



Documentation for Confluence 5.9

Contents

Confluence Administrator's Guide	6
Getting Started as Confluence Administrator	7
Manage Users	9
Add and Invite Users	10
Delete or Disable Users	15
Restore Passwords To Recover Admin User Rights	16
Edit User Details	19
Change a Username	21
Managing Site-Wide Permissions and Groups	22
Confluence Groups for Administrators	22
Adding or Removing Users in Groups	26
Global Permissions Overview	28
Setting Up Public Access	32
Configuring User Directories	33
Configuring the Internal Directory	34
Connecting to an LDAP Directory	35
Configuring the LDAP Connection Pool	43
Configuring an SSL Connection to Active Directory	44
Connecting to an Internal Directory with LDAP Authentication	54
Connecting to Crowd or JIRA for User Management	60
Reverting from Crowd or JIRA applications to Internal User Management	68
Managing Multiple Directories	71
Managing Nested Groups	75
Synchronising Data from External Directories	78
Diagrams of Possible Configurations for User Management	80
User Management Limitations and Recommendations	86
Requesting Support for External User Management	90
Disabling the Built-In User Management	92
Managing Add-ons or Plugins	92
Writing User Macros	93
User Macro Template Syntax	101
Customising your Confluence Site	106
Changing the Look and Feel of Confluence	106
Customising the Confluence Dashboard	107
Changing the Site Logo	108
Customising Colour Schemes	110
Styling Confluence with CSS	112
Basic Styling Tutorial	114
Styling Fonts in Confluence	116
Working with Themes	116
Applying a Theme to a Site	117
Creating a Theme	118
Customising Site and Space Layouts	118
Upgrading Customised Site and Space Layouts	120
Working With Decorator Macros	122
Custom Decorator Templates	123
Customising a Specific Page	124
Customising the Login Page	124
Modify Confluence Interface Text	125
Customising the eMail Templates	126
Changing the Default Behaviour and Content in Confluence	127
Administering Site Templates	127
Importing Templates	127
Changing the Site Title	129
Choosing a Default Language	130
Configuring the Administrator Contact Page	131

Configuring the Site Home Page	132
Customising Default Space Content	133
Editing the Site Welcome Message	134
Integrating Confluence with Other Applications	136
Linking to Another Application	136
Configuring Workbox Notifications	136
Integrating JIRA and Confluence	141
Registering External Gadgets	142
Configuring the Office Connector	145
Managing your Confluence License	148
Managing Confluence Data	150
Database Configuration	150
Database JDBC Drivers	151
Database Setup for Oracle	152
Configuring an Oracle Datasource in Apache Tomcat	157
Database Setup for SQL Server	159
Configuring a SQL Server Datasource in Apache Tomcat	161
Database Setup For MySQL	163
Configuring a MySQL Datasource in Apache Tomcat	167
Database Setup for PostgreSQL	169
Configuring a PostgreSQL Datasource in Apache Tomcat	172
Embedded H2 Database	174
Migrating to Another Database	175
Configuring Database Character Encoding	180
Configuring database query timeout	184
Troubleshooting External Database Connections	185
Surviving Database Connection Closures	186
Site Backup and Restore	187
Production Backup Strategy	188
Configuring Backups	190
User Submitted Backup & Restore Scripts	192
Manually Backing Up the Site	193
Restoring a Site	194
Restoring a Space	196
Restoring a Test Instance from Production	197
Restoring Data from other Backups	198
Retrieving File Attachments from a Backup	198
Troubleshooting failed XML site backups	201
Troubleshooting XML backups that fail on restore	203
Attachment Storage Configuration	207
Configuring Attachment Size	209
Hierarchical File System Attachment Storage	210
Confluence Data Model	213
Finding Unused Spaces	218
Data Import and Export	219
Import a Text File	219
Configuring a Confluence Environment	220
Confluence Home and other important directories	221
Application Server Configuration	224
Configuring URL Encoding on Tomcat Application Server	224
Managing Application Server Memory Settings	225
Switching to Apache Tomcat	225
Java Policy Settings for Enterprise or Webhosting Environments	228
Web Server Configuration	228
Configuring Web Proxy Support for Confluence	228
Running Confluence behind Apache	230
General Apache Configuration Notes	231
Using Apache with mod_proxy	232
Using Apache with virtual hosts and mod_proxy	237
Using Apache with mod_jk	239
Using mod_rewrite to Modify Confluence URLs	241
Configuring Apache to Cache Static Content via mod_disk_cache	242

