

Confluence 3.0 - 3.4 Setup

This is now an archive article.

This article has lots of room for improvement in tone and the instructions can be more clear.

- [Install Required Libraries](#)
- [PostgreSQL](#)
 - [Install PostgreSQL on Ubuntu](#)
- [Install Confluence](#)
 - [Core Package](#)
 - [Manually Setup JDK](#)
 - [Configure Confluence to Use JDK](#)
 - [Configure Data Directory](#)
 - [Configure PostgreSQL](#)
 - [Setup Database](#)
 - [Setup JDBC4 Driver](#)
- [Starting the Confluence Service](#)
- [Stopping the Service](#)
- [Administrator Tasks](#)
- [Optimize Confluence](#)
- [Connect Apache to Confluence using Mod_JK](#)
 - [Enable Tomcat AJP Connector](#)
 - [Select and Enable Context Path](#)
 - [Configure Apache Tomcat Connector](#)
- [Verify Internationalization Features](#)
- [Future Upgrade Notes](#)
 - [Database Password](#)
 - [Key Files](#)
 - [Special Modifications](#)
 - [Package Dependencies](#)
- [Disable Search Engine Indexing](#)

Install Required Libraries

As of [Confluence 3.5](#), [X11 dependencies](#) are now no longer needed.

In order to render thumbnails and related functionality you need certain x-windows libraries. Refer to the list from [the X-11 Dependencies](#) [Confluence install instructions](#) and make sure to use the list matching your version of Confluence.

Because the list changes between each release, I recommend documenting what you installed. Then when upgrading, you can take two approaches,

1. Remove all the required libraries from the previous Confluence release and then installing the required libraries for the new Confluence release
2. Diff and remove unused libraries and install the deltaWe find that due to complex dependencies approach 1 is faster and easier.

Due to complex dependencies, we find that the first approach works better.

PostgreSQL

Confluence comes with its own embedded database called hsqldb. However, for real production use you should use a full fledged database like PostgreSQL.

Install PostgreSQL on Ubuntu

Install [PostgreSQL6.0 PostgreSQL on Ubuntu](#).

Install Confluence

Core Package

Download the tar.gz install package and decompress it,

```
su - serveradmin # if you are not already
gunzip confluence-3.2-std.tar.gz
tar -xvpf confluence-3.2-std.tar
su - bhitch # log in as a staff user that can run sudo
# Use specific version numbers. This will make upgrade easier in the
future.
sudo mv home/serveradmin/confluence-3.2-std /opt/
```

Next, we make the symbolic link directory. Scripts and anything else you do will use the symbolic link directory. When the new version of Confluence comes along you have various options for upgrading.

```
su - bhitch # log in as a staff user that can run sudo
cd /opt
# make a symbolic link which scripts and commands can run against
sudo ln -s ./confluence-3.2-std ./confluence
```

Manually Setup JDK

I use a manual Java setup with an instance of Java specifically for the application.

Make sure you are install a JDK. This is a [requirement](#) of Confluence.

Unable to render {include} The included page could not be found.

Log in as your staff account which has sudo access to perform the actual move to /opt/

```
cd /home/serveradmin
sudo mv ./jdk1.6.0_16/ /opt/confluence/java # moves and renames the
directory in one step
```

Configure Confluence to Use JDK

The downloaded Confluence package is powered by Tomcat. You can specify Tomcat use a specific java directory by **modifying** the Confluence setenv.sh file as described at [Running Tomcat with a specific JRE version](#).

This can be improved by moving "how to specify java for Tomcat" out of the Tomcat 6 setup instructions into a separate page and then using an "include" here.

Configure Data Directory

Also we need to create the home directory,

```
su - serveradmin
cd ~
mkdir confluence-data # Don't use specific version numbers here because the
data will actually be upgraded
su - bhitch # to get sudo access
sudo mv /home/serveradmin/confluence-data /opt/
```

As serveradmin, set the home (data) directory by adding confluence.home to the bottom of, **/opt/confluence/confluence/WEB-INF/classes/confluence-init.properties**,

```
# confluence.home=c:/confluence/data
confluence.home=/opt/confluence-data/
```

Configure PostgreSQL

Setup Database

These are the command line version of the instructions from [Confluences wiki on setting up an external database](#).

