

Installation

DSpace System Documentation: Installation

- 1 [For the Impatient](#)
- 2 [Prerequisite Software](#)
 - 2.1 [UNIX-like OS or Microsoft Windows](#)
 - 2.2 [Oracle Java JDK 6 or later \(standard SDK is fine, you don't need J2EE\)](#)
 - 2.3 [Apache Maven 2.2.x \(Java build tool\)](#)
 - 2.3.1 [Configuring a Proxy](#)
 - 2.4 [Apache Ant 1.7 or later \(Java build tool\)](#)
 - 2.5 [Relational Database: \(PostgreSQL or Oracle\).](#)
 - 2.6 [Servlet Engine: \(Apache Tomcat 5.5 or 6, Jetty, Caucho Resin or equivalent\).](#)
 - 2.7 [Perl \(only required for \[dspace\]/bin/dspace-info.pl\)](#)
- 3 [Installation Instructions](#)
 - 3.1 [Overview of Install Options](#)
 - 3.2 [Overview of DSpace Directories](#)
 - 3.3 [Installation](#)
- 4 [Advanced Installation](#)
 - 4.1 ['cron' Jobs](#)
 - 4.2 [Multilingual Installation](#)
 - 4.3 [DSpace over HTTPS](#)
 - 4.3.1 [To enable the HTTPS support in Tomcat 5.0:](#)
 - 4.3.2 [To use SSL on Apache HTTPD with mod_jk:](#)
 - 4.4 [The Handle Server](#)
 - 4.4.1 [Updating Existing Handle Prefixes](#)
 - 4.5 [Google and HTML sitemaps](#)
 - 4.6 [DSpace Statistics](#)
- 5 [Windows Installation](#)
 - 5.1 [Pre-requisite Software](#)
 - 5.2 [Installation Steps](#)
- 6 [Checking Your Installation](#)
- 7 [Known Bugs](#)
- 8 [Common Problems](#)

For the Impatient

Since some users might want to get their test version up and running as fast as possible, offered below is an *unsupported* outline of getting DSpace to run quickly in a Unix-based environment.

Only experienced unix admins should even attempt the following without going to the detailed [Installation Instructions](#)

```
useradd -m dspace
gunzip -c dspace-1.x-src-release.tar.gz | tar -xf -
createuser -U postgres -d -A -P dspace
createdb -U dspace -E UNICODE dspace
cd [dspace-source]/dspace/config
vi dspace.cfg
mkdir [dspace]
chown dspace [dspace]
su - dspace
cd [dspace-source]/dspace
mvn package
cd [dspace-source]/dspace/target/dspace-<version>-build.dir
ant fresh_install
cp -r [dspace]/webapps/* [tomcat]/webapps
/etc/init.d/tomcat start
[dspace]/bin/dspace create-administrator
```

Prerequisite Software

The list below describes the third-party components and tools you'll need to run a DSpace server. These are just guidelines. Since DSpace is built on open source, standards-based tools, there are numerous other possibilities and setups.

Also, please note that the configuration and installation guidelines relating to a particular tool below are here for convenience. You should refer to the documentation for each individual component for complete and up-to-date details. Many of the tools are updated on a frequent basis, and the guidelines below may become out of date.

UNIX-like OS or Microsoft Windows

- UNIX-like OS (Linux, HP/UX, Mac OSX, etc.) : Many distributions of Linux/Unix come with some of the dependencies below pre-installed or easily installed via updates, you should consult your particular distributions documentation or local system administrators to determine what is already available.
- Microsoft Windows: After verifying all prerequisites below, see the [Windows Installation](#) section for Windows tailored instructions

Oracle Java JDK 6 or later (standard SDK is fine, you don't need J2EE)

DSpace now requires **Oracle** Java 6 or greater because of usage of new language capabilities introduced in 5 and 6 that make coding easier and cleaner.

Java can be downloaded from the following location: <http://java.sun.com/javase/downloads/index.jsp>

Only Oracle's Java has been tested with each release and is known to work correctly. Other flavors of Java may pose problems.

Apache Maven 2.2.x (Java build tool)

Known issue with Maven 2.0.x and Maven 3.x and DSpace 1.7.x



DSpace 1.7.x does not build properly when using Maven 2.0.x or Maven 3.x. This is a known issue. **The quick fix is to use Maven 2.2.x.** More information on this issue can be found in the following JIRA issue: [DS-788](#).

Maven is necessary in the first stage of the build process to assemble the installation package for your DSpace instance. It gives you the flexibility to customize DSpace using the existing Maven projects found in the `[dspace-source]/dspace/modules` directory or by adding in your own Maven project to build the installation package for DSpace, and apply any custom interface "overlay" changes.

Maven can be downloaded from the following location: <http://maven.apache.org/download.html>

Configuring a Proxy

You can configure a proxy to use for some or all of your HTTP requests in Maven 2.0. The username and password are only required if your proxy requires basic authentication (note that later releases may support storing your passwords in a secured keystore, in the mean time, please ensure your `settings.xml` file (usually `${user.home}/.m2/settings.xml`) is secured with permissions appropriate for your operating system).

Example:

```
<settings>
.
.
<proxies>
  <proxy>
    <active>true</active>
    <protocol>http</protocol>
    <host>proxy.somewhere.com</host>
    <port>8080</port>
    <username>proxyuser</username>
    <password>somepassword</password>
    <nonProxyHosts>www.google.com|*.somewhere.com</nonProxyHosts>
  </proxy>
</proxies>
.
.
</settings>
```

Apache Ant 1.7 or later (Java build tool)

Apache Ant is still required for the second stage of the build process. It is used once the installation package has been constructed in `[dspace-source]/dspace/target/dspace-<version>-build.dir` and still uses some of the familiar ant build targets found in the 1.4.x build process.