Starting Confluence Automatically on System Startup	243
Start Confluence Automatically on Linux	243
Start Confluence Automatically on Windows as a Service	247
Configuring Confluence	249
Viewing System Information	250
Live Monitoring Using the JMX Interface	250
Tracking Customisations Made to your Confluence Installation	252
Viewing Site Statistics	253
Viewing System Properties	255
Configuring the Server Base URL	255
Configuring the Confluence Search and Index	256
Configuring Indexing Language	256
Configuring Quick Navigation	257
Content Index Administration	257
Enabling OpenSearch	260
Rebuilding the Ancestor Table	260
Setting Up Confluence to Index External Sites	260
Setting Up an External Search Tool to Index Confluence	261
Configuring Mail	261
Configuring a Server for Outgoing Mail	261
Setting Up a Mail Session for the Confluence Distribution	263
Configuring the Recommended Updates Email Notification	264
The Mail Queue	264
Configuring Character Encoding	265
Troubleshooting Character Encodings	267
"€" Euro character not displaying properly	270
MySQL 3.x Character Encoding Problems	270
Other Settings	271
Configuring a WebDAV client for Confluence	271
Configuring HTTP Timeout Settings	276
Configuring Number Formats	276
Configuring Shortcut Links	277
Configuring Time and Date Formats	278
Enabling the Remote API	278
Enabling Threaded Comments	279
Installing a Language Pack	279
Installing Patched Class Files	280
Configuring System Properties	281
Recognised System Properties	291
Working with Confluence Logs	299
Configuring Logging	301
log4j Logging Levels	303
Troubleshooting SQL Exceptions	303
Configuring Confluence Security	304
Confluence Security Overview and Advisories	305
Confluence Cookies	307
Configuring Secure Administrator Sessions	309
Using Fail2Ban to limit login attempts	310
Securing Confluence with Apache	312
Using Apache to limit access to the Confluence administration interface	312
Trackback and External Referrers	314
Best Practices for Configuring Confluence Security	317
Hiding the People Directory	318
Configuring Captcha for Spam Prevention	318
Hiding External Links From Search Engines	319
Configuring Captcha for Failed Logins	319
Configuring XSRF Protection	321
User Email Visibility	322
Anonymous Access to Remote API	322
Running Confluence Over SSL or HTTPS	323
Connecting to LDAP or JIRA applications or Other Services via SSL	328
Configuring RSS Feeds	329

Preventing and Cleaning Up Spam	330
Scheduled Jobs	332
Configuring the Whitelist	336
Configuring the Time Interval at which Drafts are Saved	337
Performance Tuning	337
Cache Performance Tuning	340
Cache Statistics	343
Memory Usage and Requirements	344
Requesting Performance Support	345
Access Log Scripts	349
Troubleshooting Slow Performance Using Page Request Profiling	349
Identifying Slow Performing Macros	352
Compressing an HTTP Response within Confluence	353
Garbage Collector Performance Issues	354
Cluster Troubleshooting	355
Data Collection Policy	358

Confluence Administrator's Guide

About the Confluence Administrator's Guide

This guide covers features and functions that are only available to administrators.

For information on creating and administering spaces, See [Spaces](#).

This guide assumes that you are using the Confluence default theme. If your Confluence site has been customised the header may look different, and menu items appear in different locations to the examples given in this guide.

