

Lucene Reindex

Overview

This page documents the use of admin tools to rebuild the Lucene index. Indexes are created based on the hibernate-search annotations on model objects such as @Indexed, @IndexedEmbedded, @Field, @ContainedIn, etc.

Prerequisites

- MySQL database must already exist (e.g. "mappingservicecb").
 - If no data exists in the database, this tool will clear the indexes.
 - Otherwise, it will index the available data.
- MySQL database connection parameters must be defined in the properties file specified by "run.config".
- The file system directory specified by the "hibernate.search.default.indexBase" in the properties file specified by run.config must exist.
- **IMPORTANT:** The Tomcat server must not be running while this is invoked, otherwise indexes can get corrupted.

Details

This tool is used to rebuild Lucene indexes for objects indexed by the system.

Indexes are written into the directory specified by the hibernate.search.default.indexBase property in the properties file specified by "run.config".

Following are some details about the implementation of this tool:

Mojo: LuceneReindexMojo.java (in admin/mojo/src/main/resources/java/org/ihtsdot/otf/mapping/mojo)

Project: admin/lucene

Configuration Parameters

- run.config - the standard configuration file specified as a -D parameter
- indexed.objects - optional parameter that allows a comma-separated list of objects to reindex (e.g. MapRecordJpa). If omitted, all indexes are rebuilt.

Samples

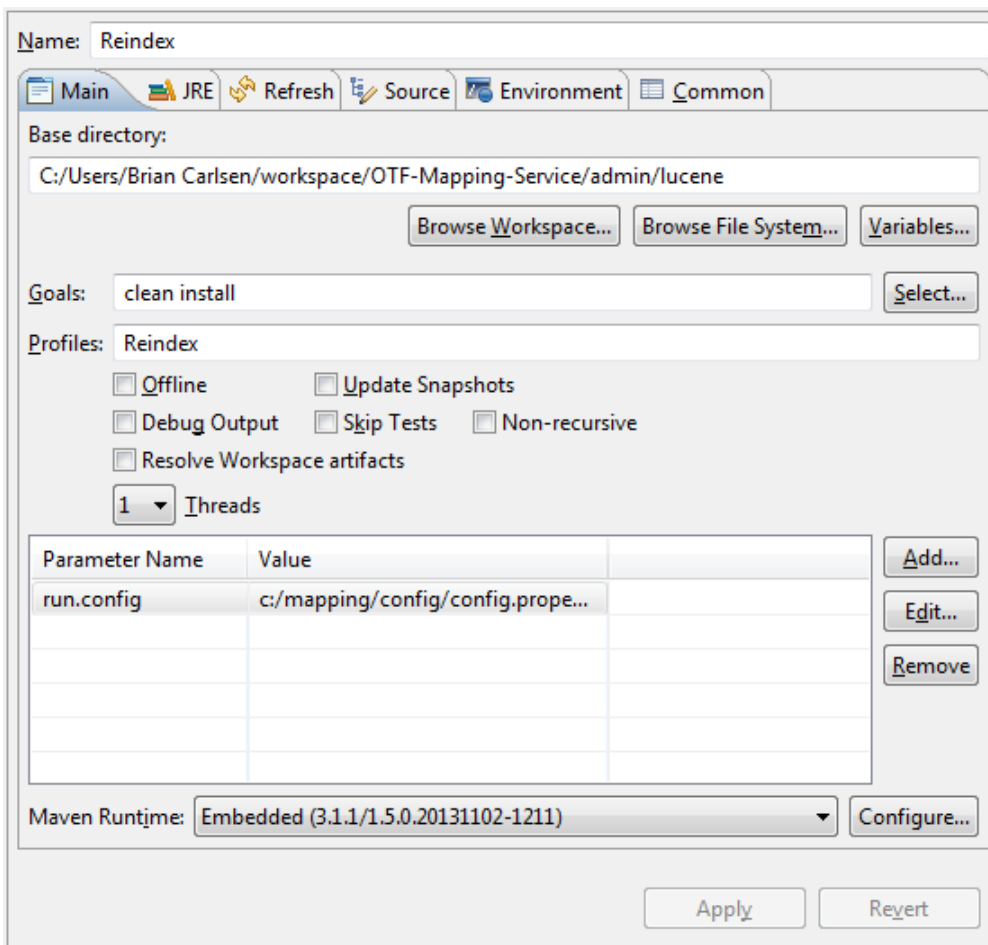
Sample command line call of the admin tool to rebuild all indexes:

```
% cd ~/code/admin/lucene
% mvn install -Preindex -Drun.config=/home/ihtsdot/config/config-load.properties
```

Sample command line call of the admin tool to reindex only certain data types:

```
# Re-indexes only map records and tracking records
% cd ~/code/admin/lucene
% mvn install -Preindex -Drun.config=/home/ihtsdot/config/config-load.properties -Dindexed.objects=MapRecordJpa,TrackingRecordJpa
```

Sample eclipse run configuration to rebuild all indexes:



Troubleshooting

Troubleshooting

Sometimes issues arise with the indexes (in particular when running admin tools while the Tomcat server is also running). Occasionally it is necessary to completely reindex all objects from scratch. When this is needed, the recommended approach is to remove the index files from the index and then run this mojo. For example, to guarantee a clean index rebuild for map records, you could do the following.

```
# stop the server
% service tomcat7 stop

# remove the index
% /bin/rm -rf /var/lib/tomcat7/indexes/lucene/indexes/*MapRecordJpa

# rebuild the index
% cd ~/code/admin/lucene
% mvn install -Preindex -Drun.config=/home/ihtsdo/config/config-load.properties -Dindexed.
objects=MapRecordJpa

# start the server
% service tomcat7 start
```

References/Links

- n/a

