Upgrading From 1.1.x to 1.2

This document refers to the install directory for your existing DSpace installation as `[dspace]`, and to the source directory for DSpace 1.2 as *[dspace-1.2-source]*. Whenever you see these path references below, be sure to replace them with the actual path names on your local system.

The process for upgrading to 1.2 from either 1.1 or 1.1.1 is the same. If you are running DSpace 1.0 or 1.0.1, you need to follow the instructions for Upgrading From 1.0.1 to 1.1 before following these instructions.

Note also that if you've substantially modified DSpace, these instructions apply to an unmodified 1.1.1 DSpace instance, and you'll need to adapt the process to any modifications you've made.

Upgrade Steps

1. Step one is, of course, to **back up all your data** before proceeding!! Include all of the contents of `[dspace]` and the PostgreSQL database in your backup.
2. Get the new DSpace 1.2 source code from the DSpace page on SourceForge and unpack it somewhere. Do not unpack it on top of your existing installation!!
3. Copy the required Java libraries that we couldn't include in the bundle to the source tree. For example:

```
cd  [dspace]/lib
cp activation.jar servlet.jar mail.jar
        [dspace-1.2-source]/lib
```

4. Stop Tomcat (or other servlet container.)
5. It's a good idea to upgrade all of the various third-party tools that DSpace uses to their latest versions:
    *   Java (note that now version 1.4.0 or later is *required*)
    *   Tomcat (Any version after 4.0 will work; symbolic links are no longer an issue)
    *   PostgreSQL (don't forget to build/download an updated JDBC driver .jar file! Also, **back up the database** first.)
    *   Ant
6. You need to add the following new parameters to your *[dspace]/dspace.cfg*:

```
##### Media Filter settings #####
# maximum width and height of generated thumbnails
thumbnail.maxwidth = 80
thumbnail.maxheight = 80
```

There are one or two other, optional extra parameters (for controlling the pool of database connections). See the version history for details. If you leave them out, defaults will be used.Also, to avoid future confusion, you might like to **remove** the following property, which is no longer required:

```
config.template.oai-web.xml =
        [dspace]/oai/WEB-INF/web.xml
```

7. The layout of the installation directory (i.e. the structure of the contents of `[dspace]`) has changed somewhat since 1.1.1. First up, your 'localized' JSPs (those in jsp/local) now need to be maintained in the *source* directory. So make a copy of them now! Once you've done that, you can remove *[dspace]/jsp* and *[dspace]/oai*, these are no longer used. (.war Web application archive files are used instead). Also, if you're using the same version of Tomcat as before, you need to **remove the lines from Tomcat's conf/server.xml file that enable symbolic links for DSpace.** These are the *<Context>* elements you added to get DSpace 1.1.1 working, looking something like this:

```
<Context path="/dspace" docBase="dspace" debug="0" reloadable="true"
        crossContext="true">
    <Resources className="org.apache.naming.resources.FileDirContext"
        allowLinking="true" />
</Context>
```

Be sure to remove the <Context> elements for both the Web UI and the OAI Web applications.

8. Build and install the updated DSpace 1.2 code. Go to the DSpace 1.2 source directory, and run:

```
ant -Dconfig= [dspace]/config/dspace.cfg update
```

9. Copy the new config files in *config* to your installation, e.g.:

```
cp  [dspace-1.2-source]/config/news-*
        [dspace-1.2-source]/config/mediafilter.cfg
        [dspace-1.2-source]/config/dc2mods.cfg
        [dspace]/config
```

10. You'll need to make some changes to the database schema in your PostgreSQL database. *[dspace-1.2-source]/etc /database_schema_ 11-12.sql* contains the SQL commands to achieve this. If you've modified the schema locally, you may need to check over this and make alterations. To apply the changes, go to the source directory, and run:

```
psql -f etc/database_schema_11-12.sql [DSpace database name] -h
        localhost
```

11. A tool supplied with the DSpace 1.2 codebase will then update the actual data in the relational database. Run it using:

```
[dspace]/bin/dsrun
        org.dspace.administer.Upgrade11To12
```

12. Then rebuild the search indexes:

```
[dspace]/bin/index-all
```

13. Delete the existing symlinks from your servlet container's (e.g. Tomcat's) *webapp* sub-directory. Copy the *.war* Web application files in *[dspace-1.2-source]/build* to the *webapps* sub-directory of your servlet container (e.g. Tomcat). e.g.:

```
cp  [dspace-1.2-source]/build/*.war
        [tomcat]/webapps
```

14. Restart Tomcat.
15. To get image thumbnails generated and full-text extracted for indexing automatically, you need to set up a 'cron' job, for example one like this:

```
# Run the media filter at 02:00 every day
0 2 * * *  [dspace]/bin/filter-media
```

You might also wish to run it now to generate thumbnails and index full text for the content already in your system.

16. **Note 1**: This update process has effectively 'touched' all of your items. Although the dates in the Dublin Core metadata won't have changed (accession date and so forth), the 'last modified' date in the database for each will have been changed. This means the e-mail subscription tool may be confused, thinking that all items in the archive have been deposited that day, and could thus send a rather long email to lots of subscribers. So, it is recommended that you **turn off the e-mail subscription feature for the next day**, by commenting out the relevant line in DSpace's cron job, and then re-activating it the next day. Say you performed the update on 08-June-2004 (UTC), and your e-mail subscription cron job runs at 4am (UTC). When the subscription tool runs at 4am on 09-June-2004, it will find that everything in the system has a modification date in 08-June-2004, and accordingly send out huge emails. So, immediately after the update, you would edit DSpace's 'crontab' and comment out the */dspace/bin/subs-daily* line. Then, after 4am on 09-June-2004 you'd 'un-comment' it out, so that things proceed normally. Of course this means, any *real* new deposits on 08-June-2004 won't get e-mailed, however if you're updating the system it's likely to be down for some time so this shouldn't be a big problem.

17. **Note 2:** After consultation with the OAI community, various OAI-PMH changes have occurred:
    - The OAI-PMH identifiers have changed (they're now of the form *oai:hostname:handle* as opposed to just Handles)
    - The set structure has changed, due to the new sub-communities feature.
    - The default base URL has changed
    - As noted in note 1, every item has been 'touched' and will need re-harvesting. The above means that, if already registered and harvested, you will need to re-register your repository, effectively as a 'new' OAI-PMH data provider. You should also consider posting an announcement to the OAI implementers e-mail list so that harvesters know to update their systems. Also note that your site may, over the next few days, take quite a big hit from OAI-PMH harvesters. The resumption token support should alleviate this a little, but you might want to temporarily whack up the database connection pool parameters in *[dspace]/config /dspace.cfg*. See the *dspace.cfg* distributed with the source code to see what these parameters are and how to use them. (You need to stop and restart Tomcat after changing them.)I realize this is not ideal; for discussion as to the reasons behind this please see relevant posts to the OAI community: post one, post two. If you really can't live with updating the base URL like this, you can fairly easily have thing proceed more-or-less as they are, by doing the following:
    - Change the value of *OAI_ID_PREFIX* at the top of the *org.dspace.app.oai.DSpaceOAICatalog* class to *hdl:*
    - Change the servlet mapping for the *OAIHandler* servlet back to */* (from */request*)
    - Rebuild and deploy _oai.war_However, note that in this case, all the records will be re-harvested by harvesters anyway, so you still need to brace for the associated DB activity; also note that the set spec changes may not be picked up by some harvesters. It's recommended you read the above-linked mailing list posts to understand why the change was made.
    Now, you should be finished!