Ant can be downloaded from the following location: <http://ant.apache.org>

Relational Database: (PostgreSQL or Oracle).

- **PostgreSQL 8.2 to 8.4** PostgreSQL can be downloaded from the following location: <http://www.postgresql.org/>. It is highly recommended that you try to work with Postgres 8.4 or greater, however, 8.2 or greater should still work. Unicode (specifically UTF-8) support must be enabled. This is enabled by default in 8.0+. Once installed, you need to enable TCP/IP connections (DSpace uses JDBC). In `postgresql.conf`: uncomment the line starting: `listen_addresses = 'localhost'`. Then tighten up security a bit by editing `pg_hba.conf` and adding this line: `host dspace dspace 127.0.0.1 255.255.255.255 md5`. Then restart PostgreSQL.
- **Oracle 10g or greater** Details on acquiring Oracle can be downloaded from the following location: <http://www.oracle.com/database/>. You will need to create a database for DSpace. Make sure that the character set is one of the Unicode character sets. DSpace uses UTF-8 natively, and it is suggested that the Oracle database use the same character set. You will also need to create a user account for DSpace (e.g. `dspace`) and ensure that it has permissions to add and remove tables in the database. Refer to the Quick Installation for more details.
 - **NOTE:** DSpace uses sequences to generate unique object IDs — beware Oracle sequences, which are said to lose their values when doing a database export/import, say restoring from a backup. Be sure to run the script `etc/update-sequences.sql`.

- For people interested in switching from Postgres to Oracle, I know of no tools that would do this automatically. You will need to recreate the community, collection, and eperson structure in the Oracle system, and then use the item export and import tools to move your content over.

Servlet Engine: (Apache Tomcat 5.5 or 6, Jetty, Caucho Resin or equivalent).

- **Apache Tomcat 5.5 or later.** Tomcat can be downloaded from the following location: <http://tomcat.apache.org>.
 - Note that DSpace will need to run as the same user as Tomcat, so you might want to install and run Tomcat as a user called 'dspace'. Set the environment variable `TOMCAT_USER` appropriately.
 - You need to ensure that Tomcat has a) enough memory to run DSpace and b) uses UTF-8 as its default file encoding for international character support. So ensure in your startup scripts (etc) that the following environment variable is set: `JAVA_OPTS="-Xmx512M -Xms64M -Dfile.encoding=UTF-8"`
 - **Modifications in `[tomcat]/conf/server.xml`:** You also need to alter Tomcat's default configuration to support searching and browsing of multi-byte UTF-8 correctly. You need to add a configuration option to the `<Connector>` element in `[tomcat]/config/server.xml`: `URIEncoding="UTF-8"` e.g. if you're using the default Tomcat config, it should read:

```
<!-- Define a non-SSL HTTP/1.1 Connector on port 8080 -->
<Connector port="8080"
    maxThreads="150"
    minSpareThreads="25"
    maxSpareThreads="75"
    enableLookups="false"
    redirectPort="8443"
    acceptCount="100"
    connectionTimeout="20000"
    disableUploadTimeout="true"
    URIEncoding="UTF-8" />
```

You may change the port from 8080 by editing it in the file above, and by setting the variable `CONNECTOR_PORT` in `server.xml`.

- **Jetty or Caucho Resin** DSpace will also run on an equivalent servlet Engine, such as Jetty (<http://www.mortbay.org/jetty/index.html>) or Caucho Resin (<http://www.caucho.com/>). Jetty and Resin are configured for correct handling of UTF-8 by default.

Perl (only required for `[dspace]/bin/dspace-info.pl`)

Installation Instructions

Overview of Install Options

With the advent of a new Apache [Maven 2](#) based build architecture (first introduced in DSpace 1.5.x), you now have two options in how you may wish to install and manage your local installation of DSpace. If you've used DSpace 1.4.x, please recognize that the initial build procedure has changed to allow for more customization. You will find the later 'Ant based' stages of the installation procedure familiar. Maven is used to resolve the dependencies of DSpace online from the 'Maven Central Repository' server.

It is important to note that the strategies are identical in terms of the list of procedures required to complete the build process, the only difference being that the Source Release includes "more modules" that will be built given their presence in the distribution package.

- **Default Release** (`dspace-<version>-release.zip`)
 - This distribution will be adequate for most cases of running a DSpace instance. It is intended to be the quickest way to get DSpace installed and running while still allowing for customization of the themes and branding of your DSpace instance.
 - This method allows you to customize DSpace configurations (in `dspace.cfg`) or user interfaces, using basic pre-built interface "overlays".
 - It downloads "precompiled" libraries for the core `dspace-api`, supporting servlets, taglibraries, aspects and themes for the `dspace-xmlui`, `dspace-xmlui` and other webservice/applications.
 - This approach exposes the parts of the application that the DSpace committers would prefer to see customized. All other modules are downloaded from the 'Maven Central Repository' The directory structure for this release is the following:
 - `[dspace-source]`
 - `dspace/` - DSpace 'build' and configuration module
 - `pom.xml` - DSpace Parent Project definition
- **Source Release** (`dspace-<version>-src-release.zip`)
 - This method is recommended for those who wish to develop DSpace further or alter its underlying capabilities to a greater degree.
 - It contains **all** `dspace` code for the core `dspace-api`, supporting servlets, taglibraries, aspects and themes for Manakin (`dspace-xmlui`), and other webservice/applications.
 - Provides all the same capabilities as the normal release. The directory structure for this release is more detailed:
 - `[dspace-source]`
 - `dspace/` - DSpace 'build' and configuration module
 - `dspace-api/` - Java API source module
 - `dspace-jspui/` - JSP-UI source module
 - `dspace-oai/` - OAI-PMH source module
 - `dspace-xmlui/` - XML-UI (Manakin) source module
 - `dspace-lni/` - Lightweight Network Interface source module
 - `dspace-sword` - SWORD (Simple Web-serve Offering Repository Deposit) deposit service source module
 - `dspace-test` - DSpace Tests (Unit and Integration Tests)
 - `pom.xml` - DSpace Parent Project definition

Overview of DSpace Directories

Before beginning an installation, it is important to get a general understanding of the DSpace directories and the names by which they are generally referred. (Please attempt to use these below directory names when asking for help on the DSpace Mailing Lists, as it will help everyone better understand what directory you may be referring to.)

