

Axiom Developer Guide

Axiom Developer Guide

1.2.12

Licensed to the Apache Software Foundation (ASF) under one or more contributor license agreements. See the NOTICE file distributed with this work for additional information regarding copyright ownership. The ASF licenses this file to you under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Table of Contents

1. Testing	1
Unit test organization	1
Testing Axiom with different StAX implementations	1
2. Release process	3
Release preparation	3
Prerequisites	8
Release	8
Post-release actions	11
References	11
3. The StAX specification	12
Semantics of the <code>setPrefix</code> method	12
The three <code>XMLStreamWriter</code> usage patterns	13
A. Appendix	15
Installing IBM's JDK on Debian Linux	15

List of Figures

2.1. Package dependencies for r944680	4
2.2. Package dependencies for r939984	6

Chapter 1. Testing

Unit test organization

Historically, all unit tests were placed in the `axiom-tests` project. One specific problem with this is that since all tests are in a common Maven module which depends on both `axiom-impl` and `axiom-dom`, it is not rare to see DOOM tests that accidentally use the LLOM implementation (which is the default). The project description in `axiom-tests/pom.xml` indicates that it was the intention to split the `axiom-tests` project into several parts and make them part of `axiom-api`, `axiom-impl` and `axiom-dom`. This reorganization is not complete yet¹. For new test cases (or when refactoring existing tests), the following guidelines should be applied:

1. Tests that validate the code in `axiom-api` and that do not require an Axiom implementation to execute should be placed in `axiom-api`. This primarily applies to tests that validate utility classes in `axiom-api`.
2. The code of unit tests that apply to all Axiom implementations and that check conformance to the specifications of the Axiom API should be added to `axiom-api` and executed in `axiom-impl` and `axiom-dom`. Currently, the recommended way is to create a base class in `axiom-api` (with suffix `TestBase`) and to create subclasses in `axiom-impl` and `axiom-dom`. This makes sure that the DOOM tests never accidentally use LLOM (because `axiom-impl` is not a dependency of `axiom-dom`).
3. Tests that check integration with other libraries should be placed in `axiom-integration`. Note that this is the only module that requires Java 1.5 (so that e.g. integration with JAXB2 can be tested).
4. Tests related to code in `axiom-api` and requiring an Axiom implementation to execute, but that don't fall into category 2 should stay in `axiom-tests`.

Testing Axiom with different StAX implementations

The following StAX implementations are available to test compatibility with Axiom:

Woodstox

This is the StAX implementation that Axiom uses by default.

Sun Java Streaming XML Parser (SJSXP)

This implementation is available as Maven artifact `com.sun.xml.stream:sjsxp:1.0.1`.

StAX Reference Implementation

The reference implementation was written by BEA and is available as Maven artifact `stax:stax:1.2.0`. The homepage is <http://stax.codehaus.org/Home>. Note that the JAR doesn't contain the necessary files to enable service discovery. Geronimo's implementation of the StAX API library will not be able to locate the reference implementation unless the following system properties are set:

```
javax.xml.stream.XMLInputFactory=com.bea.xml.stream.MXP  
javax.xml.stream.XMLOutputFactory=com.bea.xml.stream.XML
```

¹See WSCOMMONS-419 [<https://issues.apache.org/jira/browse/WSCOMMONS-419>].

XL XP-J

“XL XML Processor for Java” is IBM's implementation of StAX 1.0 and is part of IBM's JRE/JDK v6. Note that due to an agreement between IBM and Sun, IBM's Java implementation for the Windows platform is not available as a separate download, but only bundled with another IBM product, e.g. WebSphere Application Server for Developers [<http://www.ibm.com/developerworks/downloads/ws/wasdevelopers/>].

On the other hand, the JDK for Linux can be downloaded as a separate package from the developerWorks site. There are versions for 32-bit x86 (“xSeries”) and 64-bit AMD. They are available as RPMs and tarballs. To install the JDK properly on a Debian based system (including Ubuntu), follow the instructions given in the section called “Installing IBM's JDK on Debian Linux”.