The commands were derived from the [visual guide by Confluence](#).

First log into the PostgreSQL prompt as explain in [PostgreSQL on Ubuntu](#).

Create the confluence database admin account.

Make sure to change putYourPasswordHere to a secure password.

```
CREATE ROLE confluencedbuser LOGIN ENCRYPTED PASSWORD 'putYourPasswordHere'
NOINHERIT
VALID UNTIL 'infinity';
COMMENT ON ROLE confluencedbuser IS 'Account used by the Confluence
application.';
```

The servers will respond to each respective command,

```
CREATE ROLE
COMMENT
```

Create the database.

```
CREATE DATABASE confluencedb
WITH ENCODING='UTF8'
OWNER=confluencedbuser
CONNECTION LIMIT=-1;
```

If all goes well the response will be,

```
CREATE DATABASE
```

Finally quit the PostgreSQL prompt,

```
\q
```

The name of the database is different than in the confluence documentation. Rather than the database name confluence, we use confluencedb to clearly designate a database.

Setup JDBC4 Driver

As of Confluence 3.5.1 the PostgreSQL JDBC4 driver is now included. This section is still useful as a reference for future version mismatches.

Confluence 3.2 - 3.5.0 does include a JDBC driver for PostgreSQL. However, because we are using the newest version of Java, download the JDBC4 driver from the [PostgreSQL website](http://jdbc.postgresql.org/download/postgresql-9.0-801.jdbc4.jar) and place the jar file into the /opt/confluence/confluence/WEB-INF/lib/ directory. Then delete the old JDBC4 driver.

```
su - serveradmin
wget http://jdbc.postgresql.org/download/postgresql-9.0-801.jdbc4.jar
rm /opt/confluence/confluence/WEB-INF/lib/postgresql-8.4-701.jdbc3.jar
```

Starting the Confluence Service

We will now start Confluence for the first time to initialize everything.

If using virtual machine technology, now would be a good time to take a snapshot of the system.

```
cd /opt/confluence/bin/
./startup.sh
```

Next, hit the website with a browser, <http://localhost:8080/> (change localhost to ip address of the server) to start the first time configuration wizard.

If you have been following these notes on how to setup Ubuntu, your firewall may be enabled. Make sure to enable access to port 8080.

Step through the configuration wizard and make the following selections,

1. Choose Product Installation
2. Make sure PostgreSQL is selected (it should be by default) click External Database.
3. Direct JDBC
4. Fill out the form,
 - a. Database URL: Change the default line, jdbc:postgresql://localhost:5432/confluence to jdbc:postgresql://localhost:5432/confluencedb
 - b. User: confluencedbuser
 - c. Pass: the password you used earlier to create the user
5. The initial user will be called "Server Admin" with user name serveradmin

Stopping the Service

There appears to be a [problem with shutting down the Confluence server](#). According to Atlassian, this only happens on certain configuration and they are not able to determine yet why. I have personally experienced with every type of Virtualized Ubuntu system I have setup over the past few

years.

Start by trying to perform a normal shutdown,

```
su - svradm
/opt/confluence/bin/shutdown.sh
```

Next ensure that Confluence actually shut down,

```
ps -ef | grep java | grep confluence
```

The resulting output shows all your java processes. Look for you Confluence java process which looks something similar to below,

```
svradm      2364      1  0  01:00 ?          00:02:37 /opt
```

If it does not exist then the shutdown occurred properly. However, if it still exists try running ps again in a few minutes. If it still exists after that you need to kill the process,

```
kill 2364 # Note change 2364, to the process ID that you see on your
screen.
```

Administrator Tasks

This section still needs to be written to take the user through setting up the base url, creating their first space, disabling automatic backup, how to do real backups and any other important activities.

Optimize Confluence

There are some basic tasks we should do before getting other people on board.

Once logged in go to, Browser, Confluence Admin. There are tasks there that Confluence recommends, the most important to us being,

These instructions should be written with more detail. Particularly, "Setup email".

1. Disable automatic backup
2. Setup email
3. Update Confluence plugins

Connect Apache to Confluence using Mod_JK

As this install of Confluence is powered by Tomcat we can front Confluence with the Apache Web Server.