DSpace uses three separate directory trees. Although you don't need to know all the details of them in order to install DSpace, you do need to know they exist and also know how they're referred to in this document:

1. **The installation directory**, referred to as `[dspace]`. This is the location where DSpace is installed and running off of it is the location that gets defined in the `dspace.cfg` as "dspace.dir". It is where all the DSpace configuration files, command line scripts, documentation and webapps will be installed to.
2. **The source directory**, referred to as `[dspace-source]`. This is the location where the DSpace release distribution has been unzipped into. It usually has the name of the archive that you expanded such as `dspace-<version>-release` or `dspace-<version>-src-release`. It is the directory where all of your "build" commands will be run.
3. **The web deployment directory**. This is the directory that contains your DSpace web application(s). In DSpace 1.5.x and above, this corresponds to `[dspace]/webapps` by default. However, if you are using Tomcat, you may decide to copy your DSpace web applications from `[dspace]/webapps/` to `[tomcat]/webapps/` (with `[tomcat]` being wherever you installed Tomcat, also known as `$CATALINA_HOME`). For details on the contents of these separate directory trees, refer to `directories.html`. *Note that the `[dspace-source]` and `[dspace]` directories are always separate!*

Installation

This method gets you up and running with DSpace quickly and easily. It is identical in both the Default Release and Source Release distributions.

1. **Create the DSpace user**. This needs to be the same user that Tomcat (or Jetty etc.) will run as. e.g. as **root** run:

```
useradd -m dspace
```

2. **Download the latest DSpace release** There are two version available with each release of DSpace: (*dspace-1.x-release*. and *dspace-1.x-src-release.xxx*); you only need to choose one. If you want a copy of all underlying Java source code, you should download the *dspace-1.x-src-release.xxx* Within each version, you have a choice of compressed file format. Choose the one that best fits your environment.
3. **Unpack the DSpace software**. After downloading the software, based on the compression file format, choose one of the following methods to unpack your software:

- a. **Zip file**. If you downloaded *dspace-1.6-release.zip* do the following:

```
unzip dspace-1.7-release.zip
```

- b. **.gz file**. If you downloaded *dspace-1.6-release.tar.gz* do the following:

```
gunzip -c dspace-1.7-release.tar.gz | tar -xf -
```

- c. **.bz2 file**. If you downloaded *_dspace-1.6-release.tar.bz2* do the following:

```
bunzip2 dspace-1.7-release.tar.bz | tar -xf -
```

For ease of reference, we will refer to the location of this unzipped version of the DSpace release as `[dspace-source]` in the remainder of these instructions. After unpacking the file, the user may wish to change the ownership of the *dspace-1.6-release* to the 'dspace' user. (And you may need to change the group).

4. Database Setup

- **PostgreSQL:**
 - A PostgreSQL JDBC driver is configured as part of the default DSpace build. You no longer need to copy any PostgreSQL jars to get PostgreSQL installed.
 - Create a `dspace` database, owned by the `{dspace PostgreSQL user (you are still logged in at 'root')}`:

```
createuser -U postgres -d -A -P dspace
createdb -U dspace -E UNICODE dspace
```

You will be prompted for a password for the DSpace database. (This isn't the same as the *dspace* user's UNIX password.)

- **Oracle:**
 - Setting up oracle is a bit different now. You will need still need to get a Copy of the oracle JDBC driver, but instead of copying it into a lib directory you will need to install it into your local Maven repository. (You'll need to download it first from this location: <http://www.oracle.com/technetwork/database/enterprise-edition/jdbc-112010-090769.html>) Run the following command (all on one line)

```
mvn install:install-file -Dfile=ojdbc6.jar -DgroupId=com.oracle
-DartifactId=ojdbc6 -Dversion=11.2.0.2.0 -Dpackaging=jar -DgeneratePom=true
```

- Create a database for DSpace. Make sure that the character set is one of the Unicode character sets. DSpace uses UTF-8 natively, and it is required that the Oracle database use the same character set. Create a user account for DSpace (e.g. *dspace*,) and ensure that it has permissions to add and remove tables in the database.
- Edit the *[dspace-source]/dspace/config/dspace.cfg* database settings:

```
db.name      = oracle
db.url       = jdbc:oracle:thin:@//host:port/dspace
db.driver    = oracle.jdbc.OracleDriver
```

5. **Initial Configuration:** Edit *[dspace-source]/dspace/config/dspace.cfg*, in particular you'll need to set these properties:

- *dspace.dir* - must be set to the *[dspace]* (installation) directory.
- *dspace.url* - complete URL of this server's DSpace home page.
- *dspace.hostname* - fully-qualified domain name of web server.
- *dspace.name* - "Proper" name of your server, e.g. "My Digital Library".
- *db.password* - the database password you entered in the previous step.
- *mail.server* - fully-qualified domain name of your outgoing mail server.
- *mail.from.address* - the "From:" address to put on email sent by DSpace.
- *feedback.recipient* - mailbox for feedback mail.
- *mail.admin* - mailbox for DSpace site administrator.
- *alert.recipient* - mailbox for server errors/alerts (not essential but very useful!)
- *registration.notify* - mailbox for emails when new users register (optional)

You can interpolate the value of one configuration variable in the value of another one. For example, to set *feedback.recipient* to the same value as *mail.admin*, the line would look like:

```
feedback.recipient = ${mail.admin}
```

Refer to the [General Configuration](#) section for details and examples of the above.