- [Getting Started as Confluence Administrator](#)
- [Manage Users](#)
 - [Add and Invite Users](#)
 - [Delete or Disable Users](#)
 - [Restore Passwords To Recover Admin User Rights](#)
 - [Edit User Details](#)
 - [Change a Username](#)
 - [Managing Site-Wide Permissions and Groups](#)
 - [Configuring User Directories](#)
- [Managing Add-ons or Plugins](#)
- [Writing User Macros](#)
 - [User Macro Template Syntax](#)
- [Customising your Confluence Site](#)
 - [Changing the Look and Feel of Confluence](#)
 - [Changing the Default Behaviour and Content in Confluence](#)
- [Integrating Confluence with Other Applications](#)
 - [Linking to Another Application](#)
 - [Configuring Workbox Notifications](#)
 - [Integrating JIRA and Confluence](#)
 - [Registering External Gadgets](#)
 - [Configuring the Office Connector](#)
- [Managing your Confluence License](#)
- [Managing Confluence Data](#)
 - [Database Configuration](#)
 - [Site Backup and Restore](#)
 - [Attachment Storage Configuration](#)
 - [Confluence Data Model](#)
 - [Finding Unused Spaces](#)
 - [Data Import and Export](#)
 - [Import a Text File](#)
- [Configuring a Confluence Environment](#)
 - [Confluence Home and other important directories](#)
 - [Application Server Configuration](#)
 - [Web Server Configuration](#)
 - [Starting Confluence Automatically on System Startup](#)
- [Configuring Confluence](#)
 - [Viewing System Information](#)
 - [Configuring the Server Base URL](#)
 - [Configuring the Confluence Search and Index](#)
 - [Configuring Mail](#)
 - [Configuring Character Encoding](#)
 - [Other Settings](#)

Downloads



[Download the Confluence documentation in PDF format.](#)

Other Resources

[Confluence Installation and Upgrade Guide](#)

[Confluence Knowledge Base](#)

[Atlassian Answers](#)

- [Configuring System Properties](#)
- [Working with Confluence Logs](#)
- [Configuring Confluence Security](#)
- [Scheduled Jobs](#)
- [Configuring the Whitelist](#)
- [Configuring the Time Interval at which Drafts are Saved](#)
- [Performance Tuning](#)
 - [Cache Performance Tuning](#)
 - [Memory Usage and Requirements](#)
 - [Requesting Performance Support](#)
 - [Troubleshooting Slow Performance Using Page Request Profiling](#)
 - [Compressing an HTTP Response within Confluence](#)
 - [Garbage Collector Performance Issues](#)
- [Cluster Troubleshooting](#)
- [Data Collection Policy](#)

Getting Started as Confluence Administrator

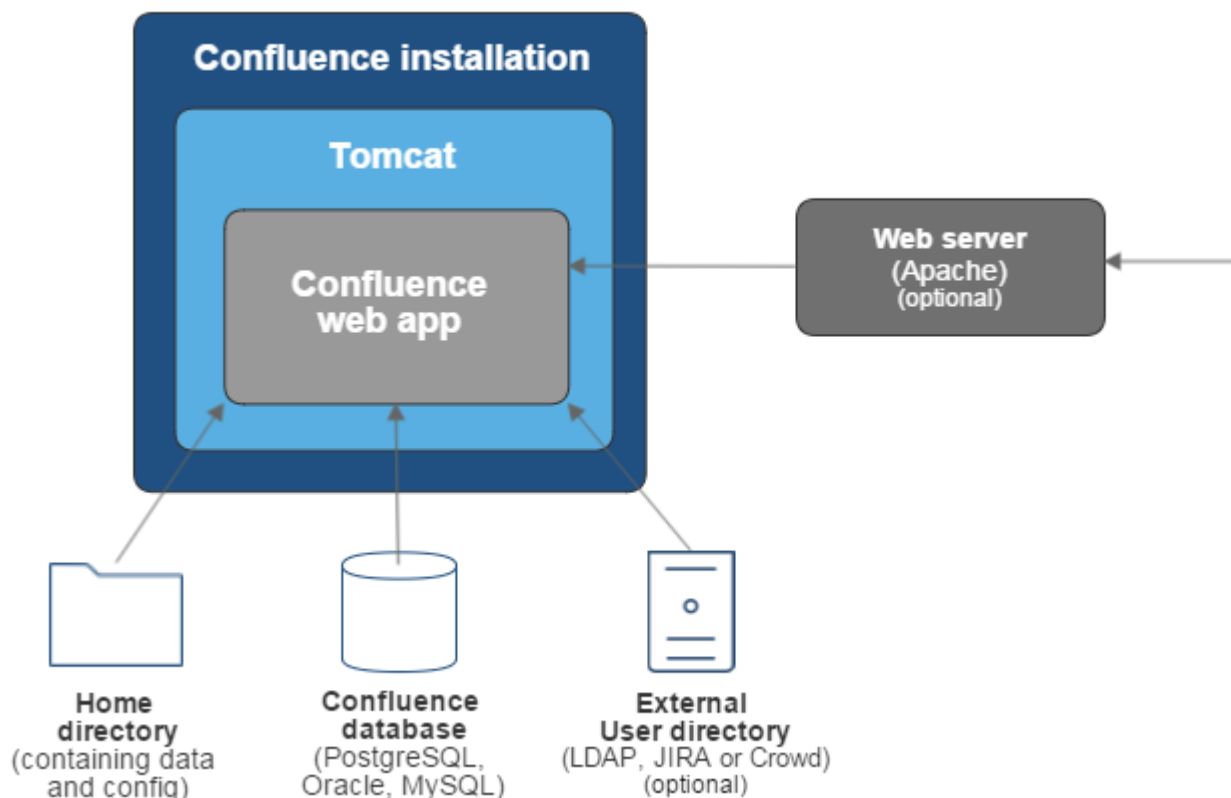
If you're just starting out as Confluence administrator, this page is for you. You'll find this page useful if your Confluence site is brand new, or if you're learning to administer an existing site.

