

milestone 6 - Installing Solr and GSearch

Islandora utilizes the Solr open source search platform to enable flexible and configurable indexing and searching. Solr uses the [Lucene](#) Java search library at its core for full-text indexing and search and offers hit highlighting, faceted search, dynamic clustering, database integration, and rich document (e.g., Word, PDF) handling as additional features.

The Fedora Generic Search Service, or GSearch, is a search service installed with Fedora that allows for automatic updating of the Lucene/Solr index. GSearch relies on JMS to receive messages that are sent when Fedora objects are ingested, modified or purged. This keeps the Lucene index in sync with the Fedora repository.

How Islandora uses Solr/Lucene and Gsearch

Islandora makes it possible to use the power of Solr/Lucene for discovery. Gsearch is used as a method for keeping indexes current. When an item is ingested, the FOXML is transformed by an XSLT file stored in Gsearch into a format that can be read by Solr's schema and returned based on the request handlers in our custom solrconfig.xml.

The Islandora Solr Search module is packaged with files that will support Islandora solution packs, but these can be modified if you are familiar with Solr. Specifically the Solr schema and corresponding Gsearch XSLT are a good starting point even if you do not use the Solr Search module. Additional information about Solr is presented in Chapter 4: Search and Discovery in Islandora (The Solr Module).

While Solr is not required to run Islandora, it is recommended.

Installing Solr and GSearch

1. [Download GSearch 2.7](#), unzip and copy fedoragsearch.war to the Tomcat webapps folder (\$CATALINA_HOME/webapps).

```
# cd /opt/downloads
# wget http://downloads.sourceforge.net/fedora-commons/fedoragsearch-2.7.zip
# unzip fedoragsearch-2.7.zip
# cd fedoragsearch-2.7
# cp -v fedoragsearch.war /usr/local/fedora/tomcat/webapps
```

2. Download [Solr 4.2.0](#) and unpack.

```
# cd /opt/downloads
# wget https://archive.apache.org/dist/lucene/solr/4.2.0/solr-4.2.0.tgz
# tar -xzf solr-4.2.0.tgz
```

3. Create a directory for Solr, and copy the contents of solr-4.2.0/example/solr. These instructions presume that it will be installed at /usr/local/fedora/solr:

```
# mkdir -p /usr/local/fedora/solr
# cp -Rv solr-4.2.0/example/solr/* /usr/local/fedora/solr
```

4. Copy (and rename) solr-4.2.0.war to \$CATALINA_HOME/webapps.

```
# cp -v solr-4.2.0/dist/solr-4.2.0.war /usr/local/fedora/tomcat/webapps/solr.war
```

5. Edit the file \$FEDORA_HOME/server/config/fedora-users.xml to add an additional user:

```
<user name="fgsAdmin" password="[password]">
  <attribute name="fedoraRole">
    <value>administrator</value>
  </attribute>
</user>
```

6. Restart Fedora to deploy GSearch and Solr

```
# /usr/local/fedora/tomcat/bin/shutdown.sh
# /usr/local/fedora/tomcat/bin/startup.sh
```

This will deploy solr and fedora

7. Now we need to edit the configuration for gsearch.

Edit the file `$(FEDORA_HOME)/tomcat/webapps/fedoragsearch/FgsConfig/fgsconfig-basic-for-islandora.properties`. The entries below are of an example setup.

```
configDisplayName=configProductionSolr
gsearchBase=http://localhost:8080
gsearchAppName=fedoragsearch

gsearchUser=fgsAdmin
gsearchPass=[password]
finalConfigPath=/usr/local/fedora/tomcat/webapps/fedoragsearch/WEB-INF/classes

logFilePath=/usr/local/fedora/server/logs
logLevel=DEBUG
namesOfRepositories=FgsRepos

namesOfIndexes=FgsIndex
fedoraBase=http://localhost:8080
fedoraAppName=fedora

fedoraUser=fedoraAdmin
fedoraPass=[password]
fedoraVersion=3.8.1

objectStoreBase=/usr/local/fedora/data/objectStore
indexEngine=Solr
indexBase=http://localhost:8080/solr

indexDir=/usr/local/fedora/solr/collection1/data/index
indexingDocXslt=foxmlToSolr
```



This Gsearch properties file is extremely sensitive to whitespace. Ensure there are no extra characters/space at the end of each line!