6. **DSpace Directory:** Create the directory for the DSpace installation (i.e. *[dspace]*). As *root* (or a user with appropriate permissions), run:

```
mkdir [dspace]
chown dspace [dspace]
```

(Assuming the *dspace* UNIX username.)

7. **Installation Package:** As the *dspace* UNIX user, generate the DSpace installation package in the *[dspace-source]/dspace* directory:

```
cd [dspace-source]/dspace/
mvn package
```

Defaults to PostgreSQL settings

Without any extra arguments, the DSpace installation package is initialized for PostgreSQL. _If you want to use Oracle instead, you should build the DSpace installation package as follows:

```
mvn -Ddb.name=oracle package
```

8. **Build DSpace and Initialize Database:** As the *dspace* UNIX user, initialize the DSpace database and install DSpace to *[dspace]_*:

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

To see a complete list of build targets, run: `ant help` *The most likely thing to go wrong here is the database connection. See the [Common Problems](#) Section.*

9. **Deploy Web Applications.** You have two choices or techniques for having Tomcat/Jetty/Resin serve up your web applications:

- **Technique A.** Simple and complete. You copy only (or all) of the DSpace Web application(s) you wish to use from the *[dspace]/webapps* directory to the appropriate directory in your Tomcat/Jetty/Resin installation. For example:

```
cp -R [dspace]/webapps/* [tomcat]/webapps* (This will copy all the web applications to Tomcat).
```

```
cp -R [dspace]/webapps/jspui [tomcat]/webapps* (This will copy only the.jspui web application to Tomcat.)
```
- **Technique B.** Tell your Tomcat/Jetty/Resin installation where to find your DSpace web application(s). As an example, in the `<Host>` section of your *[tomcat]/conf/server.xml* you could add lines similar to the following (but replace *[dspace]* with your installation location:

```
<!-- Define the default virtual host
      Note: XML Schema validation will not work with Xerces 2.2.
-->
<Host name="localhost"  appBase="[dspace]/webapps"
      ....
```

10. **Administrator Account:** Create an initial administrator account:

```
[dspace]/bin/dspace create-administrator
```

11. **Initial Startup!** Now the moment of truth! Start up (or restart) Tomcat/Jetty/Resin. Visit the base URL(s) of your server, depending on which DSpace web applications you want to use. You should see the DSpace home page. Congratulations! Base URLs of DSpace Web Applications:
- *JSP User Interface* - (e.g.) <http://dspace.myu.edu:8080/jspui>
 - *XML User Interface* (aka. Manakin) - (e.g.) <http://dspace.myu.edu:8080/xmlui>
 - *OAI-PMH Interface* - (e.g.) <http://dspace.myu.edu:8080/oai/request?verb=Identify> (Should return an XML-based response)

In order to set up some communities and collections, you'll need to login as your DSpace Administrator (which you created with `create-administrator` above) and access the administration UI in either the JSP or XML user interface.

Advanced Installation

The above installation steps are sufficient to set up a test server to play around with, but there are a few other steps and options you should probably consider before deploying a DSpace production site.

'cron' Jobs

A couple of DSpace features require that a script is run regularly – the e-mail subscription feature that alerts users of new items being deposited, and the new 'media filter' tool, that generates thumbnails of images and extracts the full-text of documents for indexing.

To set these up, you just need to run the following command as the *dspace* UNIX user:

```
crontab -e
```

Then add the following lines:

```
# Send out subscription e-mails at 01:00 every day
0 1 * * * [dspace]/bin/dspace sub-daily
# Run the media filter at 02:00 every day
0 2 * * * [dspace]/bin/dspace filter-media
# Run the checksum checker at 03:00
0 3 * * * [dspace]/bin/dspace checker -lp
# Mail the results to the sysadmin at 04:00
0 4 * * * [dspace]/bin/dspace checker-emailer -c
```

Naturally you should change the frequencies to suit your environment.

PostgreSQL also benefits from regular 'vacuuming', which optimizes the indexes and clears out any deleted data. Become the *postgres* UNIX user, run `crontab -e` and add (for example):

```
# Clean up the database nightly at 4.20am
20 4 * * * vacuumdb --analyze dspace > /dev/null 2>&1
```

In order that statistical reports are generated regularly and thus kept up to date you should set up the following cron jobs:

```
# Run stat analysis
0 1 * * * [dspace]/bin/dspace stat-general
0 1 * * * [dspace]/bin/dspace stat-monthly
0 2 * * * [dspace]/bin/dspace stat-report-general
0 2 * * * [dspace]/bin/dspace stat-report-monthly
```

Obviously, you should choose execution times which are most useful to you, and you should ensure that the ~~report~~ scripts run a short while after the analysis scripts to give them time to complete (a run of around 8 months worth of logs can take around 25 seconds to complete); the resulting reports will let you know how long analysis took and you can adjust your cron times accordingly.

Multilingual Installation

In order to deploy a multilingual version of DSpace you have to configure two parameters in `[dspace-source]/config/dspace.cfg`:

- *default.locale*, e.g. `default.locale = en`
- *webui.supported.locales*, e.g. `webui.supported.locales = en, de`

The Locales might have the form `country`, `country_language`, `country_language_variant`.

According to the languages you wish to support, you have to make sure, that all the i18n related files are available see the Multilingual User Interface Configuring MultiLingual Support section for the JSPUI or the Multilingual Support for XMLUI in the configuration documentation.

DSpace over HTTPS

If your DSpace is configured to have users login with a username and password (as opposed to, say, client Web certificates), then you should consider using HTTPS. Whenever a user logs in with the Web form (e.g. `dspace.myuni.edu/dspace/password-login`) their DSpace password is exposed in plain text on the network. This is a very serious security risk since network traffic monitoring is very common, especially at universities. If the risk seems minor, then consider that your DSpace administrators also login this way and they have ultimate control over the archive.

