

Portable Fedora Objects

Fedora objects that are intended for public distribution or migration must be careful to assure portability so that they can be ingested by a foreign repository and expected to work. Currently, the primary concern is with links or other references to resources assumed to exist locally within the fedora server.

Motivation

Links to resources within the fedora web application typically take the form `http://HOST:PORT/APP/PATH`, where `HOST` is the host on which Fedora runs, `PORT` is the port on which the server runs, and `APP` is the name of the a web application, and `PATH` is the path, and `PATH` is the rest of the path.

As an example, consider a url invoking a dissemination. Using default hosts, ports, and context names, such url might look like:

```
http://localhost:8080/fedora/get/demo:SmileyStuff/demo:Collection/list
```

In a different repository with custom host, port, and context names, the same dissemination on the same objects might look like:

```
http://example.org:9090/prod_fedora_32/get/demo:SmileyStuff/demo:Collection/list
```

Considering both URLs, we notice that the `PATH` is the only part of the url that hasn't changed.

Likewise, there is often a need to link to applications that are not in fedora, but *are* in the same servlet container. An obvious example of this is the saxon servlet, which is distributed with Fedora:

```
http://localhost:8080/saxon/SaxonServlet?...  
http://example.org:9090/saxon/SaxonServlet?...
```

In this situation, both `APP` and `PATH` have remained the same.

In fedora objects, there are sometimes needs to link to a dissemination of another object, or invoke a servlet installed alongside Fedora such as the 'saxon' servlet. Here are some examples:

In a datastream (based on demo:SmileyStuff demo object)

```
<foxml:datastream CONTROL_GROUP="E" ID="LIST" STATE="A" VERSIONABLE="true">  
  <foxml:datastreamVersion ID="LIST.0"  
    LABEL="Result of list dissemination" MIMETYPE="text/xml">  
    <foxml:contentLocation  
      REF="http://example.org:9090/prod_fedora_32/get/demo:SmileyStuff/demo:Collection/list" TYPE="URL"/>  
    </foxml:datastreamVersion>  
  </foxml:datastream>
```

In the WSDL of a Service Deployment (SDep) object (based on demo:CollectionImpl):

```
<wsdl:service name="ImageCollection">  
  <wsdl:port binding="this:ImageCollection_http" name="ImageCollection_port">  
    <http:address location="http://example.org/9090/" />  
  </wsdl:port>  
</wsdl:service>  
<wsdl:binding name="ImageCollection_http" type="this:ImageCollectionPortType">  
  <http:binding verb="GET" />  
  <wsdl:operation name="view">  
    <http:operation location="/saxon/SaxonServlet?source=(LIST)&style=(XSLT)&clear-stylesheet-cache=(CLEAR_CACHE)" />  
    ...  
  <wsdl:operation name="list">  
    <http:operation location="/prod_fedora_32/risearch?type=(TYPE)&lang=(LANG)&format=(FORMAT)&query=(QUERY)" />  
    ...
```

If fedora objects are truly written this way and ingested into a repository with a different hostname, port, or fedora app context, they will fail since the links will be broken in that environment

Solution

As a solution, Fedora uses a form of text substitution that allows local links to be represented in a portable way. Two variables are defined:

```
http://local.fedora.server/ = http://HOST:PORT/  
http://local.fedora.server/fedora/ = http://HOST:PORT/FEDORA_APP/
```

This substitution applies only to URLs for external reference datastreams (control groups R and E), and to the content of WSDL datastreams in SDep objects. At present, it does NOT occur in the content of datastreams (WSDL being the exception).

Let us re-write the FOXML examples above into a portable form.

In a datastream (based on demo:SmileyStuff demo object)

```
<foxml:datastream CONTROL_GROUP="E" ID="LIST" STATE="A" VERSIONABLE="true">  
  <foxml:datastreamVersion ID="LIST.0"  
    LABEL="Result of list dissemination" MIMETYPE="text/xml">  
    <foxml:contentLocation  
      REF="http://local.fedora.server/fedora/get/demo:SmileyStuff/demo:Collection/list" TYPE="URL"/>  
    </foxml:datastreamVersion>  
  </foxml:datastream>
```

At runtime, this will be interpreted as real url http://example.org:9090/prod_fedora_32/get/demo:SmileyStuff/demo:Collection/list or <http://localhost:8080/fedora/get/demo:SmileyStuff/demo:Collection/list>, etc depending on the local Fedora configuration.

As far as the WSDL from (based on demo:CollectionImpl), we re-write it as follows:

```
<wsdl:service name="ImageCollection">  
  <wsdl:port binding="this:ImageCollection_http" name="ImageCollection_port">  
    <http:address location="LOCAL"/>  
  </wsdl:port>  
</wsdl:service>  
<wsdl:binding name="ImageCollection_http" type="this:ImageCollectionPortType">  
  <http:binding verb="GET"/>  
  <wsdl:operation name="view">  
    <http:operation location="http://local.fedora.server/saxon/SaxonServlet?source=(LIST)&style=(XSLT)&clear-stylesheet-cache=(CLEAR_CACHE)"/>  
    ...  
  <wsdl:operation name="list">  
    <http:operation location="http://local.fedora.server/fedora/risearch?type=(TYPE)&lang=(LANG)&format=(FORMAT)&query=(QUERY)"/>  
    ...  
  </wsdl:operation>  
</wsdl:binding>
```

Notice, we made a minor structural change in order to allow the entire URL to be written in a contiguous fashion.

Use and Practice

- By default, Fedora will store objects internally in their portable form. It is only on runtime access that the URLs are translated to concrete values. Thus, changing the port or app server context of an already-installed Fedora instance with loaded objects should work fine
- Objects can be exported in portable form from an existing repository only when exported in the "Migrate" context. See [ingest & export](#) documentation for more details