

Ldap Synchronization Connector (LSC) - User Guide

Jonathan Clarke

18/01/2010

Contents

| | |
|---|----------|
| Introduction | 2 |
| Basic principles | 3 |
| Entries and attributes | 3 |
| Pivot attributes | 3 |
| One-to-one entry matching | 3 |
| One-way only | 3 |
| Four operations available | 4 |
| Transformations | 4 |
| Obtaining and installing LSC | 4 |
| Choosing the right version | 4 |
| Version numbering | 4 |
| Current version | 5 |
| Downloading and installing the software | 5 |
| Dependencies | 5 |
| Downloading | 5 |
| Installing | 6 |

| | |
|---|----------|
| Configuring LSC | 6 |
| Source connector configuration | 6 |
| Destination connector configuration | 6 |
| Data transformation configuration | 6 |
| Logging configuration | 6 |
| Running LSC | 6 |
| Integrating LSC to a running system (including best practices) | 7 |
| Troubleshooting and getting help | 7 |

The latest version of this document is available online at <http://lsc-project.org/user-guide/>. You can also download a PDF version.

This guide covers versions **1.2.*** of LSC. Other versions are available.

Introduction

Ldap Synchronization Connector, or LSC, is a tool for synchronizing data repositories, with a particular focus on identity information.

It can read from a variety of different sources, perform limitless transformations of data on-the-fly and write the results to a LDAP directory. The tool is designed to be fast, and extensible in all areas.

It is an open source project, released under the BSD license.

This document is a hands-on guide to all aspects of using LSC:

1. Basic principles
2. Obtaining and installing LSC
3. Configuring LSC
4. Running LSC
5. Integrating LSC to a running system (including best practices)
6. Troubleshooting and getting help

This guide aims to be exhaustive. As such, the reader is not expected to read it from cover to cover, but is encouraged to skip to the sections he's interested in. A quick-start guide is also available for the impatient.

Basic principles

Entries and attributes

All data is split up into two elements:

- **Entries**, which are containers. It can be considered equivalent to an *account* (“this account exists or it does not”)
- **Attributes**, which qualify the entry. These are also commonly called fields, elements or keys (“this account is called jdoe”)

Pivot attributes

To synchronize data, LSC must be able to match up entries in the source and in the destination. This is an important prerequisite to synchronizing your data: a **unique identifier** must be available in all entries and in both your source and your destination.

This unique identifier may be a combination of several fields (for example: surname and first name).

Commonly used pivot attributes include email addresses, or serial numbers (such as UUIDs).

One-to-one entry matching

LSC will try and match every entry in the source to one, and only one entry in the destination. Matching is done by searching on the pivot attributes.

Synchronizing several entries to one, or one entry to several could be achieved by implementing multiple synchronizations, but is a hack.

One-way only

LSC synchronizations are **one-way only**: from a source to a destination.

It usually only makes sense to copy data in one direction. However, if you need to copy a different set of data back, you can simply implement a separate synchronization.

Four operations available

For every entry, LSC may perform one of the following operations (or nothing):

1. Add a new entry (if an entry from the source was not found in the destination)
2. Delete an old entry (if an entry in the destination was not found in the source)
3. Update an existing entry (if attributes vary between the destination and the calculated transformations from the source)
4. Rename an existing entry (if the LDAP DN should be changed)

Please note that one and only one of these operations is performed on each run. So, if an entry is renamed, but some of its attributes also need updating, this will only happen over two runs.

Transformations

... Work in progress ...

Obtaining and installing LSC

Choosing the right version

Several versions of LSC are available. As with many Open Source projects, the main development focus is on one branch, while maintenance is performed on the previous branch. Choosing the right version is not always easy when starting out with new software. This section offers guidance.

Version numbering

A full version number looks something like **1.1.0**. LSC version numbers are in two parts:

- The first two digits represent the **major** branch number, for example **1.1**.
- The third digit represents a **minor** release within that branch, similar to a patch level.

A minor update (patch level increase, such as 1.1.0 to 1.1.1) never changes functionality, and is only released to fix bugs.

Current version

At the time of writing, the **1.2** branch is the most recent version available, and is fully supported.

The **1.1** branch is no longer actively developed, and will no longer be supported in the near future. The latest and most stable versions of LSC are always shown on the project's website at <http://lsc-project.org/>.

Development work is currently transitioning between improving this branch, **1.2**, and implementing new features for future versions.

Downloading and installing the software

Dependencies

To run LSC you need Sun's Java Virtual Machine (JVM), version 1.5 or higher. Download and install Sun's Java Runtime Edition (JRE), if you don't already have it installed. Please note that it may be simpler to use a pre-packaged version from your operating system.

All other dependencies required by LSC are included in the download.

Downloading

In the past, two separate components were required to use LSC (*lsc-core* and *lsc-sample*). This is no longer the case. LSC is distributed like most programs, as an archive to download and install.

There are several formats available. They all lead to the same result, just choose your preference from the following:

- Graphical installer based in a JAR file. After download, run this with `java -jar filename`.
- Compressed archive in a gzipped tarball.
- Compressed archive in a zip archive.

See all downloads at <http://tools.lsc-project.org/files>.

Note: .bz2 files can be decompressed using 7-Zip on Windows.

Installing

Depending on the packaging you chose, either run the graphical installer as a Java executable (`java -jar filename`), or extract the compressed archive to the location of your choice.

In either case, LSC will be installed in a self-contained directory, including all dependencies required to run.

Once installed, you should see the following folders in the installation directory:

- *bin* (programs to launch and interact with LSC)
- *etc* (configuration files)
- *lib* (software dependencies)

Configuring LSC

There are several different aspects to configuring a connector using LSC:

1. Source connector configuration
2. Destination connector configuration
3. Data transformation configuration
4. Logging configuration

Within the installation directory, all configuration files are located in the **etc/** directory. Most configuration is done in the **lsc.properties** file, except for logging which is setup in **logback.xml**. Each source connector may use additional configuration files - such as is the case for the database source connector used in **db2ldap** connectors.

Source connector configuration

Destination connector configuration

Data transformation configuration

Logging configuration

Running LSC

... Work in progress ...

Command line options. Logs.

Integrating LSC to a running system (including best practices)

... Work in progress ...

Running make-lsc-archive script. Cron and logrotate integration. Log integration. Nagios supervision.

Troubleshooting and getting help

... Work in progress ...

IRC, Mailing lists, common errors...