

Configuring Low Level Storage

Introduction

Often abbreviated "LLStore", the *Low Level Storage Interface* is a critical component of Fedora. It stores and provides access to the authoritative copy of all digital object XML (FOXML) and datastreams managed by a Fedora repository.

LLStore is an internal Java interface and is not intended to be accessed directly by user applications. Instead, applications are expected to interact with Fedora through the web-based repository APIs.

Fedora's default LLStore module stores digital object XML and datastreams as individual files in a regular filesystem. However, there are several other options that can be configured.

General Configuration

If you don't want to use the default implementation or configuration options, you must configure the LLStore module before starting Fedora for the first time. Configuration is done within the `server/config/fedora.fcfg` file, by modifying the **class** and **param** values as appropriate for the LLStore implementation you are plugging in:

```
<module role="fedora.server.storage.lowlevel.ILowlevelStorage"
  class="org.example.SomeLLStoreModule">
  <param name="someParam" value="someValue"/>
  <!-- etc -->
</module>
```

Consult the documentation for each plug-in listed below for the specific class name and configuration options supported. Depending on the plug-in you are installing, you may also need to add one or more .jar files to the Fedora webapp's classpath (e.g. WEB-INF/lib).

Plug-Ins Provided by Fedora Commons

Akubra Low Level Storage (New)

[Akubra](#) is a newly-developed, generic "blob" storage API. The [AkubraLowlevelStorage](#) plug-in is included with Fedora 3.2 and acts as a bridge between Fedora and Akubra, so that Akubra implementations can be used to support low-level storage in Fedora.

Because it offers an improved file storage abstraction and is pluggable itself, we plan to use Akubra as the default low-level storage option for future releases of Fedora. However, it is also a very new API; we would like to gain more experience with it and hear more user feedback before making it the default option.

To use this plug-in, replace the existing LowlevelStorage module in `fedora.fcfg` with the following:

```
<module role="fedora.server.storage.lowlevel.ILowlevelStorage"
  class="fedora.server.storage.lowlevel.akubra.AkubraLowlevelStorageModule" />
```

Then modify the `akubra-llstore.xml` file as appropriate. This is a Spring bean configuration file and is where you tell Fedora which Akubra BlobStore implementation to use and how it should be configured.

Default Akubra-LLStore Configuration

The default Akubra-LLStore configuration is similar to the Filesystem Low Level Storage module traditionally used by Fedora in that it stores object xml and datastreams as individual files on your local filesystem. However, it differs in a couple important ways:

- It does not require the use of database tables to "look up" the path to each file.
- It stores files in a deterministic location based on an md5 hash of the unique id of each file.

Customizing the Akubra-LLStore Configuration

By editing the [default akubra-llstore.xml](#) file in a couple simple ways, you can control:

- The *base directory* of the object xml and datastream storage areas. Unless you are just testing out Fedora with Akubra, you should definitely change these values to a permanent storage location. Look for `/tmp/datastreamStore` and `/tmp/objectStore` in the file and change the values as desired.
- The *shape of the directory tree* in which to store the files. You control this by changing the "###" values in the configuration file. The default values of "###" ensure that there are at most 256 directories in each store, all at the top level, and files will be distributed fairly evenly among them. For more information on what "###" means and how it controls the directory structure, see [HashPathIdMapper](#) in the Fedora 3.2 Javadocs.

For more extensive information on Akubra, please refer to the [Akubra Wiki Space](#).

Filesystem Low Level Storage (Mature)

This is Fedora's default storage option and requires no additional setup to use. It comes pre-configured with your Fedora installation to store all objects in `$FEDORA_HOME/data/objects/`, and all datastreams in `$FEDORA_HOME/data/datastreams/`. Paths are allocated based on the date the item was first created. For example, a datastream created on May 8th, 2009 might be located in the `2009/0508/20/48/` directory.

To change these locations, consult the `fedora.fcfg` file that came with your Fedora installation and change the indicated param values as desired.

S3 Low Level Storage (Experimental)

This plugin stores all object XML and datastream content on Amazon's Simple Storage System (S3) and serves as a good example of how a `LowlevelStorage` implementation can be built. However, we do not recommend the use of this plugin in production environments because:

- It does not support rebuilds (see "Rebuilder Support" below)
- We plan to replace it with an Akubra S3 implementation in the near future.

For an overview of how this plugin works and how it can be configured to work with Fedora, [please see this document](#).

The S3 Low Level Storage is not distributed in binary form. You may obtain it from our legacy Subversion repository via:

```
svn co https://svn.fedora-commons.svn.sourceforge.net/svnroot/fedora-commons/incubator/AmazonS3Storage
```

Third-Party Plug-Ins

iRODs Low Level Storage

This plugin was developed by the [DICE Group](#) and allows Fedora to use [iRODs](#) to store digital objects and datastreams. It is based on the SRB plugin developed by the DART Project (see below).

For more information on this plugin, please visit <https://www.irods.org/index.php/Fedora>

SRB Low Level Storage

This plugin was developed by the [DART Project](#) and allows Fedora to use the [Storage Resource Broker \(SRB\)](#) to store digital objects and datastreams.

For more information on this plugin, please visit <http://www.itee.uq.edu.au/~eresearch/projects/dart/outcomes/FedoraDB.php>

Sun Honeycomb Low Level Storage

This plugin was developed by [Sun Microsystems](#) and allows Fedora to use the [StorageTek 5800 System](#) to store digital objects and datastreams.

For more information on this plugin, please visit <http://opensolaris.org/os/project/honeycomb/>

Rebuilder Support

Fedora's Rebuild Utility ([fedora-rebuild](#)) provides repository administrators with an automated way to reconstitute Fedora's higher-level indexes (the SQL database and the Resource Index) when upgrading Fedora, migrating to another SQL database, or recovering from inconsistencies.

Historically, the only Low Level Storage implementation that supported this utility was the Filesystem Low Level Storage module that came bundled with Fedora. As of Fedora 3.2, the rebuilder now works with any `LLStore` implementation that supports the new [IListable interface](#).

Currently, the only `LLStore` implementations that support this interface (and thus can be used with the rebuild utility) are the Akubra and Filesystem `LLStore` modules released with Fedora 3.2.