Confluence is a Java-based web application. For the [supported environments](#), there is an installer that will set up an application server and copy the application files to the designated directories on your server machine. If you prefer, you can install Confluence from a zip file. See the [Confluence Installation Guide](#) for details.

On this page:

- [Quick access to administrative functions via Confluence search](#)
- [How to administer and configure Confluence](#)
- [Getting started on a new Confluence site](#)
- [Getting to know an existing Confluence site](#)

Diagram: A Confluence installation

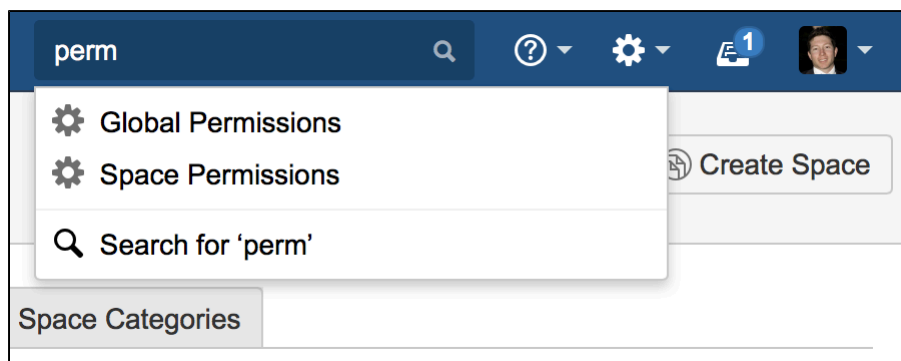


Quick access to administrative functions via Confluence search

Quick tip for getting to administration screens: Start typing what you want to do into the Confluence search box at top right of the screen. The matching administrative functions will appear with a cog icon at the top of the dropdown search results.

Even faster via /: Press / on your keyboard then continue typing the action you want. Notes about finding admin options via quick search:

- Pressing / puts your cursor in the search field (as it does in JIRA applications).
- System admin, Confluence admin, and space admin options may appear in the search results.
- Confluence permissions determine the admin options that appear in search results. You'll only see the options you're allowed to perform.



How to administer and configure Confluence

After installing Confluence, you will perform the initial configuration via a web interface called the [Confluence Setup Wizard](#). **Introducing the Confluence Administration Console:** From this point onwards, many of the

admin functions are available from the Confluence Administration Console, which is part of the Confluence web interface. If you have administrative permissions, you'll have access to the Confluence Administration Console via your web browser, using the standard Confluence URL for your site.

To access the Confluence Administration Console:

1. Open your Confluence URL in your web browser
2. Choose



> **General Configuration** in the header

For further configuration options, you can edit the XML and properties files that are part of your Confluence installation directory. To get started, take a look at the [Confluence Home and other important directories](#). The Confluence administration guide will lead you through tasks such as [configuring the log files](#) and [configuring system properties](#).