Chapter 2. Release process

Release preparation

The following items should be checked before starting the release process:

- Check for the latest Apache parent POM version (artifact `org.apache:apache`) and if necessary, change the parent of the Axiom root POM.
- Check the dependencies between Java packages in the `axiom-api` module. The `org.apache.axiom.util` package (including its subpackages) is specified to contain utility classes that don't depend on higher level APIs. More precisely, `org.apache.axiom.util` should only have dependencies on `org.apache.axiom.ext`, but not e.g. on `org.apache.axiom.om`. SonarJ [<http://www.hello2morrow.com/products/sonarj>] can be used to check these dependencies. The following figure shows the expected structure:

Figure 2.1. Package dependencies for r944680

In contrast, the following figure shows an earlier trunk version of `axiom-api` with incorrect layering and cyclic dependencies involving `org.apache.axiom.util`:

Figure 2.2. Package dependencies for r939984

The check can also be done using `jdepend-maven-plugin` [<http://mojo.codehaus.org/jdepend-maven-plugin/>]. To do this, execute the following command in the `axiom-api` module:

```
mvn jdepend:generate
```

Then open `target/site/jdepend-report.html` and go to the "Cycles" section. The report should not show any package cycles involving `org.apache.axiom.mime`, `org.apache.axiom.util` and `org.apache.axiom.ext`.

- Check that the generated Javadoc contains the appropriate set of packages. In particular, unit test related classes should be excluded, except for the test suite classes in `org.apache.axiom.ts` (we don't need to hide the fact that we have a reusable test suite...).
- Check that all dependencies and plugins are available from standard repositories. To do this, clean the local repository and execute **`mvn clean install`** followed by **`mvn site`**.
- Check that the set of license files in the `legal` directory is complete and accurate.
- Check that the Maven site conforms to the latest version of the Apache Project Branding Guidelines [<http://apache.org/foundation/marks/pmcs>].
- Check that the `apache-release` profile can be executed properly. To do this, issue the following command:

```
mvn clean install -Papache-release -Dmaven.test.skip=true
```

You may also execute a dry run of the release process:

```
mvn release:prepare -DdryRun=true
```

After this, you need to clean up using the following command:

```
mvn release:clean
```

- Prepare the release note. This should include a description of the major changes in the release as well as a list of resolved JIRA issues. Note that both `index.apx` and `RELEASE-NOTE.txt` need to be updated.
- Add an entry for the release to the `download.xml.vm` file and change the links for older releases so that they point to `archive.apache.org` (Since the Axiom project doesn't use branches and produces releases directly from the trunk, there should only be a single mirrored release).
- Preview and validate the changes that will be done by the release plugin to the POM files. In order to do this, execute the following command:

```
mvn release:prepare -DdryRun=true -Dmaven.test.skip=true
```

Next, compare the `pom.xml.tag` files to the corresponding `pom.xml` files:

```
for pom in $(find . -name "pom.xml"); do diff $pom $pom.tag; done
```

The differences should be limited to `version` and `scm` tags. If necessary, change the original POM files to avoid spurious changes. After that, clean up using:

```
mvn release:clean
```

Prerequisites

The following things are required to perform the actual release:

- A PGP key that conforms to the requirement for Apache release signing [<http://www.apache.org/dev/release-signing.html>]. To make the release process easier, the passphrase for the code signing key should be configured in `${user.home}/.m2/settings.xml`:

```
<settings>
...
<profiles>
  <profile>
    <id>apache-release</id>
    <properties>
      <gpg.passphrase> <!-- YOUR KEY PASSPHRASE --> </gpg.passphrase>
    </properties>
  </profile>
</profiles>
...
</settings>
```

- The release process uses a Nexus staging repository. Every committer should have access to the corresponding staging profile in Nexus. To validate this, login to `repository.apache.org` and check that you can see the `org.apache.ws` staging profile. The credentials used to deploy to Nexus should be added to `settings.xml`:

```
<servers>
...
<server>
  <id>apache.releases.https</id>
  <username><!-- ASF username --></username>
  <password><!-- ASF LDAP password --></password>
</server>
...
</servers>
```

Release

In order to prepare the release artifacts for vote, execute the following steps:

1. Update the release date in `download.xml.vm` and `index.apt`.
2. Temporarily disable the Jenkins build(s) for Axiom, in order to avoid accidental deployment of the release candidate to the local repository of a Jenkins executor if the release process fails somewhere in the middle and/or a Jenkins build starts at the wrong moment.
3. Start the release process with the following command:

```
mvn release:prepare
```

When asked for the "SCM release tag or label", override the default value (`axiom-x.y.z`) by entering a tag in the form `x.y.z`, which is compatible with the tag names used for previous releases.

The above command will create a tag in Subversion and increment the version number of the trunk to the next development version. It will also create a `release.properties` file that will be used in the next step.

