

Upgrading From 3.x to 4.x

These instructions are valid for any of the following upgrade paths:



- Upgrading 3.0, 3.1, or 3.2 to 4.0 or 4.1

For more information about specific fixes released in each 4.x version, please refer to the appropriate release notes:

- [DSpace Release 4.0 Status](#)

In the notes below `[dSPACE]` refers to the install directory for your existing DSpace installation, and `[dSPACE-source]` to the source directory for DSpace 4.0. Whenever you see these path references, be sure to replace them with the actual path names on your local system.

Backup your DSpace

Before you start your upgrade, it is strongly recommended that you create a backup of your DSpace instance. Backups are easy to recover from; a botched install/upgrade is very difficult if not impossible to recover from. The DSpace specific things to backup are: configs, source code modifications, database, and assetstore. On your server that runs DSpace, you might additionally consider checking on your cron/scheduled tasks, servlet container, and database.

Make a complete backup of your system, including:

- Database: Make a snapshot/dump of the database. For the PostgreSQL database use Postgres' `pg_dump` command. For example:

```
pg_dump -U [database-user] -f [backup-file-location] [database-name]
```

- Assetstore: Backup the directory (`[dSPACE]/assetstore` by default, and any other assetstores configured in the `[dSPACE]/config/dSPACE.cfg` "assetstore.dir" and "assetstore.dir.#" settings)
- Configuration: Backup the entire directory content of `[dSPACE]/config`.
- Customizations: If you have custom code, such as themes, modifications, or custom scripts, you will want to back them up to a safe location.

Update Prerequisite Software (as necessary)

DSpace 4.0 requires the following versions of prerequisite software (see [Prerequisite Software](#) section of "Installing DSpace" for more details):

- Java 7 (Oracle or OpenJDK)
- Maven 3.x or above
- Database
 - PostgreSQL 8.4 to 9.1, OR
 - Oracle 10g or above
- Tomcat 7 or above

Upgrade Steps

1. **Download DSpace 4.0:** Either download DSpace 4.0 from [DSpace.org](#) or check it out directly from the [Github repository](#). If you downloaded DSpace do not unpack it on top of your existing installation. Refer to [Installation Instructions, Step 3](#) for unpacking directives.
2. **Merge any customizations (if needed).** If you have made any local customizations to your DSpace installation they *may* need to be migrated over to the new DSpace. For minor upgrades, there are often fewer changes, but it is still worth checking to see if any effect your customizations. Customizations are typically housed in one of the following places:
 - JSPUI modifications: `[dSPACE-source]/dSPACE/modules/jspui/src/main/webapp/`
 - XMLUI modifications: `[dSPACE-source]/dSPACE/modules/xmlui/src/main/webapp/`
 - Config modifications: `[dSPACE]/config`
3. **Edit the build.properties file (if needed)** (`[dSPACE-source]/build.properties`). Any settings changed in this `build.properties` file are automatically copied over to the final `dSPACE.cfg` file during the "Build DSpace" process (in the next step). For more information on the `build.properties` file, see "[The build.properties Configuration Properties File](#)" section of the [Configuration Reference](#) documentation.
4. **Build DSpace.** Run the following commands to compile DSpace :

```
cd [dSPACE-source]/dSPACE/  
mvn -U clean package
```

In DSpace 4.0, the above command must be run from `[dSPACE-source]`



In the DSpace 4.0 release, the above "mvn -U clean package" command **must** be run from the root source directory (i.e. `[dSPACE-source]`), otherwise you will receive build errors. This was a small (but annoying) bug in our Maven build process, which is fixed in the 4.1 release (see [DS-1867](#))

You will find the result in `[dSPACE-source]/dSPACE/target/dSPACE-[version]-build.dir`. Inside this directory is the compiled binary distribution of DSpace. Before rebuilding DSpace ('package'), the above command will clean out any previously compiled code ('clean') and ensure that your local DSpace JAR files are updated from the remote maven repository.

