

Saxon - Bug #4706

jar scheme and xsl:include uniqueness

2020-09-02 09:57 - Air Quick

Status:	Closed	Start date:	2020-09-02
Priority:	Normal	Due date:	
Assignee:	Norm Tovey-Walsh	% Done:	100%
Category:	Saxon extensions	Estimated time:	0:00 hour
Sprint/Milestone:		Spent time:	0:00 hour
Legacy ID:		Fixed in Maintenance Release:	10.3
Applies to branch:	10, 9.9, trunk	Platforms:	
Fix Committed on Branch:	10, trunk		

Description

If you extract stylesheets from test.jar and run compile-xslt-tests.xml, it works without any warning:

```
$ jar xf test.jar
$ java -jar SaxonHE10-2J/saxon-he-10.2.jar -it -xsl:io/xspec/xspec/impl/src/compiler/compile-xslt-tests.xml
<?xml version="1.0" encoding="UTF-8"?>good
```

However, if you run the same stylesheet directly with jar: scheme, an "included more than once" warning happens followed by a "duplicate named function" error:

```
$ java -jar SaxonHE10-2J/saxon-he-10.2.jar -it -xsl:'jar:file:/tmp/test.jar!/io/xspec/xspec/impl/src/compiler/compile-xslt-tests.xml' -u
Warning at xsl:stylesheet on line 6 column 31 of compile-scenario.xml:
  Stylesheet module compile/compile-scenario.xml is included or imported more than once.
  This is permitted, but may lead to errors or unexpected behavior
Error at xsl:function on line 8 column 48 of compile-scenario.xml:
  XTSE0770 Duplicate named function (see line 8 of
  jar:file:/tmp/test.jar!/io/xspec/xspec/impl/src/compiler/base/compile/compile-scenario.xml)
Error at xsl:call-template on line 9 column 52 of main.xml:
  XTSE0650 Cannot find a template named compile-scenario
Errors were reported during stylesheet compilation
```

Must xsl:include/@href be literally unique when loaded with jar: scheme?

I searched the documentation but couldn't figure it out, except that [a change history](#) says, "it is possible to hold stylesheet modules in a JAR file and resolve xs:include and xs:import references between files within the same JAR archive."

History

#1 - 2020-09-02 10:49 - Michael Kay

The problem appears to be as follows: URI resolution (of the relative URI against the base URI) is actually done twice. First we do it within the Saxon code, to decide whether this is a reference to a stylesheet module that we already know about. If it's not, then we call the URIResolver to fetch the stylesheet - and because the URIResolver is supplied with the relative URI and base URI separately, it does it again.

The URI resolution process within the URIResolver includes code that recognises the JAR: scheme specially (these are not actually standard RFC URIs, and they need custom handling: this is done by using the java.net.URL class in place of java.net.URI). But the code that computes a document key to access the stylesheet module cache is doing a straightforward call on new URI(baseURI).resolve(href), and this gives the answer defined by RFC2396, which is not useful (the URI is classified as an "opaque" URI, so the rules of RFC2396 essentially mean that the relative URI is returned unchanged). The effect is that two xsl:include references using the same relative URI are considered to refer to the same stylesheet module when this is not actually the case.

We need to change the code used to compute a document key so it uses the same algorithm as the standard URI resolver

#2 - 2020-09-02 10:53 - Michael Kay

- Tracker changed from Support to Bug

- Category set to Saxon extensions

- Status changed from New to In Progress
- Assignee set to Norm Tovey-Walsh
- Priority changed from Low to Normal
- Applies to branch 10, 9.9, trunk added

#9 - 2020-09-03 12:03 - Norm Tovey-Walsh

Hello. Thank you for the bug report. I believe I've isolated the cause and expect to have a fix for the next patch release of 10.

#10 - 2020-09-10 15:38 - Norm Tovey-Walsh

- Status changed from In Progress to Resolved
- Fix Committed on Branch 10, trunk added

The problem has been identified and patched for the next maintenance release. Thank you!

#11 - 2020-10-28 18:57 - O'Neil Delpratt

Bug fix applied in the Saxon 10.3 maintenance release

#12 - 2020-10-28 19:13 - O'Neil Delpratt

- Status changed from Resolved to Closed
- % Done changed from 0 to 100
- Fixed in Maintenance Release 10.3 added

Files

test.jar	6.56 KB	2020-09-02	Air Quick
----------	---------	------------	-----------