Getting started on a new Confluence site

Is this a new Confluence site? Here are some things to get started with:

- Decide whether you want to allow public (anonymous) access to your site. See [Setting Up Public Access](#).
- Add a space and some content. See [Create a Space](#) then [Pages and Blogs](#).
- Invite some users to your site. See [Add and Invite Users](#).
- Decide whether you will manage your users in Confluence or hook up an external LDAP directory. See [Configuring User Directories](#).
- Make sure you have set up an email server. The above task list will include this step, but it is worth mentioning it here again. Email notifications are an important part of collaborating on Confluence. See [Configuring a Server for Outgoing Mail](#).

Now you can continue getting to know your site, as described in the next section.

Getting to know an existing Confluence site

Has the site been around a while, but you are new to Confluence administration? Take a look at these topics:

- Understand the Confluence permission scheme. See [Permissions and Restrictions](#).
- Get to know the power of add-ons (also called plugins), for extending and customising your Confluence site. See [Managing Add-ons or Plugins](#).
- Investigate more ways of customising Confluence. See [Customising your Confluence Site](#).

Manage Users

A Confluence user is a person who can read or update a Confluence site. You can choose whether your Confluence site is accessible to anonymous users (people who have not logged in) or only to logged-in users. See [Setting Up Public Access](#).

Confluence user management

You can add users to Confluence, and then assign them permissions that determine their access to the content and administrative functions in your Confluence site. You can also collect users into groups, and assign the permissions to groups for easier management. See the following topics:

- [Add and Invite Users](#)
- [Delete or Disable Users](#)
- [Managing Site-Wide Permissions and Groups](#)

On this page:

- [Confluence user management](#)
- [Authentication](#)
 - [Seraph](#)
 - [XML-RPC and SOAP authentication](#)
 - [Password authentication](#)
- [Earlier user management frameworks](#)

Related pages:

- [Configuring Confluence Security](#)

By default, Confluence stores its users and groups in the Confluence database. This is called the internal directory. You can choose to connect Confluence to an external userbase instead, such as Microsoft Active Directory or another LDAP server. You can also use Atlassian [Crowd](#) and [JIRA](#) applications as directory managers. When you add a user or group to Confluence, it will be added to the external directory too, based on your configuration options. See [Configuring User Directories](#).

Authentication

Seraph

Almost all authentication in Confluence (and JIRA applications) is performed through [Seraph](#), Atlassian's open source web authentication framework. The goal of Seraph is to provide a simple, extensible authentication system that we can use on any application server.

Seraph is implemented as a servlet filter. Its sole job is, given a web request, to associate that request with a particular user (or no user if the request is anonymous). It supports several methods of authentication, including HTTP Basic Authentication, form-based authentication, and looking up credentials already stored in the user's session.

Seraph itself performs no user management functions. It merely checks the credentials of the incoming request and delegates any user management functions (looking up a user, checking a user's password) to Confluence's user management system.

If you want to integrate Confluence with your own single sign-on (SSO) infrastructure, you would do so by installing [Atlassian Crowd](#) or by writing a [custom Seraph authenticator](#). See our developer documentation on [HTTP authentication with Seraph](#).

XML-RPC and SOAP authentication

Normally, requests for [Confluence's remote API](#) will include an authentication token as the first argument. With this method of authentication, XML-RPC and SOAP authentication requests are checked directly against the user management framework, and tokens are assigned directly by the remote API subsystem. These requests do not pass through Seraph authenticators.

However, if the token argument is blank, Seraph will be used as a fallback authentication method for remote API requests. So, to use a custom Seraph authenticator with XML-RPC or SOAP requests, ensure that you pass an empty string as the authentication token to remote API methods.

Password authentication

By default, password authentication is delegated from Seraph to the user management system. This is not necessary, however. Single sign-on systems may have no password authentication at all, and get all the necessary credentials from the SSO provider.

Earlier user management frameworks

- **Atlassian-User – now behind the scenes.** Atlassian-User is a user and group management framework developed by Atlassian. It provides user, group and profile management services to Confluence. In earlier versions of Confluence, you needed to configure your user directories by editing the `atlassian-user.xml` file directly. In Confluence 3.5 and later this is no longer necessary, nor is it possible. Please refer to the documentation for [Confluence 3.4](#) or earlier, if you need details of this framework. Refer to the [Confluence 3.5 Upgrade Notes](#) for details of the automatic migration that will occur during the upgrade process. *Not applicable to Confluence Cloud.*
- **OSUser – obsolete.** [OpenSymphony User](#) was Confluence's core user management framework before Atlassian-User. Please refer to the documentation for Confluence 3.4 or earlier, if you need details of this framework.

