessary.

6. (0 points.) Once your DTD and schema are working, you may, for the sake of marginally better performance, disable validation by killing Tomcat (at the same terminal in which you've been running Tomcat) and then restarting it, after executing the following command.²⁶

```
setenv CATALINA_OPTS "$D_LOCAL_HOST $D_VALIDATION_OFF"
```

After all, `catalog.xml` is static, so there's no need to validate it repeatedly.

Checking In.

7. (30 points.) Let's get those virtual aisles up and running.

²⁴ Windows users should instead execute the below.

```
set CATALINA_OPTS=%D_LOCAL_HOST% %D_VALIDATION_XSD%
```

²⁵ Well, some applications might. But hopefully not yours. :-)

²⁶ Windows users should instead execute the below.

```
set CATALINA_OPTS=%D_LOCAL_HOST% %D_VALIDATION_OFF%
```

Complete the implementation of Scamazon.com's online catalog by editing `catalog.xml` and `Catalog.java` as you see fit. For now, you needn't worry about any interaction with the site's shopping cart. (However, you're welcome to develop the catalog in lockstep with the `Cart` servlet, whose requirements are discussed later in this document.) Although the design of this catalog is largely up to you, we do ask that you adhere to the following guidelines.

- i. The catalog must somehow display the values of every `label` attribute and every text node in `catalog.xml`; it's not necessary, however, to display the values of any key attributes. Rather than display images' file names, however, you should display the images themselves, all of which can be found in `project4-6.0/webapps/scamazon/images/`.²⁷ Though you probably thought on `catalog.xml`'s structure during development of your DTD and schema, be sure you understand it fully before attempting to display the file's contents.²⁸

It's fine to display the catalog's entire contents in one (large) page (*i.e.*, "aisle"), but you're welcome to partition the catalog into multiple pages. In fact, it'd be neat if you allowed the user to browse the catalog by category; but such is not required.

- ii. Needless to say, your site's aisle(s) should be dynamically generated by some transformation of `catalog.xml` and not defined statically in XHTML files.

Although works of art are not required, they are certainly encouraged!

Checking Out.

8. (30 points.) Let's grab a shopping cart now and take it out for a spin.

Implement, in `Cart.java`, some kind of shopping cart and modify your virtual aisle(s) so that a user can somehow add one or more items to that cart via hyperlinks (that reference the `Cart` servlet and include parameters) or via XHTML forms (which can be posted to the `Cart` servlet). Technically speaking, a shopping cart is little more than some mechanism for remembering (for the duration of the current session) which items a user would like to buy.

Although the design of this shopping cart is largely up to you, we ask that you adhere to the following guidelines.

- i. The `Cart` servlet must allow a user to add items to and remove items from his or her shopping cart. Moreover, it must be possible somehow to update the quantities of items in the shopping cart (either by submitting specific values or, perhaps more simply, by adding or removing items).

²⁷ Feel free to edit or resize these images or even replace them with nicer ones.

²⁸ If, after significant thought, you're still puzzled by some aspect of `catalog.xml`, feel free to contact the staff or fellow students.

- ii. It must be possible to view the contents of one's shopping cart. Not only should the shopping cart display the quantity of each item in the cart, it must also display each item's category, name, description, descriptor, and price. The shopping cart must also display the total cost of the items therein. (If you'd like to be a geek and compute shipping and tax, feel free. Heck, you can even accept coupons.) All of this information must be assembled by the Cart servlet into a String of XML (structured however you see fit), which must then be passed to `cart.xml` (in `project4-6.0/webapps/scamazon/xml/`) for transformation into XHTML. So that we know the structure of your String of XML, create in `project4-6.0/webapps/scamazon/dtd/cart.dtd` a DTD or in `project4-6.0/webapps/scamazon/xsd/cart.xsd` an XML schema. (You needn't create both.) Though your DTD or schema must be correct, you needn't configure your Cart servlet validate your XML; after all, since the XML is for Scamazon.com's internal use, it's presumably always valid.
- iii. While browsing Scamazon.com's aisle(s), it must be possible to view the contents of one's shopping cart, perhaps via some hyperlink.
- iv. While viewing one's shopping cart, it must be possible to return to Scamazon.com's aisle(s), perhaps via some hyperlink.
- v. When viewing one's shopping cart, it must be possible to check out. That is, it must be possible to submit the requisite contents of that shopping cart in the form of a PO element (*i.e.*, a String of XML whose root element is PO) to the warehouse's Purchasing service. That String must be submitted to the web service by way of the Proxy class's `processPO(.,.,.)` proxy method.²⁹ Upon receipt of an acknowledgement from the web service in the form of a PO-ACK element (*i.e.*, a String of XML whose root element is PO-ACK), the Cart servlet should display a confirmation of the user's order (again displaying each purchased item's category, name, description, descriptor, and price). However, the Cart servlet should take care to inform the user of how many of each item will actually be shipped (*i.e.*, whether some item was backordered as a result of insufficient availability at the warehouse). This confirmation should be in the form of a webpage, generated dynamically by `confirm.xml` (in `project4-6.0/webapps/scamazon/xml`); as before, the Cart servlet must apply this stylesheet to a String of XML, structured however you see fit. So that we know the structure of your String of XML, create in `project4-6.0/webapps/scamazon/dtd/confirm.dtd` a DTD or in `project4-6.0/webapps/scamazon/xsd/confirm.xsd` an XML schema. (You needn't create both.) Also, so that we know the structure of your PO element, create in