5. **Stop Tomcat.** Take down your servlet container. For Tomcat, use the `$CATALINA_HOME/shutdown.sh` script. (Many Unix-based installations will have a startup/shutdown script in the `/etc/init.d` or `/etc/rc.d` directories.)
6. **Update DSpace.**
 - a. Update the DSpace installed directory with the new code and libraries. Issue the following commands:

```
cd [dspace-source]/dspace/target/dspace-[version]-build.dir
ant update
```

- b. The database schema has changed in 4.0. So, you will need to update your existing DSpace 3.x database. Please use the appropriate command and SQL script to update your database:

- PostgreSQL:
 - *Upgrade database to 4.0:* `psql --user [dspace-dbms-user] -f [dspace-source]/dspace/etc/postgres/database_schema_3-4.sql [dspace-database]`
You should be prompted for the database password.
 - *Upgrade database to 4.7:* `psql --user [dspace-dbms-user] -f [dspace-source]/dspace/etc/postgres/database_schema_4-47.sql [dspace-database]`
You should be prompted for the database password.
 - *Upgrade database to 4.8:* `psql --user [dspace-dbms-user] -f [dspace-source]/dspace/etc/postgres/database_schema_4-48.sql [dspace-database]`
You should be prompted for the database password.
- Oracle:
 - *Upgrade database to 4.0:* `sqlplus [dspace-dbms-user]/[database password] [dspace-source]/dspace/etc/oracle/database_schema_3-4.sql`
 - *Upgrade database to 4.7:* `sqlplus [dspace-dbms-user]/[database password] [dspace-source]/dspace/etc/oracle/database_schema_4-47.sql`
 - *Upgrade database to 4.8:* `sqlplus [dspace-dbms-user]/[database password] [dspace-source]/dspace/etc/oracle/database_schema_4-48.sql`

[dspace-dbms-user] will be the value of `db.username` in `config/dspace.cfg`. The database password will be the value of `db.password`. [dspace-database] will be the part of `db.url` following the last slash.

7. **Update your DSpace Configurations.** You should review your configuration for new and changed configurations in DSpace 4.0. In the specific case of `dspace.cfg` it is recommended to start with a fresh copy of the file from the *new* version and copy your site-specific settings from the old file. Read the new file carefully to see if you need (or want) other alterations. **Please notice** that the default search and browse support has changed from the old Lucene/DBMS-based method to [Discovery](#).
8. **Deploy Web Applications.** If necessary, copy the web applications files from your `[dspace]/webapps` directory to the subdirectory of your servlet container (e.g. Tomcat):

```
cp -R [dspace]/webapps/* [tomcat]/webapps/
```

See [the installation guide](#) for full details.

9. **Restart servlet container.** Now restart your Tomcat/Jetty/Resin server program and test out the upgrade.
10. **Upgrade Solr Indexes.** After upgrading DSpace, you should optimize your Solr indexes to ensure that they are in the current format. Solr can upgrade recent formats to its current one. Doing this whenever you upgrade DSpace should keep the index formats current. You can optimize your indexes with a command line tool such as `wget` (you have to run it locally from your DSpace server and use the correct port where Solr runs):

Optimizing Solr indexes

```
wget http://localhost:8080/solr/<core>/update?optimize=true
```

Depending on your set-up, some or all of the following cores may exist: `search`, `statistics`, `oai`. To see what solr cores exist in your set-up, look for the directories in `[dspace]/solr` - each directory corresponds to a solr core.

If you have been through several DSpace upgrades and have not done this, there is a chance that your indexes are in a format too old for the most recent Solr to convert:

Format error



Caused by: `org.apache.lucene.index.IndexFormatTooOldException: Format version is not supported (resource: segment _386q in resource ChecksumIndexInput(MMapIndexInput(path="/space/dspace/solr/statistics/data/index/segments_37m6"))): 2.x`. This version of Lucene only supports indexes created with release 3.0 and later.

You may be able to use your older DSpace installation **before** upgrading it, to upgrade your indexes enough that a second optimization after upgrading DSpace will succeed. You can also use an external tool with its own version of the required libraries, such as [Luke](#). You will need a version of Lucene (the indexing library used by Solr) which can convert the earlier version into the later one. So, for example, you could use the Solr webapp. in DSpace 1.8.x or 3.x, or Luke 3.5, to upgrade version 2 indexes to version 3. Once you have version 3 index files, you should be able to upgrade to DSpace 4.x, and the next optimization you do will upgrade the indexes again to version 4.

11. **Refresh Browse and Search Indexes.** DSpace 4 relies on SOLR based Discovery for both search and browse purposes. To update the Discovery indexes, run the following command from your DSpace install directory as the `dspace` user.

```
[dspace]/bin/dspace index-discovery -f
```

12. **Check your cron / Task Scheduler jobs.** The index maintenance commands' names have changed to make them clearer. You will need to update your scripts.