4. Perform the release using the following command:

```
mvn release:perform
```

This will upload the release artifacts to the Nexus staging repository.

5. Log in to the Nexus repository (<https://repository.apache.org/>) and close the staging repository. The name of the staging profile is `org.apache.ws`. See <http://maven.apache.org/developers/release/apache-release.html> for a more thorough description of this step.
6. Generate and deploy the Maven site on a public Web server. You may use `people.apache.org` for this.
7. Start the release vote by sending a mail to `dev@ws.apache.org`. The mail should mention the following things:
 - The list of issues solved in the release (by linking to the relevant JIRA view).
 - The location of the Nexus staging repository.
 - The location where source and binary packages may be downloaded. This can be a reference to the location inside the staging repository.
 - A link to the preview of the Maven site.
8. Reenable the Jenkins build(s).

If the vote passes, execute the following steps:

1. Promote the artifacts in the staging repository. See <http://maven.apache.org/developers/release/apache-release.html> for detailed instructions for this step.
2. Log in to `people.apache.org` and publish the release distributions to `www.apache.org`. The `etc/dist.py` script can be used for that:

```
cd /www/www.apache.org/dist/ws/commons/axiom
umask 0002
python path_to_etc/dist.py version
```

`version` is the release version, e.g. 1.2.9.

If not yet done, export your public key and append it to the `KEYS` file located in `/www/www.apache.org/dist/ws/commons/axiom`. The command to export a public key is as follows:

```
gpg --armor --export key_id
```

3. Check out the release tag from Subversion and generate the Maven site. Also check out the existing site:

```
svn co https://svn.apache.org/repos/asf/webservices\
/axiom/site axiom-site
```

The existing site needs to be overwritten with the site generated for the new release. This can be done easily using the `etc/syncsite.py` script (requires Python 2.6). It will copy the files and execute any necessary Subversion commands to add new files and to delete files that no longer exist. After executing the script, commit the changes to Subversion.



When generating the site, please use a recent JDK version, because some older versions generate broken links.

4. Log in to `people.apache.org` and update the site:

```
cd /www/ws.apache.org/axiom/  
umask 0002  
svn update
```

The umask setting makes sure that other members of the ws group will be able to update the site as well.

It may take several hours before all the updates have been synchronized to the relevant ASF systems. Before proceeding, check that

- the Maven artifacts for the release are available from the Maven central repository;
- the Maven site has been synchronized to <http://ws.apache.org/axiom/>;
- the binary and source distributions can be downloaded from <http://ws.apache.org/axiom/download.cgi>.

Once everything is in place, send announcements to `users@ws.apache.org` and `announce@apache.org`. Since the two lists have different conventions, audiences and moderation policies, to send the announcement separately to the two lists.

Sample announcement:

Apache Axiom Team is pleased to announce the release of Axiom x.y.z. The release is available for download at:

<http://ws.apache.org/axiom/download.cgi>

Apache Axiom is a StAX-based, XML Infoset compliant object model which supports on-demand building of the object tree. It supports a novel "pull-through" model which allows one to turn off the tree building and directly access the underlying pull event stream. It also has built in support for XML Optimized Packaging (XOP) and MTOM, the combination of which allows XML to carry binary data efficiently and in a transparent manner. The combination of these is an easy to use API with a very high performant architecture!

Developed as part of Apache Axis2, Apache Axiom is the core of Apache Axis2. However, it is a pure standalone XML Infoset model with novel features and can be used independently of Apache Axis2.

Highlights in this release:

- ...
- ...

Resolved JIRA issues:

- [WSCOMMONS-513] Behavior of `insertSiblingAfter` and `insertSiblingBefore` is not well defined for orphan nodes
- [WSCOMMONS-488] The sequence of events produced by `OMStAXWrapper` with `inlineMTOM=false` is inconsistent

For `users@ws.apache.org`, the subject (“Axiom x.y.z released”) should be prefixed with “[ANN][Axiom]”, while for `announce@apache.org` “[ANN]” is enough. Note that mail to `announce@apache.org` must be sent from an `apache.org` address.

Post-release actions

- Update the DOAP file (see `etc/axiom.rdf`) and add a new entry for the release.
- Update the status of the release version in the AXIOM project in JIRA.
- Remove archived releases from `/www/www.apache.org/dist/ws/commons/axiom` on `people.apache.org`.