8. Edit `$(FEDORA_HOME)/tomcat/webapps/fedoragsearch/FgsConfig/fgsconfig-basic.xml` (replace 'fgsconfig-basic.properties' with 'fgsconfig-basic-for-islandora.properties'):

```
-- <property file="fgsconfig-basic.properties" />
-- <copy file="fgsconfig-basic.properties" tofile="fgsconfig-basic- $\{\text{configDisplayName}\}$ .properties"
preserverlastmodified="true" />
-- <copy file="fgsconfig-basic.properties" tofile=" $\{\text{templateConfigPath}\}/\{\text{configDisplayName}\}/\{\text{finalConfigName}\}$ 
/fgsconfig-basic- $\{\text{configDisplayName}\}$ .properties" preserverlastmodified="true" />

++ <property file="fgsconfig-basic-for-islandora.properties" />
++ <copy file="fgsconfig-basic-for-islandora.properties" tofile="fgsconfig-basic- $\{\text{configDisplayName}\}$ .
properties" preserverlastmodified="true" />
++ <copy file="fgsconfig-basic-for-islandora.properties" tofile=" $\{\text{templateConfigPath}\}/\{\text{configDisplayName}\}$ 
/ $\{\text{finalConfigName}\}$ /fgsconfig-basic- $\{\text{configDisplayName}\}$ .properties" preserverlastmodified="true" />
```

9. Navigate to `$(FEDORA_HOME)/tomcat/webapps/fedoragsearch/FgsConfig/` and run Apache ant on fgsconfig-basic.xml (you may need to install ant)

```
# apt-get install ant
# cd /usr/local/fedora/tomcat/webapps/fedoragsearch/FgsConfig/
# ant -f fgsconfig-basic.xml
```

10. Set up solr to use the newly generated schema. Start by backing up the existing schema:

```
# mv -v /usr/local/fedora/solr/collection1/conf/schema.xml $FEDORA_HOME/solr/collection1/conf/schema.bak
```

11. Copy the new, ant-generated Solr schema into place as schema.xml

```
# cp -v /usr/local/fedora/tomcat/webapps/fedoragsearch/WEB-INF/classes/fgsconfigFinal/index/FgsIndex/conf/schema-4.2.0-for-fgs-2.7.xml $FEDORA_HOME/solr/collection1/conf/schema.xml
```

12. Create a context file for Solr at /usr/local/fedora/tomcat/conf/Catalina/localhost/solr.xml

```
<?xml version="1.0" encoding="UTF-8"?>
<Context docBase="/usr/local/fedora/tomcat/webapps/solr.war" debug="0" crossContext="true">
  <Environment name="solr/home" type="java.lang.String" value="/usr/local/fedora/solr" override="true"/>
</Context>
```

13. Restart Fedora.

```
# /usr/local/fedora/tomcat/bin/shutdown.sh
# /usr/local/fedora/tomcat/bin/startup.sh
```

Test your GSearch installation

- Go to <http://localhost:8080/fedoragsearch/rest>
- Login using fedoraAdmin
- Click on updateindex
- Click on updateIndex fromFoxmlFiles

fedora **Admin Client for Fedora Generic Search Service**

You are logged in as fedoraAdmin using configForIslandora

[updateindex](#) [gfindObjects](#) [browseIndex](#) [getRepositoryInfo](#) [getIndexInfo](#) [enduserSearch](#) (29 milliseconds)

updateIndex

Index name: Fgsindex	Inserted number of index documents: 0	Updated number of index documents: 0	Deleted number of index documents: 0	Resulting number of index documents: 334	Number of warnings (see log): 0
--------------------------------	---------------------------------------	--------------------------------------	--------------------------------------	--	---------------------------------

Foxml files path:

Pid:

Pid:

Copyright © 2010, 2011, 2012 Technical University of Denmark



If you can not access the rest page please retrace your steps. Be sure to not have extra characters or white space in the /fgsconfig-basic-for-islandora.properties files as emphasized on step 7.



GSearch and Solr may need to be further configured to index specific datastreams, depending on the enhancement modules you install. Sample GSearch and Solr configurations can be found [here](#) and [here](#).

Installing & Configuring the Solr Search Module

The Islandora Solr search module allows you to search the Solr index. The Islandora Solr Sample Configuration module provides default display profiles to the module. The module makes four new blocks available; two for search and two for display. The other block is called the *Advanced Search Block* and does fielded searches against the Solr index. Both blocks would use whatever request handler is configured in the module settings. For information on how to configure the Solr module, see [Islandora Solr Search](#).