Add and Invite Users

There are a number of ways to add users to Confluence:

- **By user signup:** If [user signup](#) is enabled on your Confluence site,

- people can add themselves as users of the site.
- **Via an invitation link:** You can [invite people to sign up](#) by sending them an invitation link. You can copy and paste the link, or prompt Confluence to send the link in an email message.
- **By adding users manually:** If you have Administrator or System Administrator [permission](#), you can [manually add new users](#).
- **Via an external user directory:** See [Configuring User Directories](#).

You may also be interested in information about [allowing anonymous users access to your site](#). Anonymous users don't count against your Confluence license totals.

Allow user signup

If you enable user signup, a 'Sign Up' option will appear on the Confluence screens. The option will be on the login screen, and also in the header on public sites. People can choose the option to create their own usernames on Confluence.

On this page:

- [Allow user signup](#)
- [Manage user signup notifications](#)
- [Invite people to sign up](#)
- [Reset the invitation link](#)
- [Add users manually](#)
- [Notes](#)

Related pages:


- [Manage Users](#)
- [Setting Up Public Access](#)
- [Configuring a Server for Outgoing Mail](#)

You can restrict the signup to people whose email addresses are within a given domain or domains. This is useful if you want to ensure that only people within your organisation can add their own usernames.

You will still be able to add or invite users manually, whether user signup is enabled or not.

You need Confluence Administrator or System Administrator [permissions](#) to change the signup options.

To set the user signup options:


1. Choose 
 - > **User management**
2. Select the **User Signup Options** tab
3. Choose **Allow people to sign up to create their account**
4. Choose one of the following options:
 - **Restricted by domain(s)** – Note: You need to set up a [mail server](#) for Confluence before you can configure domain restricted signup. When you choose this option, you'll see a text box. Enter one or more domains, separated by commas. People will only be able to sign up if their email address belongs to one of the domains specified here. Confluence will send the person an email message, asking them to click a link to confirm their email address. For example: [mydomain.com](#), [mydomain.net](#)
 - **No restrictions** – Anyone will be able to sign up to Confluence. Confluence will not send any email message requesting confirmation.
5. Choose **Notify administrators by email when an account is created** if you want Confluence to send an email message to all administrators (people with Confluence Administrator or System Administrator permissions) every time someone signs up to Confluence