The solution is to use *HTTPS* (HTTP over SSL, i.e. Secure Socket Layer, an encrypted transport), which protects your passwords against being captured. You can configure DSpace to require SSL on all "authenticated" transactions so it only accepts passwords on SSL connections.

The following sections show how to set up the most commonly-used Java Servlet containers to support HTTP over SSL.

To enable the HTTPS support in Tomcat 5.0:

1. **For Production use:** Follow this procedure to set up SSL on your server. Using a "real" server certificate ensures your users' browsers will accept it without complaints. In the examples below, `$CATALINA_BASE` is the directory under which your Tomcat is installed.
 - a. Create a Java keystore for your server with the password *changeit*, and install your server certificate under the alias *"tomcat"*. This assumes the certificate was put in the file *server.pem*:

```
$JAVA_HOME/bin/keytool -import -noprompt -v -storepass changeit
                        -keystore $CATALINA_BASE/conf/keystore -alias tomcat -file
myserver.pem
```

- b. Install the CA (Certifying Authority) certificate for the CA that granted your server cert, if necessary. This assumes the server CA certificate is in *ca.pem*:

```
$JAVA_HOME/bin/keytool -import -noprompt -storepass changeit
                        -trustcacerts -keystore $CATALINA_BASE/conf/keystore -alias ServerCA
                        -file ca.pem
```

- c. Optional – ONLY if you need to accept client certificates for the X.509 certificate stackable authentication module See the configuration section for instructions on enabling the X.509 authentication method. Load the keystore with the CA (certifying authority) certificates for the authorities of any clients whose certificates you wish to accept. For example, assuming the client CA certificate is in *client1.pem*:

```
$JAVA_HOME/bin/keytool -import -noprompt -storepass changeit
                        -trustcacerts -keystore $CATALINA_BASE/conf/keystore -alias client1
                        -file client1.pem
```

- d. Now add another Connector tag to your *server.xml* Tomcat configuration file, like the example below. The parts affecting or specific to SSL are shown in bold. (You may wish to change some details such as the port, pathnames, and keystore password)

```
<Connector port="8443"
            maxThreads="150" minSpareThreads="25"
            maxSpareThreads="75"
            enableLookups="false"
            disableUploadTimeout="true"
            acceptCount="100" debug="0"
            scheme="https" secure="true" sslProtocol="TLS"
            keystoreFile="conf/keystore" keystorePass="changeit" clientAuth="true" - ONLY if using
            client X.509 certs for authentication!
            truststoreFile="conf/keystore" trustedstorePass="changeit" />
```

Also, check that the default Connector is set up to redirect "secure" requests to the same port as your SSL connector, e.g.:

```
<Connector port="8080"
            maxThreads="150" minSpareThreads="25"
            maxSpareThreads="75"
            enableLookups="false"
            redirectPort="8443"
            acceptCount="100" debug="0" />
```

2. **Quick-and-dirty Procedure for Testing:** If you are just setting up a DSpace server for testing, or to experiment with HTTPS, then you don't need to get a real server certificate. You can create a "self-signed" certificate for testing; web browsers will issue warnings before accepting it but they will function exactly the same after that as with a "real" certificate. In the examples below, *\$CATALINA_BASE* is the directory under which your Tomcat is installed.

a. Optional – ONLY if you don't already have a server certificate. Follow this sub-procedure to request a new, signed server certificate from your Certifying Authority (CA):

- Create a new key pair under the alias name *"tomcat"*. When generating your key, give the Distinguished Name fields the appropriate values for your server and institution. CN should be the fully-qualified domain name of your server host. Here is an example:

```
$JAVA_HOME/bin/keytool -genkey -alias tomcat -keyalg RSA -keysize
1024 \
-keystore $CATALINA_BASE/conf/keystore -storepass changeit
-validity 365 \
-dname 'CN=dspace.myuni.edu, OU=MIT Libraries, O=Massachusetts
Institute of Technology, L=Cambridge, S=MA, C=US'
```

- Then, create a CSR (Certificate Signing Request) and send it to your Certifying Authority. They will send you back a signed Server Certificate. This example command creates a CSR in the file *tomcat.csr*

```
$JAVA_HOME/bin/keytool -keystore $CATALINA_BASE/conf/keystore
-storepass changeit \
-certreq -alias tomcat -v -file tomcat.csr
```

- Before importing the signed certificate, you must have the CA's certificate in your keystore as a *trusted certificate*. Get their certificate, and import it with a command like this (for the example *mitCA.pem*):

```
$JAVA_HOME/bin/keytool -keystore $CATALINA_BASE/conf/keystore
-storepass changeit \
-import -alias mitCA -trustcacerts -file mitCA.pem
```

- Finally, when you get the signed certificate from your CA, import it into the keystore with a command like the following example: (cert is in the file *signed-cert.pem*)

```
$JAVA_HOME/bin/keytool -keystore $CATALINA_BASE/conf/keystore
-storepass changeit \
-import -alias tomcat -trustcacerts -file signed-cert.pem
```

Since you now have a signed server certificate in your keystore, you can, obviously, skip the next steps of installing a signed server certificate and the server CA's certificate.

b. Create a Java keystore for your server with the password *changeit*, and install your server certificate under the alias *"tomcat"*. This assumes the certificate was put in the file *server.pem*:

```
$JAVA_HOME/bin/keytool -genkey -alias tomcat -keyalg RSA -keystore
$CATALINA_BASE/conf/keystore -storepass changeit
```

When answering the questions to identify the certificate, be sure to respond to "First and last name" with the fully-qualified domain name of your server (e.g. *test-dspace.myuni.edu*). The other questions are not important.

c. Optional – ONLY if you need to accept client certificates for the X.509 certificate stackable authentication module See the configuration section for instructions on enabling the X.509 authentication method. Load the keystore with the CA (certifying authority) certificates for the authorities of any clients whose certificates you wish to accept. For example, assuming the client CA certificate is in *client1.pem*:

```
$JAVA_HOME/bin/keytool -import -noprompt -storepass changeit
-trustcacerts -keystore $CATALINA_BASE/conf/keystore -alias client1
-file client1.pem
```

d. Follow the procedure in the section above to add another Connector tag, for the HTTPS port, to your *server.xml* file.