References

The following documents are useful when preparing and executing the release:

- ASF Source Header and Copyright Notice Policy [<http://www.apache.org/legal/src-headers.html>]
- Apache Project Branding Guidelines [<http://apache.org/foundation/marks/pmcs>]
- DOAP Files [<http://projects.apache.org/doap.html>]
- Publishing Releases [<http://www.apache.org/dev/release-publishing.html>]

Chapter 3. The StAX specification

The StAX specification comprises two parts: a specification document titled “Streaming API For XML JSR-173 Specification” and a Javadoc describing the API. Both can be downloaded from the JSR-173 page [<http://jcp.org/en/jsr/detail?id=173>]. Since StAX is part of Java 6, the Javadocs can also be viewed online [<http://java.sun.com/javase/6/docs/api/javax/xml/stream/package-summary.html>].

Semantics of the `setPrefix` method

Probably one of the more obscure parts of the StAX specifications is the meaning of the `setPrefix`¹ method defined by `XMLStreamWriter`. To understand how this method works, it is necessary to look at different parts of the specification:

- The Javadoc of the `setPrefix` method.
- The table shown in the Javadoc of the `XMLStreamWriter` class in Java 6².
- Section 5.2.2, “Binding Prefixes” of the specification.
- The example shown in section 5.3.2, “`XMLStreamWriter`” of the specification.

In addition, it is important to note the following facts:

- The terms *defaulting prefixes* used in section 5.2.2 of the specification and *namespace repairing* used in the Javadocs of `XMLStreamWriter` are synonyms.
- The methods writing namespace qualified information items, i.e. `writeStartElement`, `writeEmptyElement` and `writeAttribute` all come in two variants: one that takes a namespace URI and a prefix as arguments and one that only takes a namespace URI, but no prefix.

The purpose of the `setPrefix` method is simply to define the prefixes that will be used by the variants of the `writeStartElement`, `writeEmptyElement` and `writeAttribute` methods that only take a namespace URI (and the local name). This becomes clear by looking at the table in the `XMLStreamWriter` Javadoc. Note that a call to `setPrefix` doesn't cause any output and it is still necessary to use `writeNamespace` to actually write the necessary namespace declarations. Otherwise the produced document will not be well formed with respect to namespaces.

The Javadoc of the `setPrefix` method also clearly defines the scope of the prefix bindings defined using that method: a prefix bound using `setPrefix` remains valid till the invocation of `writeEndElement` corresponding to the last invocation of `writeStartElement`. While not explicitly mentioned in the specifications, it is clear that a prefix binding may be masked by another binding for the same prefix defined in a nested element.

An aspect that may cause confusion is the fact that in the example shown in section 5.3.2 of the specifications, the calls to `setPrefix` (and `setDefaultNamespace`) all appear immediately before a call to `writeStartElement` or `writeEmptyElement`. This may lead people to incorrectly believe that a prefix binding defined using `setPrefix` only applies to the next element written³. This interpretation is clearly in contradiction with the `setPrefix` Javadoc, unless one assumes that “the current `START_ELEMENT` / `END_ELEMENT` pair” means the element opened by a call to

¹For simplicity, we only discuss `setPrefix` here. The same remarks also apply to `setDefaultNamespace`.

³Another factor that contributes to the confusion is that in SAX, prefix mappings are always generated before the corresponding `startElement` event and that their scope ends with the corresponding `endElement` event. This is so because the `ContentHandler` interface specifies that “all `startPrefixMapping` events will occur immediately before the corresponding `startElement` event, and all `endPrefixMapping` events will occur immediately after the corresponding `endElement` event”.

`writeStartElement` immediately following the call to `setPrefix`. This however would be a very arbitrary interpretation of the Javadoc⁴.

The correctness of the comments in the previous paragraph can be checked using the following code snippet:

```
XMLOutputFactory f = XMLOutputFactory.newInstance();
XMLStreamWriter writer = f.createXMLStreamWriter(System.out);
writer.writeStartElement("root");
writer.setPrefix("p", "urn:ns1");
writer.writeEmptyElement("urn:ns1", "element1");
writer.writeEmptyElement("urn:ns1", "element2");
writer.writeEndElement();
writer.flush();
writer.close();
```

This produces the following output⁵:

```
<root><p:element1/><p:element2/></root>
```

Since the code doesn't call `writeNamespace`, the output is obviously not well formed with respect to namespaces, but it also clearly shows that the scope of the prefix binding for `p` extends to the end of the root element and is not limited to `element1`.

