

Triplestore Setup

Fuseki

Fuseki is very easy to setup

1. Download latest jena-fuseki distribution tarball from <http://www.apache.org/dist/jena/binaries/>
2. Unpack the download archive
3. Start fuseki-server

```
curl -O http://www.apache.org/dist/jena/binaries/jena-fuseki-1.0.1-distribution.tar.gz
tar xvfz jena-fuseki-1.0.1-distribution.tar.gz
cd jena-fuseki-1.0.1
./fuseki-server --update --mem /test
```

Note: The "fuseki-server" command above runs completely in-memory. No triples will be persisted. This is great for test, but not recommended for a production installation. See: http://jena.apache.org/documentation/serving_data/index.html

Fuseki should be running at <http://localhost:3030/>

If you want to change the port on which Fuseki runs, you can add the "--port" option to the run command (e.g. ./fuseki-server --port 3131 --update --mem /test).

To access the SPARQL query form, go to <http://localhost:3030/control-panel.tpl> and select your datasource (in this case /test).

- This allows you to perform SPARQL queries or updates, or upload an RDF file to replace the contents of the triplestore.

Sesame

Sesame requires a little more setup to run with the tests, since by default it uses the same port as Fedora. To setup Sesame with Tomcat running on an alternate port:

- Download Tomcat from <http://tomcat.apache.org/download-70.cgi> and unpack the downloaded archive.

```
$ curl -L -O https://dist.apache.org/repos/dist/release/tomcat/tomcat-7/v7.0.59/bin/apache-tomcat-7.0.59.tar.gz
$ tar xvfz apache-tomcat-7.0.59.tar.gz
```

- Edit the Tomcat configuration file (apache-tomcat-7.0.59/conf/server.xml) to use a different port (e.g., 8081 instead of 8080).
- Download Sesame from <http://sourceforge.net/projects/sesame/files/Sesame%202/>
- Unpack Sesame and move the Sesame WAR file into the Tomcat webapps directory, then start Tomcat.

```
$ curl -L -O http://sourceforge.net/projects/sesame/files/Sesame%202/2.7.13/openrdf-sesame-2.7.13-sdk.tar.gz
$ tar xvfz openrdf-sesame-2.7.13-sdk.tar.gz
$ cp openrdf-sesame-2.7.13/war/openrdf-sesame.war apache-tomcat-7.0.59/webapps/
$ apache-tomcat-7.0.59/bin/startup.sh
```

Sesame should now be running at: <http://localhost:8080/openrdf-sesame/system/overview.view>

Before you can load data or perform queries, you must create a repository. You can do this using OpenRDF Workbench, a companion webapp to Sesame which provides the ability to manage repositories, perform SPARQL queries, etc. To deploy this webapp, copy the OpenRDF Workbench WAR file to the Tomcat webapps directory:

```
$ cp openrdf-sesame-2.7.13/war/openrdf-workbench.war apache-tomcat-7.0.59/webapps/
```

Then you can create a new repository at <http://localhost:8081/openrdf-workbench/repositories/NONE/create>, and query the test repository at <http://localhost:8081/openrdf-workbench/repositories/test/query>.

Alternatively, you can use the Sesame console application to create a repository from the command line:

```
$ openrdf-sesame-2.7.13/bin/console.sh
> connect http://localhost:8081/openrdf-sesame.
> create native.
Repository ID [native]: test
Repository title [Native store]: test store
Triple indexes [spoc,posc]: spoc,posc
Repository created
> exit.
Disconnecting from http://localhost:8081/openrdf-sesame/
```

Notes

- You may need to set the Sesame data directory in your Tomcat7 configuration (e.g. /etc/default/tomcat7)

```
JAVA_OPTS="${JAVA_OPTS} -Dinfo.aduna.platform.appdata.basedir=/tmp/sesame"
```

Access Control

By default, Sesame will allow any connection to perform any operation, including adding and deleting triples or Sesame repositories. Since Sesame runs in a standard Java Servlet container, you can configure access control

- Tomcat supports a [variety of authentication systems](#). For this example, we'll edit the configuration file of the default authentication system that uses a plaintext XML file to store usernames and passwords. Edit `apache-tomcat-7.0.59/conf/tomcat-users.xml`, adding this line inside the `tomcat-users` element (making sure it's not in the commented-out examples):

```
<user username="user" password="password" roles="sesame-admin"/>
```

- Restart Tomcat for the username and password changes to take effect.
- When creating repositories using the Sesame console, specify the username and password when connecting:

```
> connect http://localhost:8081/openrdf-sesame user password.
```

- Edit `apache-tomcat-7.0.59/webapps/openrdf-sesame/WEB-INF/web.xml` and add the following security configuration inside the `web-app` element, replacing "test" in the `url-pattern` element with the name of your repository:

```
<login-config>
  <auth-method>BASIC</auth-method>
  <realm-name>Sesame</realm-name>
</login-config>
<security-role>
  <description>Role required to delete/update triples or change configuration</description>
  <role-name>sesame-admin</role-name>
</security-role>
<security-constraint>
  <web-resource-collection>
    <web-resource-name>Sesame System Config Restricted access</web-resource-name>
    <url-pattern>/repositories/test/*</url-pattern>
    <url-pattern>/repositories/SYSTEM/statements</url-pattern>
    <url-pattern>/repositories/SYSTEM/namespaces</url-pattern>
    <url-pattern>/repositories/SYSTEM/namespaces/*</url-pattern>
    <http-method>POST</http-method>
    <http-method>PUT</http-method>
    <http-method>DELETE</http-method>
  </web-resource-collection>
  <auth-constraint>
    <role-name>sesame-admin</role-name>
  </auth-constraint>
</security-constraint>
```

- Read operations will work without authorization, but creating/deleting Sesame repositories, or adding/delete triples, will require a username /password.

See [Basic Security with HTTP authentication](#) for more info on using container access control with Sesame, including requiring a separate account for read operations.

Documentation

- [Sesame Documentation](#)