To use SSL on Apache HTTPD with mod_jk:

If you choose [Apache HTTPD](#) as your primary HTTP server, you can have it forward requests to the [Tomcat servlet container](#) via [Apache Jakarta Tomcat Connector](#). This can be configured to work over SSL as well. First, you must configure Apache for SSL; for Apache 2.0 see [Apache SSL/TLS Encryption](#) for information about using `mod_ssl`.

If you are using X.509 Client Certificates for authentication: add these configuration options to the appropriate `httpd` configuration file, e.g. `ssl.conf`, and be sure they are in force for the virtual host and namespace locations dedicated to DSpace:

```
## SSLVerifyClient can be "optional" or
"require"
SSLVerifyClient optional
SSLVerifyDepth 10
SSLCACertificateFile
path-to-your-client-CA-certificate
SSLOptions StdEnvVars ExportCertData
```

Now consult the [Apache Jakarta Tomcat Connector](#) documentation to configure the `mod_jk` (note: **NOT** `mod_jk2`) module. Select the AJP 1.3 connector protocol. Also follow the instructions there to configure your Tomcat server to respond to AJP.

To use SSL on Apache HTTPD with `mod_webapp` consult the DSpace 1.3.2 documentation. Apache have deprecated the `mod_webapp` connector and recommend using `mod_jk`.

To use Jetty's HTTPS support consult the documentation for the relevant tool.

The Handle Server

First a few facts to clear up some common misconceptions:

- You don't **have** to use CNRI's Handle system. At the moment, you need to change the code a little to use something else (e.g PURLs) but that should change soon.
- You'll notice that while you've been playing around with a test server, DSpace has apparently been creating handles for you looking like `hdl:123456789/24` and so forth. These aren't really Handles, since the global Handle system doesn't actually know about them, and lots of other DSpace test installs will have created the same IDs. They're only really Handles once you've registered a prefix with CNRI (see below) and have correctly set up the Handle server included in the DSpace distribution. This Handle server communicates with the rest of the global Handle infrastructure so that anyone that understands Handles can find the Handles your DSpace has created. If you want to use the Handle system, you'll need to set up a Handle server. This is included with DSpace. Note that this is not required in order to evaluate DSpace; you only need one if you are running a production service. You'll need to obtain a Handle prefix from [the central CNRI Handle site](#).

A Handle server runs as a separate process that receives TCP requests from other Handle servers, and issues resolution requests to a global server or servers if a Handle entered locally does not correspond to some local content. The Handle protocol is based on TCP, so it will need to be installed on a server that can broadcast and receive TCP on port 2641.

1. To configure your DSpace installation to run the handle server, run the following command:

```
[dspace]/bin/dspace make-handle-config [dspace]/handle-server
```

Ensure that `[dspace]/handle-server` matches whatever you have in `dspace.cfg` for the `handle.dir` property.

2. Edit the resulting `[dspace]/handle-server/config.dct` file to include the following lines in the "server_config" clause:

```
"storage_type" = "CUSTOM"
"storage_class" = "org.dspace.handle.HandlePlugin"
```

This tells the Handle server to get information about individual Handles from the DSpace code.

3. Once the configuration file has been generated, you will need to go to <http://hdl.handle.net/4263537/5014> to upload the generated sitebndl.zip file. The upload page will ask you for your contact information. An administrator will then create the naming authority/prefix on the root service (known as the Global Handle Registry), and notify you when this has been completed. You will not be able to continue the handle server installation until you receive further information concerning your naming authority.
4. When CNRI has sent you your naming authority prefix, you will need to edit the `config.dct` file. The file will be found in `[dspace]/handle-server`. Look for "300:0.NA/YOUR_NAMING_AUTHORITY". Replace `YOUR_NAMING_AUTHORITY` with the assigned naming authority prefix sent to you.
5. Now start your handle server (as the dspace user):

```
[dspace]/bin/start-handle-server
```

Note that since the DSpace code manages individual Handles, administrative operations such as Handle creation and modification aren't supported by DSpace's Handle server.

Updating Existing Handle Prefixes

If you need to update the handle prefix on items created before the CNRI registration process you can run the `[dspace]/bin/dspace update-handle-prefix script`. You may need to do this if you loaded items prior to CNRI registration (e.g. setting up a demonstration system prior to migrating it to production). The script takes the current and new prefix as parameters. For example:

```
[dspace]/bin/dspace update-handle-prefix 123456789 1303
```

This script will change any handles currently assigned prefix 123456789 to prefix 1303, so for example handle 123456789/23 will be updated to 1303/23 in the database.

Google and HTML sitemaps

To aid web crawlers index the content within your repository, you can make use of sitemaps. There are currently two forms of sitemaps included in DSpace: Google sitemaps and HTML sitemaps.

Sitemaps allow DSpace to expose its content without the crawlers having to index every page. HTML sitemaps provide a list of all items, collections and communities in HTML format, whilst Google sitemaps provide the same information in gzipped XML format.

To generate the sitemaps, you need to run `[dspace]/bin/dspace generate-sitemaps`. This creates the sitemaps in `[dspace]/sitemaps/`.

