

Confluence 3.5+ Setup

- Introduction
- PostgreSQL
 - Install PostgreSQL on Ubuntu
- Install Confluence
 - Core Package
 - Manually Setup JDK
 - Configure Confluence to Use JDK
 - Configure Confluence to Only Start and Stop with ServerAdmin
 - Configure Data Directory
 - Configure PostgreSQL
 - Setup Database
 - Setup JDBC4 Driver
- Starting the Confluence Service
- Stopping the Service
- Administrator Tasks
- Setup Confluence to Send Emails
- Connect Apache to Confluence using Mod_JK
 - Enable Tomcat AJP Connector
 - Select and Enable Context Path
 - Configure Apache Tomcat Connector
 - Change Base Url in Confluence
- Verify Internationalization Features
- Future Upgrade Notes
 - Backup Your Data
 - Database Password
 - Key Files
 - Special Modifications
 - Plug-Ins
- Disable Search Engine Indexing
- System Profile

Introduction

This article takes you through a Confluence 3.5+ setup using the following software stack,

1. Confluence Standalone (includes its own embedded Tomcat) for Linux
2. Java Development Kit (JDK)
3. PostgreSQL
4. Apache Web Server (optional)

It is recommended that you finish reading the article before starting.

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-3.5.7-std.tar.gz  
tar -xvpf confluence-3.5.7-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-3.5.7-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-3.5.7-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 and not a JRE. This is a [requirement](#) of Confluence.

Go to the [Java website](#), choose the 64-bit JRE (Java Runtime Engine). Download the **.tar.gz** file and upload the file to the serveradmin home folder. For older versions of Java go to the [Java Archives site](#).

Use the account that will be launching the Java process. In this example it will be serveradmin,

```
su - serveradmin # If you are not already serveradmin  
cd ~  
wget --no-check-certificate --no-cookies --header "Cookie:  
oraclelicense=accept-securebackup-cookie"  
http://download.oracle.com/otn/java/jdk/7u7-b10/jre-7u7-linux-x64.tar.gz
```

The steps for JRE and JDK are both the same. Here is an example of a **JRE** setup,

```
su - serveradmin # If you are not already serveradmin  
cd ~ # Switch to the serveradmin home directory  
tar -xvpf jre-7u7-linux-x64.tar.gz
```

The result will be an uncompressed jre directory using the same name as the package. In this example the folder name would be, **jre-7u7-linux-x64**.

If you plan to use multiple versions of Java, we recommend keeping the folder name with the version number information and using symbolic links. If you are only using one version of Java, then simply rename the folder.

For the server example, we will rename the folder,

```
mv jre-7u7-linux-x64 java
```

You may be interested in how to [Zero Footprint Java on Windows](#).

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 **confluence/bin/setenv.sh** file and test with **confluence/bin/version.sh** 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 Confluence to Only Start and Stop with ServerAdmin

To prevent read errors and for security purposes we want Confluence to only run as serveradmin.

Modify **/opt/confluence/bin/startup.sh** and **/opt/confluence/bin/shutdown.sh** as instructed in "Only Allow serveradmin to Run Tomcat" at [Portable Tomcat 6.x & Instances](#).

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 Confluence's wiki on setting up an external database.

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

First log into the PostgreSQL prompt as explained 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'  
NOCREATEDB  
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](#) 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://krypton:8080> 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. Makes sure **PostgreSQL** is selected (it should be by default) click **External Database**.
3. Click **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 confluencedb user
 - d. Click **Next**

At this point if you happen to be monitoring your system you should see a spike in CPU and confluence slowly taking up more memory. Be patient as this step can take a few minutes as the tables and contents of the database are generated.

If all goes well the you will be prompted to Load Content. You can opt to load some default content, start with a blank slate or restore from backed up data if you are rebuilding from a previous installation of Confluence.

The intial 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 your 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.

Setup Confluence to Send Emails

Confluence can send email notifications for updates and also used to perform user setup and password resets.

[Setup your server to allow Email Delivery.](#)

Verify Confluence can send emails using Confluence Administration, Mail Servers and Send test email link.

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 Confluence to send emails
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 9009 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,

```
useURIValidationHack="false" URIEncoding="UTF-8" />

<Connector port="8009" protocol="AJP/1.3" URIEncoding="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](#).

Change Base Url in Confluence

Put instructions here about changing the base url if fronted by Apache.

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.

Backup Your Data

It goes without saying, backup your data. With Cloud and Virtual computing take a snap shot. With manual backups you can use Confluence's built in backup utility (provide link to instructions) or manually make the following backups,

- TAR the Confluence data home folder
- Backup the database using your database management tools

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/bin/shutdown.sh
- /opt/confluence/bin/startup.sh

Optionally, if there was a version mismatch with the jdbc driver,

- /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

Special Modifications

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

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

Plug-Ins

Also, keep a page listing any plug-ins 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/
```

System Profile

Using htop, initial setup without any spaces created is 396MB. The process itself uses VIRT=913M and RES=338M.

```
#!/bin/sh

# Bonsaiframework - Modification Start
# -----
if [ "$LOGNAME" != "svrdm" ]; then
echo "This service should only managed with the user svradm"
exit 1
fi
# -----
# Bonsaiframework - Modification End
```