Manage user signup notifications

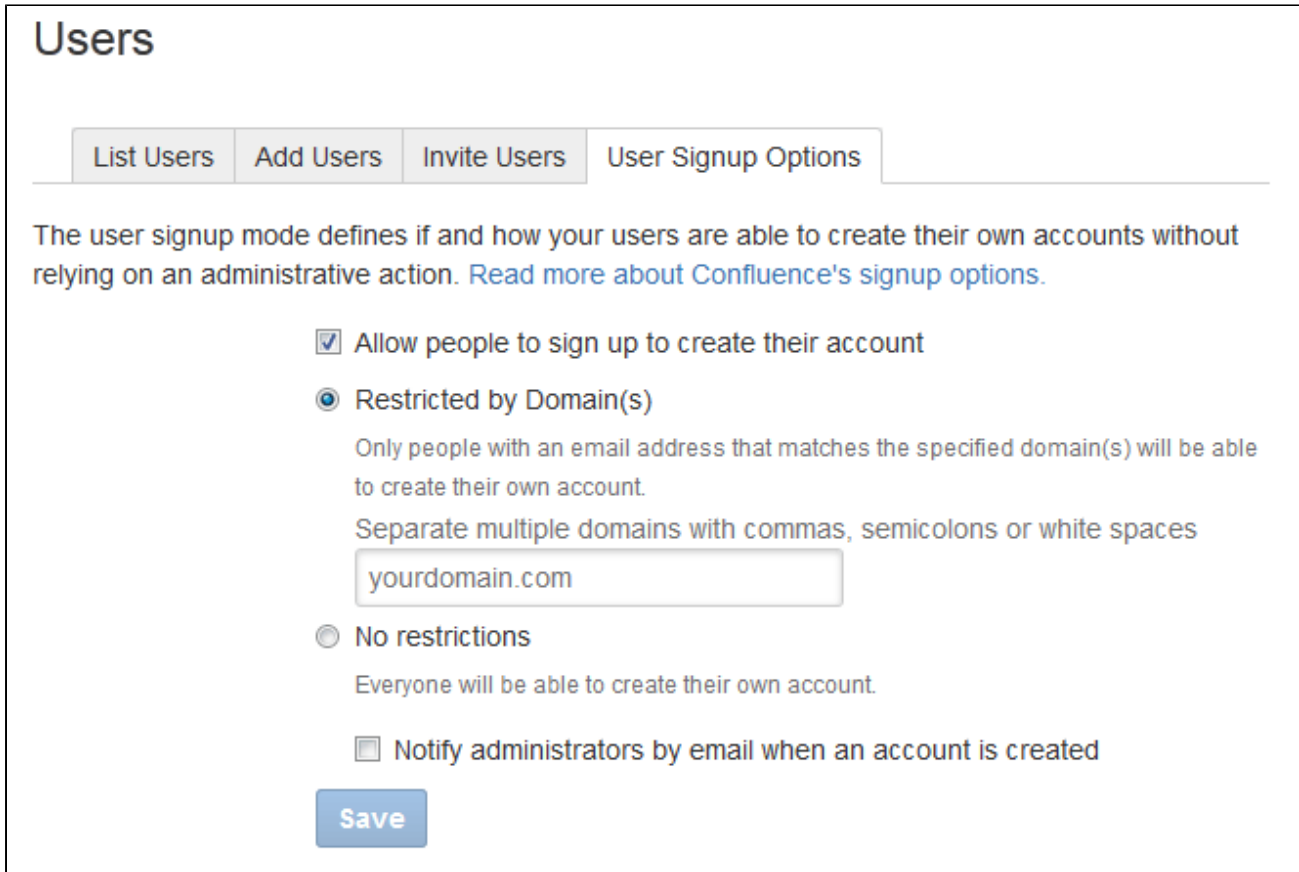
By default, Confluence will send an email notification to all Confluence administrators whenever someone signs up to your Confluence site. The administrators (people with Confluence Administrator or System Administrator permissions) will receive this message when someone signs up either by clicking the 'Sign Up'

link or by clicking the invitation URL sent by an administrator.

To disable this notification:

1. Choose  > **User management**
2. Select the **User Signup Options** tab
3. Remove the tick from **Notify administrators by email when an account is created**
4. Choose **Save**

Screenshot: User signup options




Invite people to sign up

You can invite new users to the site by sending them a signup URL, called an 'invitation link'. You can copy the invitation link and paste it onto a page or into an email message, or you can prompt Confluence to send an email message containing the same link.

The option to send invitations is independent of the signup options. You can send invitations if signup is open to all, restricted by domain, or disabled entirely. Even if signup is restricted or disabled, a person who has received an invitation will be able to sign up.

When someone visits the invitation link in a browser, a Confluence signup screen will appear.

To invite people to sign up:

1. Choose  > **User management**
2. Select the **Invite Users** tab
3. Do either of the following:
 - Copy the **Invitation Link** and paste it into an email message, or onto a page on your intranet, for example

- Alternatively, prompt Confluence to send an email message for you:
 - a. Enter one or more email addresses in the field labelled **Email To**
Separate the addresses with commas. For example: `john@example.com, sarah@example.com`
 - b. Change the **Message** if you want to
 - c. Choose **Send**

Reset the invitation link

The invitation link includes a security token, like this:

```
http://confluence.example.com/signup.action?token=d513a04456312c47
```

This security token is a shared token – individual invitations don't have unique tokens. Anyone who obtains this token will be able to sign up to Confluence.

You can change the token at any time, by choosing **Reset**. The previous invitation link will then become unusable.

Screenshot: Inviting users

Users

[List Users](#) [Add Users](#) [Invite Users](#) [User Signup Options](#)

Inviting users is the easiest way to get your team collaborating on Confluence. Just share the link below with your team, either by copying and pasting it, or by using the email option below.

