

Confluence Install Cookbook

- [Requirements](#)
 - [Atlassian](#)
 - [OS and Hardware Settings](#)
 - [Mysql Settings](#)
 - [Mail](#)
- [Installation](#)
 - [Prepare database](#)
 - [Install binaries](#)
- [Postinstall](#)
 - [confluence.cfg.xml](#)
 - [setenv.sh](#)
 - [server.xml](#)
 - [Confluence running as Confluence user](#)
- [Other](#)
 - [Proxy passing](#)
 - [Time And Date Setup](#)
- [Logfiles](#)
 - [Logrotate](#)

Requirements

Atlassian

See <https://confluence.atlassian.com/display/DOC/System+Requirements> for newest requirements

OS and Hardware Settings

Virtual: Why not - the benefits are huge

Preferred OS: Ubuntu LTS 64-bit (this Cookbook is 90% Linux orientated)

Hardware: At least 2 vCPU's and at least 3 GB Ram

Mysql Settings

Preferred Database: MySQL (or Postgres)

Settings for MySQL (ref: [MySQL](#))

my.cnf

```
[mysqld]
transaction-isolation = READ-COMMITTED
log-bin=mysql-bin
binlog_format=row
default-storage-engine=innodb
max_allowed_packet=64M

[mysql]
default-character-set=utf8

[mysqldump]
max_allowed_packet=64M
```

Notice that max_allowed_packet=64M, where as [this link](#) says 32M, I have seen Gliffy updates fail with 32M

On large installations or where indexes are huge, changing innodb_lock_wait_timeout can be an option (Ref: [JIRAKB](#))

my.cnf

```
[mysqld]
innodb_lock_wait_timeout = 300
```

Mail

SMTP Access for sending mail

IMAP Access for picking up mail

Installation

Prepare database

Create the Database

```
mysql> create database confluence CHARACTER SET utf8 COLLATE utf8_bin;
mysql> GRANT SELECT,INSERT,UPDATE,DELETE,CREATE,DROP,ALTER,INDEX on confluence.* to confluence@localhost
identified by '*****';
Query OK, 0 rows affected (0.00 sec)
```

Install binaries

Download Confluence from <http://www.atlassian.com/software/confluence/download>

Download the MySQL JDBC from <http://dev.mysql.com/downloads/connector/j/>

For the binary installer, JAVA is integrated, for EAR or WAR files this must be downloaded and installed separately.

```
sh ./atlassian-confluence-5.1.1-x64.bin
```

Follow the required steps (remember path to **Confluence_Install** and **Confluence_Home**) and connect to the tomcat instance started; if the MySQL Driver can be found, restart the Tomcat.

Postinstall

This is where the tweaking comes into place, to avoid common problems:

confluence.cfg.xml

This file is found in **Confluence_Home**

Change the number of Database connections to higher than standard

```
<property name="hibernate.c3p0.max_size">50</property>
```

Make sure that the *?autoReconnect=true* is on the jdbc connection

```
<property name="hibernate.connection.url">jdbc:mysql://localhost/confluence?autoReconnect=true&
sessionVariables=storage_engine%3DInnoDB</property>
```

setenv.sh

This file is found in the **Confluence_Install/bin**

Read [Garbage Collector Performance Issues](#) for settings

Add support for UTF-8 File system by adding *-Dfile.encoding=UTF-8* to the JAVA_OPTS:

```
JAVA_OPTS="-Xms1024m -Xmx1024m -XX:MaxPermSize=512m $JAVA_OPTS -Djava.awt.headless=true -XX:NewSize=512m -Dfile.encoding=UTF-8"
export JAVA_OPTS
```

Tweaking of memory and usage comes in play here, I prefer this for a 4 GB Server:

```
JAVA_OPTS="-Xms2048m -Xmx2048m -XX:MaxPermSize=512m -Djava.awt.headless=true -XX:NewSize=700m -XX:+UseParallelGC -Dsun.rmi.dgc.client.gcInterval=900000 -Dsun.rmi.dgc.server.gcInterval=900000 -XX:+DisableExplicitGC -Dfile.encoding=UTF-8"
export JAVA_OPTS
```

For saving Garbage Collection in /pack/confluence/logs/gc.log for debugging

```
JAVA_OPTS="-Xms2048m -Xmx2048m -XX:MaxPermSize=512m -Djava.awt.headless=true -verbose:gc -Xloggc:/pack/confluence/logs/gc.log -XX:+PrintGCTimeStamps -XX:+PrintGCDetails -XX:NewSize=700m -XX:+UseParallelGC -Dsun.rmi.dgc.client.gcInterval=900000 -Dsun.rmi.dgc.server.gcInterval=900000 -XX:+DisableExplicitGC -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/backup/hdump -Dfile.encoding=UTF-8"
export JAVA_OPTS
```