To avoid unexpected results and keep the code maintainable, it is in general advisable to keep the calls to `setPrefix` and `writeNamespace` aligned, i.e. to make sure that the scope (in `XMLStreamWriter`) of the prefix binding defined by `setPrefix` is compatible with the scope (in the produced document) of the namespace declaration written by the corresponding call to `writeNamespace`. This makes it necessary to write code like this:

```
writer.writeStartElement("p", "element1", "urn:ns1");
writer.setPrefix("p", "urn:ns1");
writer.writeNamespace("p", "urn:ns1");
```

As can be seen from this code snippet, keeping the two scopes in sync makes it necessary to use the `writeStartElement` variant which takes an explicit prefix. Note that this somewhat conflicts with the purpose of the `setPrefix` method; one may consider this as a flaw in the design of the StAX API.

The three `XMLStreamWriter` usage patterns

Drawing the conclusions from the previous section and taking into account that `XMLStreamWriter` also has a “namespace repairing” mode, one can see that there are in fact three different ways to use `XMLStreamWriter`. These usage patterns correspond to the three bullets in section 5.2.2 of the StAX specification⁶:

1. In the “namespace repairing” mode (enabled by the `javax.xml.stream.isRepairingNamespaces` property), the writer takes care of all

⁴Early versions of XL XP-J were based on this interpretation of the specifications, but this has been corrected. Versions conforming to the specifications support a special property called `javax.xml.stream.XMLStreamWriter.isSetPrefixBeforeStartElement`, which always returns `Boolean.FALSE`. This allows to easily distinguish the non conforming versions from the newer versions. Note that in contrast to what the usage of the reserved `javax.xml.stream` prefix suggests, this is a vendor specific property that is not supported by other implementations.

⁵This has been tested with Woodstox 3.2.9, SJSXP 1.0.1 and version 1.2.0 of the reference implementation.

⁶The content of this section is largely based on a reply posted by Tatu Saloranta on the Axiom mailing list [<http://markmail.org/message/olsdl3p3gcicqeqob>]. Tatu is the main developer of the Woodstox project.

namespace bindings and declarations, with minimal help from the calling code. This will always produce output that is well-formed with respect to namespaces. On the other hand, this adds some overhead and the result may depend on the particular StAX implementation (though the result produced by different implementations will be equivalent).

In repairing mode the calling code should avoid writing namespaces explicitly and leave that job to the writer. There is also no need to call `setPrefix`, except to suggest a preferred prefix for a namespace URI. All variants of `writeStartElement`, `writeEmptyElement` and `writeAttribute` may be used in this mode, but the implementation can choose whatever prefix mapping it wants, as long as the output results in proper URI mapping for elements and attributes.

2. Only use the variants of the writer methods that take an explicit prefix together with the namespace URI. In this usage pattern, `setPrefix` is not used at all and it is the responsibility of the calling code to keep track of prefix bindings.

Note that this approach is difficult to implement when different parts of the output document will be produced by different components (or even different libraries). Indeed, when passing the `XMLStreamWriter` from one method or component to the other, it will also be necessary to pass additional information about the prefix mappings in scope at that moment, unless it is acceptable to let the called method write (potentially redundant) namespace declarations for all namespaces it uses.

3. Use `setPrefix` to keep track of prefix bindings and make sure that the bindings are in sync with the namespace declarations that have been written, i.e. always use `setPrefix` immediately before or immediately after each call to `writeNamespace`. Note that the code is still free to use all variants of `writeStartElement`, `writeEmptyElement` and `writeAttribute`; it only needs to make sure that the usage it makes of these methods is consistent with the prefix bindings in scope.

The advantage of this approach is that it allows to write modular code: when a method receives an `XMLStreamWriter` object (to write part of the document), it can use the namespace context of that writer (i.e. `getPrefix` and `getNamespaceContext`) to determine which namespace declarations are currently in scope in the output document and to avoid redundant or conflicting namespace declarations. Note that in order to do so, such code will have to check for an existing prefix binding before starting to use a namespace.

Appendix A. Appendix

Installing IBM's JDK on Debian Linux

1. Make sure that `fakeroot` and `java-package` are installed:

```
# apt-get install fakeroot java-package
```

2. Download the `.tgz` version of the JDK from <http://www.ibm.com/developerworks/java/jdk/linux/download.html>.
3. Edit `/usr/share/java-package/ibm-j2sdk.sh` and (if necessary) add an entry for the particular version of the IBM JDK downloaded in the previous step.
4. Build a Debian package from the tarball:

```
$ fakeroot make-jpkg xxxx.tgz
```

5. Install the Debian package.