²⁹ Know that you can avoid hardcoding the URL of the warehouse's Purchasing service into your code. In fact, you can obtain dynamically the service's hostname by executing the below.

```
System.getProperty("local.host");
```

And you can obtain the service's port number by executing (from within some servlet) the below.

```
request.getServerPort();
```

- `project4-6.0/webapps/scamazon/dtd/po.dtd` a DTD or in `project4-6.0/webapps/scamazon/xsd/po.xsd` an XML schema. (You needn't create both.) Though your designs must be correct, you needn't validate your XML.
- vi. Upon checkout and confirmation thereof, it must be possible for a user to obtain a receipt for his or her order as a PDF, perhaps by clicking a hyperlink. That receipt should contain at least as much information as is required for the XHTML-based confirmation screen. Know that the `Receipt` servlet is already functional; it uses FOP 0.20.5's API. In fact, try it out on some XSL-FO file or some XML file and XSLT file which together produce XSL-FO. However, at present, the servlet only allows for the processing of files on the local file system. More convenient, perhaps, would be a tool that takes a `String` of XML as input (in addition to the name of some XSLT file on the local system), lest your `Cart` servlet need to record a user's order on disk for the sole purpose of receipts. (After all, this whole site's a scam; you shouldn't have to remember users' orders!) How you implement receipts, then, is up to you. (You're welcome to overhaul `Receipts.java` or even delete it altogether, relocating some its functionality to some other servlet.) No matter your approach, upon Tomcat's termination, there should be no evidence on disk of your users' orders.³⁰ Dang it, this is a scam, alright?!

Think carefully about how much information your shopping cart has to remember about each item. Per this document's forthcoming discussion of the warehouse's `Purchasing` service, your `Cart` servlet need submit to that service, upon the user's request, only the values of items' key attributes. After all, recall the structure of `inventory.xml`. However, your shopping cart's going to need to display more than just the values of key attributes. It's probably more efficient to tuck a good amount of data away in a `Session` object or, if it's system-wide (*i.e.*, not user-specific) data, in a class variable, than it is to store only the values of key attributes in a `Session` object and constantly re-parse `catalog.xml`. In fact, you may find it helpful to define in `Cart.java` an inner class or two, instantiations of which can be tossed into a `Session` object for safekeeping. But, again, such implementation details are up to you.

The (Non-)Fulfillment Partner.

9. (30 points.) Let's make sure that warehouse is able to (pretend to) fulfill those orders that the `Cart` servlet wants to submit in the form of `PO` elements.

Notice that `Purchasing.java` is already designed to apply `project4-6.0/webapps/warehouse/xsl/po-ack.xsl` to whatever `String` of XML is passed into its `processPO(·)` method. Bearing in mind the structure you chose for the `Cart` servlet's dynamically generated `PO` element, augment `po-ack.xsl` so that it transforms a `PO` element into a `PO-ACK` element, whose structure is up to you. However, as you know from this document's discussion of the `Cart` servlet, that `PO-ACK` element had better contain enough information so that the `Cart` servlet can confirm (or deny) a user's order effectively.

³⁰ The mere existence of `project4-6.0/webapps/scamazon/fo/` does not imply that you need make use of it (let alone forever save users' receipts); it is present to facilitate trial of the staff's `Receipt` servlet.

Of course, `po-ack.xml` should respect the actual quantities of items in stock, as defined in `project4-6.0/webapps/warehouse/xml/inventory.xml`, ultimately reporting which items are in stock and which are backordered. It is not necessary to update `inventory.xml` dynamically based on POs processed; you may assume that the quantities specified therein are eternal.

You're welcome to modify `Purchasing.java` as you see fit.

So that we know the structure of your `PO-ACK` element, create in `project4-6.0/webapps/warehouse/dtd/po-ack.dtd` a DTD or in `project4-6.0/webapps/warehouse/xsd/po-ack.xsd` an XML schema. (You needn't create both.) Though your DTD or schema must be correct, you needn't configure your service to validate your XML; nor need you configure your service to validate the PO elements that it receives from Scamazon.com's Cart servlet.