If You are planning to run under the same Fully Qualified Domain - review <https://jira.atlassian.com/browse/JRA-8726>

server.xml

This file is found in [Confluence_Install/conf](#)

Binding to a fixed IP Address, add the *address=* to the connector. Also the *port=* can be changed (Under Linux only root can bond to 0-1023):

```
<Connector className="org.apache.coyote.tomcat4.CoyoteConnector" port="8080" address="10.0.0.10" minProcessors="5"
maxProcessors="75"
enableLookups="false" redirectPort="8443" acceptCount="10" debug="0" connectionTimeout="20000"
useURISValidationHack="false" />
```

To secure correct UTF-8 Handling, add *URIEncoding="UTF-8"* to the connector:

```
<Connector className="org.apache.coyote.tomcat4.CoyoteConnector" port="8080" address="10.0.0.10" minProcessors="5"
maxProcessors="75"
enableLookups="false" redirectPort="8443" acceptCount="10" debug="0" connectionTimeout="20000"
useURISValidationHack="false" URIEncoding="UTF-8" />
```

If behind a traffic Manager or Apache Proxy, add *scheme=*, *proxyName=* and *proxyPort=* to the context t (See [Apache2 Proxy Passing](#) or Proxy Passing section below):

```
<Context path="" docBase="../confluence" debug="0" reloadable="false" useHttpOnly="true" scheme="https"
proxyName="jira.example.com" proxyPort="443">
```

If the Tomcat needs to travel through symbolic links on the filesystem, add the *allowLinking="true"* to the context:

Tomcat 7.X - This is in server.xml

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

Tomcat 8.X - This is in context.xml

```
<Context>
  <Resources allowLinking="true" />
</Context>
```

If the Confluence instance need to run below root /, change the *path=* parameter:

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

Its possible to increase maxPostSize for faster page rendering - <https://confluence.atlassian.com/display/CONFKB/Slow+Page+Rendering+of+Large+Pages+Due+to+HTTP+POST+Limitations>

Confluence running as Confluence user



Please notice that only root can bind to ports below 1024, so running as a normal user requires port usage above or stuff like a Apache Proxy, StingRay Manager or other port-switching tool

Make sure Confluence is not running as root (for security reasons); look at [Confluence_Install/bin/user.sh](#) for a username:

```
# START INSTALLER MAGIC ! DO NOT EDIT !
CONF_USER="confluence" ##
# END INSTALLER MAGIC ! DO NOT EDIT !
export CONF_USER ##
```

And check the startup script /etc/init.d/confluence look like this (with the right path):

```
#!/bin/bash
# Confluence Linux service controller script
cd "/pack/confluence/bin"
case "$1" in
  start)
    ./start-confluence.sh
    ;;
  stop)
    ./stop-confluence.sh
    ;;
  restart)
    ./stop-confluence.sh
    ./start-confluence.sh
    ;;
  *)
    echo "Usage: $0 {start|stop|restart}"
    exit 1
    ;;
esac
```

Finally, make sure the confluence user has access:

```
cd Confluence_Install
sudo chown -R confluence:confluence confluence

cd Confluence_Home
sudo chown -R confluence:confluence confluence-data
```

Other

Proxy passing

There are good reasons for using an Apache or Traffic Manager in front of the Confluence Installation, some are:

- No port changing (non-root users can assign to ports below 1024)
- Use of URL Rewrite
- Use of URL Blocking
- Use of Allow/Denial
- SSL offloading/handling outside the Confluence

See my example in [Apache2 Proxy Passing](#)

Time And Date Setup

Time and Date should be set up according to

<https://confluence.atlassian.com/display/DOC/Configuring+Time+and+Date+Formats>

My formats for danish is:

Time Format	HH:mm
Date Time Format	dd-MM-yyyy HH:mm
Date Format	dd-MM-yyyy

Logfiles

Logrotate

Set up logrotate to avoid ever growing catalina.out log file. Here [Confluence_Install](#) is /opt/confluence, logs are rotated daily and kept for 7 days:

/etc/logrotate.d/confluence

```
/opt/confluence/logs/catalina.out {  
    daily  
    rotate 7  
    compress  
    copytruncate  
    delaycompress  
    missingok  
    size 10M  
    notifempty  
}
```