The sitemaps can be accessed from the following URLs:

- <http://dspace.example.com/dspace/sitemap> - Index sitemap
 - <http://dspace.example.com/dspace/sitemap?map=0> - First list of items (up to 50,000)
 - <http://dspace.example.com/dspace/sitemap?map=n> - Subsequent lists of items (e.g. 50,0001 to 100,000) etc...
- HTML sitemaps follow the same procedure:
- <http://dspace.example.com/dspace/htmlmap> - Index HTML based sitemap
 - etc...

When running `[dspace]/bin/dspace generate-sitemaps` the script informs Google that the sitemaps have been updated. For this update to register correctly, you must first register your Google sitemap index page (`/dspace/sitemap`) with Google at <http://www.google.com/webmasters/sitemaps/>. If your DSpace server requires the use of a HTTP proxy to connect to the Internet, ensure that you have set `http.proxy.host` and `http.proxy.port` in `[dspace]/config/dspace.cfg`.

The URL for pinging Google, and in future, other search engines, is configured in `[dspace-space]/config/dspace.cfg` using the `sitemap.engineurls` setting where you can provide a comma-separated list of URLs to 'ping'.

You can generate the sitemaps automatically every day using an additional cron job:

```
# Generate sitemaps

0 6 * * * [dspace]/bin/dspace generate-sitemaps
```

DSpace Statistics

DSpace uses the Apache Solr application underlying the statistics. There is no need to download any separate software. All the necessary software is included. To understand all of the configuration property keys, the user should refer to 5.2.35 DSpace Statistic Configuration for detailed information.

1. **DSpace Configuration for Accessing Solr.** In the `dspace.cfg` file review the following fields to make sure they are uncommented:

```
solr.log.server = ${dspace.baseUrl}/solr/statistics
solr.dbfile = ${dspace.dir}/config/GeoLiteCity.dat
solr.spiderips.urls = http://iplists.com/google.txt, \
                     http://iplists.com/inktomi.txt, \
                     http://iplists.com/lycos.txt, \
                     http://iplists.com/infoseek.txt, \
                     http://iplists.com/altavista.txt, \
                     http://iplists.com/excite.txt, \
                     http://iplists.com/misc.txt, \
                     http://iplists.com/non_engines.txt
```

2. **DSpace logging configuration for Solr.** If your DSpace instance is protected by a proxy server, in order for Solr to log the correct IP address of the user rather than of the proxy, it must be configured to look for the X-Forwarded-For header. This feature can be enabled by ensuring the following setting is uncommented in the logging section of `dspace.cfg`:

```
useProxies = true
```

3. **Configuration Control.** In the *dspace.cfg* set the following property key: `_statistics.item.authorization.admin=true`. This will require the user to sign on to see that statistics. Setting the statistics to "false" will make them publicly available.
4. Final steps.
 - Perform the following step:

```
cd [dspace-source]/dspace
mvn package
cd [dspace-source]/dspace/target/dspace-<version>-build.dir
ant -Dconfig=[dspace]/config/dspace.cfg update
cp -R [dspace]/webapps/* [TOMCAT]/webapps
```

If you only need to build the statistics, and don't make any changes to other web applications, you can replace the copy step above with:

```
cp -R [dspace]/webapps/solr [TOMCAT]/webapps
```

- Restart your webapps (Tomcat/Jetty/Resin)

Windows Installation

Pre-requisite Software

If you are installing DSpace on Windows, you will still need to install all the same [Prerequisite Software](#), as listed above.

- If you install PostgreSQL, it's recommended to select to install the pgAdmin III tool. It provides a nice User Interface for interacting with PostgreSQL databases.

Installation Steps

1. Download the DSpace source from [SourceForge](#) and unzip it ([WinZip](#) will do this)
2. Ensure the PostgreSQL service is running, and then run pgAdmin III (Start -> PostgreSQL 8.0 -> pgAdmin III). Connect to the local database as the postgres user and:
 - Create a 'Login Role' (user) called *dspace* with the password *dspace*
 - Create a database called *dspace* owned by the user *dspace*, with UTF-8 encoding
3. Update paths in *[dspace-source]/dspace/config/dspace.cfg*
 - **Note:** Use forward slashes / for path separators, though you can still use drive letters, e.g.: *dspace.dir = C:/DSpace*. Also, make sure you change all of the parameters with file paths to suit, specifically:

```
dspace.dir
config.template.log4j.properties
config.template.log4j-handle-plugin.properties
config.template.oaicat.properties
assetstore.dir
log.dir
upload.temp.dir
report.dir
handle.dir
```

4. Create the directory for the DSpace installation (e.g. *C:/DSpace*)
5. Generate the DSpace installation package by running the following from command line (cmd) from your *[dspace-source]/dspace/* directory:

```
mvn package
```

- **Note #1:** This will generate the DSpace installation package in your *[dspace-source]/dspace/target/dspace-[version]-build.dir/* directory.
- **Note #2:** Without any extra arguments, the DSpace installation package is initialized for PostgreSQL. If you want to use Oracle instead, you should build the DSpace installation package as follows:

```
mvn -Ddb.name=oracle package
```

6. Initialize the DSpace database and install DSpace to *[dspace]* (e.g. *C:/DSpace*) by running the following from command line from your *[dspace-source]/dspace/target/dspace-[version]-build.dir/* directory:

```
ant fresh_install
```

- *Note:* to see a complete list of build targets, run: `ant help`
7. Create an administrator account, by running the following from your *[dSPACE]* (e.g. *C:\DSpace*) directory:

```
[dSPACE]\bin\dspace create-administrator
```

8. Copy the Web application directories from *[dSPACE]\webapps* to Tomcat's webapps dir, which should be somewhere like *C:\Program Files\Apache Software Foundation\Tomcat\webapps*
 - Alternatively, Tell your Tomcat installation where to find your DSpace web application(s). As an example, in the *<Host>* section of your *[tomcat/conf/server.xml]* you could add lines similar to the following (but replace *[dSPACE]* with your installation location):

```
<!-- DEFINE A CONTEXT PATH FOR DSpace JSP User Interface -->
<Context path="/jspui" docBase="[dSPACE]\webapps\jspui" debug="0"
        reloadable="true" cachingAllowed="false"
        allowLinking="true"/>

<!-- DEFINE A CONTEXT PATH FOR DSpace OAI User Interface -->
<Context path="/oai" docBase="[dSPACE]\webapps\oai" debug="0"
        reloadable="true" cachingAllowed="false"
        allowLinking="true"/>
```

9. Start the Tomcat service
10. Browse to either <http://localhost:8080/jspui> or <http://localhost:8080/xmlui>. You should see the DSpace home page for either the JSPUI or XMLUI, respectively.