10. (0 points.) Notice again that the `Purchasing` class makes no reference to its being a web service. Such is the beauty of a SOAP toolkit like Axis: it can transform Java classes into web services "magically"—that is, without your having to write a skeleton.

For simplicity, this project makes use of a handwritten proxy method with which your Cart servlet can access the web service.³¹ However, we could have generated a stub by first generating (or handwriting) the `Purchasing` service's WSDL, thereafter processing it with a SOAP toolkit. In fact, visit the URL below to see the WSDL that Axis generates for your `Purchasing` class.

`http://$ICEBOX.fas.harvard.edu:n/warehouse/services/Purchasing?wsdl`

Let's generate a Java stub from your `Purchasing` service's WSDL. Go ahead and execute the following command (which, unfortunately, wraps onto a second line).³²

```
java org.apache.axis.wsdl.WSDL2Java  
"http://$ICEBOX.fas.harvard.edu:n/warehouse/services/Purchasing?wsdl"
```

Check out the stub code that Axis emitted into

```
./edu/harvard/fas/$ICEBOX/warehouse/services/Purchasing/!
```

Now let's generate a Java skeleton from your `Purchasing` service's WSDL. Go ahead and execute the following command (which also, unfortunately, wraps onto a second line).³³

```
java org.apache.axis.wsdl.WSDL2Java --server-side --skeletonDeploy true  
"http://$ICEBOX.fas.harvard.edu:n/warehouse/services/Purchasing?wsdl"
```

³¹ Trust us, we thought through the logistics of having you generate skeletons and stubs and decided you it wasn't worth the additional complexity; launching your servlets and service within the confines of a single application server would no longer be as simple as executing ``ant`` followed by ``tomcat``.

³² On `nice.fas.harvard.edu`, we've aliased this command (sans URL) to ``WSDL2JavaStub``.

³³ On `nice.fas.harvard.edu`, we've aliased this command (sans URL) to ``WSDL2JavaSkeleton``.

Check out the skeleton code Axis has added to

`./edu/harvard/fas/$ICEBOX/warehouse/services/Purchasing/.`

You're welcome to try out this stub and skeleton code. But, if you do, be sure to reconfigure your project to use the provided proxy class prior to submission; do not submit a project that utilizes a dynamically generated stub or skeleton.

For additional details on dynamic generation of stubs and skeletons, follow the link to the User's Guide for Axis at <http://ws.apache.org/axis/>.

Scamazon.com. Amazon.com

11. (30 points.) It's time to add one last feature to your infrastructure. It turns out that your application's namesake, Amazon.com, offers access to its wares via Amazon Web Services (AWS)! Your final challenge, then, is to incorporate AWS's E-Commerce Service (ECS) 4.0 into Scamazon.com using SOAP! Perhaps you'd like to allow visitors to Scamazon.com to search Amazon.com's books by author, keyword, or title (so that they have somewhere to turn for their purchase once they realize your site's a scam). Or perhaps you'd like to allow visitors to browse Amazon.com's DVDs along with Scamazon.com's own "products." The possibilities are endless! Well, not quite, but it may feel that way when you peruse AWS's documentation.

In fact, surf on over to the URL below.

<http://www.amazon.com/webservices/>

Poke around the site for a bit, particularly the leftmost links entitled "FAQs" and "Amazon E-Commerce Service." Proceed to "Create an Account" in order to obtain a "subscription ID," making note of it upon receipt. You needn't download anything from Amazon, as we've included all that you need in this project's distribution. But do skim the "Online Documentation" available at the URL below, paying particular attention the section entitled "API Reference."

<http://www.awszone.com/>

Notice that Amazon allows you to access ECS not only via SOAP but also via REST (XML over HTTP) and XSLT.

Also take a peek at ECS's WSDL (for developers within the United States) at the URL below.

<http://webservices.amazon.com/AWSECommerceService/US/AWSECommerceService.wsdl>

Unfortunately, it's probably not immediately apparent from these resources alone how to employ the kit. So we'll get you started.

Notice that, in `project4-6.0/src/`, we've provided source code for `com.amazon.soap.AWSECommerceService.*`. We generated this code by executing, within `project4-6.0/src/`, the following command (which, unfortunately, wraps onto three lines).

```
java org.apache.axis.wsdl.WSDL2Java -v -W -p  
com.amazon.soap.AWSECommerceService  
http://webservices.amazon.com/AWSECommerceService/US/AWSECommerceService.wsdl
```

Throughout this project, Ant has been compiling the 168 source files that compose `com.amazon.soap.AWSECommerceService.*`, storing the resulting bytecodes in `project4-6.0/webapps/scamazon/WEB-INF/classes/`.