Enable Tomcat AJP Connector

First we must enable the the Connector. In this example we will use the standard port 8009 in the Confluence Tomcat. Edit `/opt/confluence/conf/server.xml`.

Add the **Connector** directive `<Connector port="8009" protocol="AJP/1.3 "URIEncoding="UTF-8" />` just above the **Engine** directive,

```
useURISValidationHack="false" URLEncoder="UTF-8" />

    <Connector port="8009" protocol="AJP/1.3" URLEncoder="UTF-8" />

    <Engine name="Standalone"
```

The UTF-8 is important as Confluence is developed to post using UTF-8.

Without this setting [errors occur](#) when working with attachments containing special characters. Further details are covered in the [Bonsai Tomcat Connector](#) article.

Select and Enable Context Path

Often you want to use the Apache Web Server to front more than one web application. In order to do this you want each web application to work within its own context. This is pretty standard in most Java applications. As such, we will pick a context path, in this case "wiki". As such the original url, <http://localhost:8080/> changes to <http://localhost:8080/wiki>

Modify the bundled Confluence Tomcat **Context** directive and change the attribute **path=""** to **path="/wiki"**,

```
<Context path="/wiki" docBase="../confluence" debug="0" reloadable="false"
useHttpOnly="true">
```

This change does not take effect until the Confluence service stopped and started again. We will do that in the next section.

Configure Apache Tomcat Connector

Still need to fill in clearer instructions on this topic.

Once port 8009 is enabled, mod_jk can be enabled by following the [Apache Tomcat Connector instructions](#).

Verify Internationalization Features

Even if you do not plan to have the wiki work for multiple languages, you should verify the internationalization is working properly just to handle special characters.

Visit the `http://<host address>:<port>/admin/encodingtest.action` url on your system and follow the instructions.

Future Upgrade Notes

As part of planning for future upgrades here is a checklist to keep in mind.

Database Password

Make sure you know the admin password for your database.

Key Files

- /opt/confluence/bin/setenv.sh
- /opt/confluence/conf/server.xml
- /opt/confluence/confluence/WEB-INF/classes/confluence-init.properties
- /opt/confluence/confluence/WEB-INF/lib/postgresql-8.4-701.jdbc4.jar - add

- /opt/confluence/confluence/WEB-INF/lib/postgresql-8.4-701.jdbc3.jar - remove
- /opt/confluence/bin/shutdown.sh
- /opt/confluence/bin/startup.sh

Special Modifications

You may choose to modify special files. For example, I [modify the footer](#),

- /opt/confluence/confluence/decorators/includes/footer-content.vm

Package Dependencies

The package dependencies change between confluence versions. You should check if there are differences. For example look at the difference between 3.1 and 3.2,

As of [Confluence 3.5](#), [X11 dependencies](#) are now no longer needed.

confluence-3.1

- libice-dev - no longer used
- libsm-dev - no longer used
- libx11-dev - used
- libxext-dev - no longer used
- libxp-dev - no longer used
- libxt-dev - no longer used
- libxtst-dev - used

confluence-3.2

- libx11-6 - new
- libx11-dev - carry over
- libxt6 - new
- libxt6-dbg - new
- libxext6 - new
- libxtst-dev - carry over
- libxtst6 - new
- xlibs-dbg - new
- xlibs-dev - new

Also, keep a page listing any plugins you install. Upgrades generally do not require you to reinstall the plugins but they might be deprecated. After the upgrade is done, login as Administrator on Confluence and ensure plug-ins are up to date.

Disable Search Engine Indexing

In the case of a private site, you may not want Google or any other search engine to crawl your wiki. Even in the case of a password protected site this will still happen and take up unnecessary bandwidth, processing and bump up the amount of memory consumed by Confluence.

This content needs to be refined and should in fact link to a detailed article on robots.txt and other ways of preventing crawling.

If you had been following the steps outline by the Bonsai Framework the Confluence configuration will be,

- Fronted by Apache
- The root of the website is located at /home/www.krypton.com/www
- The mapping from Apache to Confluence is /wiki/

As such, robots.txt will be placed in the root of the website and look like this,

```
User-agent: *
Disallow: /wiki/
```