Checking Your Installation

The administrator needs to check the installation to make sure all components are working. Here is list of checks to be performed. In brackets after each item, it the associated component or components that might be the issue needing resolution.

- System is up and running. *User can see the DSpace home page. [Tomcat/Jetty, firewall, IP assignment, DNS]*
- Database is running and working correctly. *Attempt to create a user, community or collection [PostgreSQL, Oracle]Run the test database command to see if other issues are being report:[dSPACE]/bin/dSPACE test-database*
- Email subsystem is running. The user can issue the following command to test the email system. *t attempts to send a test email to the email address that is set in dSPACE.cfg (mail.admin). If it fails, you will get messages informing you as to why, will refer you to the DSpace documentation. [dSPACE]/bin/test-email*

Known Bugs

In any software project of the scale of DSpace, there will be bugs. Sometimes, a stable version of DSpace includes known bugs. We do not always wait until every known bug is fixed before a release. If the software is sufficiently stable and an improvement on the previous release, and the bugs are minor and have known workarounds, we release it to enable the community to take advantage of those improvements.

The known bugs in a release are documented in the *KNOWN_BUGS* file in the source package.

Please see the [DSpace bug tracker](#) for further information on current bugs, and to find out if the bug has subsequently been fixed. This is also where you can report any further bugs you find.

Common Problems

In an ideal world everyone would follow the above steps and have a fully functioning DSpace. Of course, in the real world it doesn't always seem to work out that way. This section lists common problems that people encounter when installing DSpace, and likely causes and fixes. This is likely to grow over time as we learn about users' experiences.

- **Database errors occur when you run `ant fresh_install`:** There are two common errors that occur.
 - If your error looks like this:

```
[java] 2004-03-25 15:17:07,730 INFO
        org.dspace.storage.rdbms.InitializeDatabase @ Initializing Database
[java] 2004-03-25 15:17:08,816 FATAL
        org.dspace.storage.rdbms.InitializeDatabase @ Caught exception:
[java] org.postgresql.util.PSQLException: Connection refused. Check
        that the hostname and port are correct and that the postmaster is
        accepting TCP/IP connections.
[java]      at
        org.postgresql.jdbc1.AbstractJdbc1Connection.openConnection(AbstractJd
        bc1Connection.java:204)
[java]      at org.postgresql.Driver.connect(Driver.java:139)
```

it usually means you haven't yet added the relevant configuration parameter to your PostgreSQL configuration (see above), or perhaps you haven't restarted PostgreSQL after making the change. Also, make sure that the *db.username* and *db.password* properties are correctly set in *[dspace-source]/config/dspace.cfg*. An easy way to check that your DB is working OK over TCP/IP is to try this on the command line:

```
psql -U dspace -W -h localhost
```

Enter the *dspace* database password, and you should be dropped into the psql tool with a *dspace=>* prompt.

- Another common error looks like this:

```
[java] 2004-03-25 16:37:16,757 INFO
      org.dspace.storage.rdbms.InitializeDatabase @ Initializing Database
[java] 2004-03-25 16:37:17,139 WARN
      org.dspace.storage.rdbms.DatabaseManager @ Exception initializing DB
      pool
[java] java.lang.ClassNotFoundException: org.postgresql.Driver
[java]     at java.net.URLClassLoader$1.run(URLClassLoader.java:198)
[java]     at java.security.AccessController.doPrivileged(Native
      Method)
[java]     at
      java.net.URLClassLoader.findClass(URLClassLoader.java:186)
```

This means that the PostgreSQL JDBC driver is not present in *[dspace-source]/lib*. See above.

- **Tomcat doesn't shut down:** If you're trying to tweak Tomcat's configuration but nothing seems to make a difference to the error you're seeing, you might find that Tomcat hasn't been shutting down properly, perhaps because it's waiting for a stale connection to close gracefully which won't happen.
 - To see if this is the case, try running: `ps -ef | grep java` and look for Tomcat's Java processes. If they stay around after running Tomcat's *shutdown.sh* script, trying running `kill` on them (or `kill -9` if necessary), then starting Tomcat again.
- **Database connections don't work, or accessing DSpace takes forever:** If you find that when you try to access a DSpace Web page and your browser sits there connecting, or if the database connections fail, you might find that a 'zombie' database connection is hanging around preventing normal operation.
 - To see if this is the case, try running: `ps -ef | grep postgres`
 - You might see some processes like this:

```
dspace 16325  1997  0  Feb 14  ?           0:00 postgres: dspace dspace  127.0.0.1 idle in
transaction
```

This is normal. DSpace maintains a 'pool' of open database connections, which are re-used to avoid the overhead of constantly opening and closing connections. If they're 'idle' it's OK; they're waiting to be used.

- However sometimes, if something went wrong, they might be stuck in the middle of a query, which seems to prevent other connections from operating, e.g.:

```
dspace 16325  1997  0  Feb 14  ?           0:00 postgres: dspace dspace  127.0.0.1 SELECT
```

This means the connection is in the middle of a *SELECT* operation, and if you're not using DSpace right that instant, it's probably a 'zombie' connection. If this is the case, try running `kill` on the process, and stopping and restarting Tomcat.