Installation Steps:

1. Verify that fedoraGSearch and Solr are both installed and running. Go to <http://localhost:8080/solr>

The screenshot shows the Apache Solr Admin UI. On the left is a navigation menu with 'Dashboard' selected. The main content area is divided into several sections:

- Instance:** Shows start time (14 days ago), host (islandora-test.scinet.utoronto.ca), CWD (/usr/local/fedora), Instance path (/usr/local/fedora/solr/collection1), Data path (/usr/local/fedora/solr/collection1/data), and Index path (/usr/local/fedora/solr/collection1/data/index).
- Versions:** Lists solr-spec (4.2.0.2013.03.06.22.32.13), solr-impl (4.2.0 1453694 - rmuir - 2013-03-06 22:32:13), and lucene-sped (2.0).
- JVM:** Shows runtime (Oracle Corporation Java HotSpot(TM) 64-Bit Server VM), processors (1), and various JVM arguments like -Djava.io.tmpdir, -Dcatalina.home, -Dcatalina.base, etc.
- System:** Displays physical memory (49.3% used, 3.95 GB / 8.00 GB), swap space (0.0%), file descriptor count (3.6%, 146 / 4096), and JVM memory (69.7% used, 1.35 GB / 1.93 GB).

At the bottom, there are links for Documentation, Issue Tracker, IRC Channel, Community forum, and Solr Query Syntax.

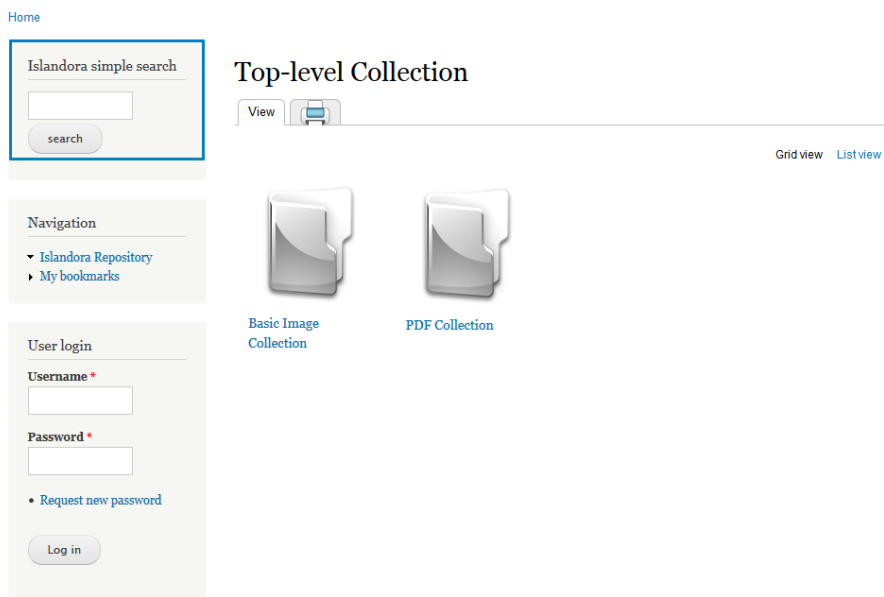
2. Solr does not automatically load existing content upon installation. Go to **Core Admin -> Reload** tab
3. Download the [Islandora Solr Search Module](#) and install as a Drupal module, as you did in the Starter Kit modules
4. Log in to your Drupal site to enable the islandora_solr_search module. Administer > Modules and enable Islandora Solr Search
5. Access the **[Structure]:[Blocks]** tab, go to the bottom of the blocks list in the **Disabled** section, find the **"+ Islandora simple search"** block and drag it up to the **Sidebar first** section.
6. Replace the original Drupal search block with the one from Islandora on the sidebar

The screenshot shows the Drupal 'Structure: Blocks' configuration page. The 'Structure' tab is selected in the top navigation. The page displays a table of blocks organized into regions:

BLOCK	REGION	OPERATIONS
Featured		
No blocks in this region		
Content		
Main page content	Content	configure
Relation Entity Collector	Content	configure
Sidebar first		
Islandora simple search	Sidebar first	configure
Navigation	Sidebar first	configure
User login	Sidebar first	configure
Sidebar second		
No blocks in this region		

The 'Islandora simple search' block is highlighted with a blue box, and its region is set to 'Sidebar first'.

7. You should now be able to see the Islandora search box in the on the sidebar.



8. To allow other users than Admin (user 1) to see the Islandora simple Search box, or search the solr index, grant them the 'Search the Solr Index' Drupal Permission at /admin/people/permissions.

PERMISSION	ANONYMOUS USER	AUTHENTICATED USER	ADMINISTRATOR
Islandora Solr			
Search the Solr index Search for objects in the repository using Solr. This permission exposes the search blocks and allows you to see search results.	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Administer Solr Administer settings for the Solr client.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Testing the Solr installation

Just enter a keyword in the Islandora search field related to some of the test objects you ingest in the Starter Kit collections. If they come up on the list everything is working as it should

! If nothing comes back please retrace your steps. Make sure you're using the Islandora search block, and not the original one from Drupal.

i From this point on you should be in good shape for the enhanced features of Islandora, by installing other solution packs in modules.