It's time to put those bytecodes to use. First, take a look at `project4-6.0/src/cscie259/project4/scamazon/ECS.java` and `project4-6.0/webapps/scamazon/ecs.jsp`. Then, go ahead and paste in your subscription ID for AWS into the former, per the file's comments, and surf on over to the URL below for a demonstration of ECS.

[http://\\$ICEBOX.fas.harvard.edu:n/scamazon/servlet/ecs](http://$ICEBOX.fas.harvard.edu:n/scamazon/servlet/ecs)

Needless to say, this servlet and JSP demonstrate ECS by querying Amazon.com for a "Large" set of details on the course's recommended books. Note that this demonstration happens to make use of Apache's Standard Taglib 1.1.2, an implementation of the JSP Standard Tag Library (JSTL), jars for which have been provided in `project4-6.0/webapps/scamazon/WEB-INF/lib/`. With JSTL are we able to confine this demonstration's logic to a servlet and its aesthetics to a JSP. You're welcome, but not required, to make use of JSTL in your own work; additional detail is available at the URL below.

<http://java.sun.com/products/jsp/jstl/reference/docs/>

For the curious, SOAP messages corresponding to this demonstration's request and response can be found in `project4-6.0/webapps/scamazon/xml/cache/`.

How you proceed to incorporate ECS into Scamazon.com is up to you, so long as you interact with it via SOAP, as we have done with our demonstration. (Though we offer this demonstration as a JSP, you might wish to incorporate ECS into one or more of your servlets.) But allow us to advise that simply incorporating into Scamazon.com some strangers' opinions on the course's recommended texts isn't terribly interesting (and certainly not worth 30 points).

Consider this exposure to ECS an opportunity to discover for yourself how to incorporate a real-world web service into code of your own. Unfortunately, ECS's documentation isn't terrific, but the API Reference, coupled with the Javadoc we've provided in `project4-6.0/webapps/ROOT/docs/`, should provide sufficient direction. But be sure to look to `project4-6.0/src/cscie259/project4/scamazon/ECS.java` and, optionally, `project4-6.0/webapps/scamazon/ecs.jsp` for hints on usage.

Discussion on the course's listserv of ECS's features and usage is both permissible and encouraged!

README.

12. (10 points.) In `project4-6.0/README.txt`, provide the teaching staff with an overview of your implementation of Scamazon.com and its warehouse.³⁴ Specifically, summarize how you implemented the project's catalog, shopping cart, and fulfillment partner, highlighting interesting implementation details. Also summarize how you incorporated ECS into your application. Your discussion should be at least three paragraphs and should well prepare the staff for a closer reading of your source code.

Submitting Project 4.

13. (0 points.) If you have not done your work on `nice.fas.harvard.edu`, transfer via SFTP your `project4-6.0/` 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.

Particularly if you developed on another machine, ensure that your application builds and executes properly on `nice.fas.harvard.edu`; misconfigured submissions may be ineligible for full credit.

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

```
ant clean
```

from within your `project4-6.0/` directory.³⁵

Next, ensure that the structure of your `project4-6.0/` directory on `nice.fas.harvard.edu` is the following.

```
project4-6.0/  
  conf/  
  src/  
    com/  
      amazon/  
        soap/  
          AWSECommerceService/  
cscie259/  
  project4/  
    scamazon/  
      warehouse/
```

³⁴ The MIME type of this file, so to speak, should remain `text/plain`.

³⁵ Know that this command deletes not only your bytecodes but also Tomcat's logs, runtime directories, and temporary files.

```
temp/  
webapps/  
  ROOT/  
    docs/  
      com/  
        amazon/  
          soap/  
            AWSECommerceService/  
cscie259/  
  project4/  
    scamazon/  
      warehouse/  
        WEB-INF/  
scamazon/  
  dtd/  
  fo/  
  images/  
  WEB-INF/  
    lib/  
  xml/  
    cache/  
  xsd/  
  xsl/  
warehouse/  
  dtd/  
  images/  
  WEB-INF/  
  xml/  
  xsd/  
  xsl/
```

Moreover, ensure that each directory contains exactly those files required by this specification; do not submit files which are unnecessary for an evaluation of your work, such as outdated source files, copies of original files, unnecessary output files, *etc.*

However, to facilitate your work's evaluation, please do not remove your subscription ID for AWS from your source code.

Finally, submit your work electronically by executing the following command from within your `project4-6.0/` directory.

```
cscie259submit project4
```

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 resubmission 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.

14. (0 points.) On to the final project!³⁶

³⁶ Note to self: use fewer footnotes.