Invitation Link <http://localhost:8090/signup.action?token=25bee249fe9fe7e1>

[Reset](#)

Resetting the invitation URL will stop people from signing up via a previously-sent URL.

Email To

Message

Hi,

Join me and the rest of the team on Sample Site!

It's the one place to create, share, discuss and store our ideas, projects and documents.

You will send less email, have fewer meetings and get more done.

Please accept this invitation to get started!

[Send](#)

Add users manually

To add a new user:

1. Choose



> **User management**

2. Select the **Add Users** tab
3. Enter the user's details
4. Choose whether Confluence should send an **email** message informing the person of their new username
The email message will contain a link that the person can use to reset their password.
5. Choose **Create**

Screenshot: Adding users

Users

List Users Add Users Invite Users User Signup Options

Add a User

Username*
Usernames must be lower case.

Full Name*

Email*

☐ Send an email to the user you have just created, which will allow them to set up their password.

Password*

Confirm Password*

Add

Notes

- **Multiple directories** – You can define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you could use the default Confluence **internal directory** and connect to an **LDAP** directory server. In that case, you can define the **directory order** to determine where Confluence looks first when processing users and groups.

Here is a summary of how the directory order affects the processing:

- The order of the directories is the order in which they will be searched for users and groups.
- Changes to users and groups will be made only in the first directory where the application has permission to make changes.

See [Managing Multiple Directories](#).

- **Email server required for domain restricted signup and for invitations** – You need to set up a [mail server](#) for Confluence, before you can configure domain restricted signup or send email invitations to

users.

- **Are the user management options not visible?** If you have external user management turned on, internal user management is disabled. To configure external user management, go to



> **General Configuration** > **Security Configuration**. See [Disabling the Built-In User Management](#).

Delete or Disable Users

You can **delete** a user from Confluence if they haven't yet added or edited any content on the site. Content includes pages and blog posts, and edits and comments on existing pages.

If a user has contributed content, you should **disable** their user account. Disabling a user account won't remove the content they've created.

Related pages:

- [Manage Users](#)
- [Configuring User Directories](#)

Someone with a disabled user account can no longer log in to Confluence, and they don't add to your license count.

To delete or disable a user:

1. Choose



> **User management**

2. Search for a user or click **Show all users**
3. Select the user
4. Do either of the following:
 - Choose **Delete** if the user hasn't contributed any content
 - Choose **Disable** to deactivate a user account if the user has contributed content

Screenshot: Administering a user

View User: ewan

« [Back to Users](#)

User	ewan
Full Name	Ewan User
Email	sample@email.com.au
Directory	Crowd Repository
Created	Apr 24, 2013 10:36
Last Updated	Apr 24, 2013 10:36
Login	Last Login: Apr 26, 2013 15:21 Current Failed Login Count: 0
Groups	confluence-users

[View Profile](#)
[Edit Groups](#)
[Edit Details](#)
[Set Password](#)
[Delete](#)
[Disable](#)

Notes

- The **Administer User** link is only visible if you are logged in as an administrator.
- You can also delete or disable users using the [Administration Console](#).
- You can [edit the groups](#) that a user belongs to, to change their permissions without completely preventing their access to Confluence.
- **Multiple user directories** – You can define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you could use the default Confluence **internal directory** and connect to an **LDAP** directory server. In that case, you can define the **directory order** to determine where Confluence looks first when processing users and groups.

Here is a summary of how the directory order affects the processing:

- The order of the directories is the order in which they will be searched for users and groups.
- Changes to users and groups will be made only in the first directory where the application has permission to make changes.

See [Managing Multiple Directories](#).

- **Number of users and your license** – The Confluence '[License Details](#)' screen tells you how many users your Confluence instance is licensed to support, and how many are currently registered. The number of registered users includes only users who have the 'Can Use' global permission. Deactivated users, as described above, aren't included. Choose **Refresh** to make sure you see the latest count.
- If the user resides in a Read Only external directory, the option to disable the user won't appear.

Restore Passwords To Recover Admin User Rights

If you're unable to log in to Confluence as an administrator (for example, you've lost the administrator password) you can follow these steps to recover admin user rights.

These instructions will not work for you if:

- **Confluence is configured for SSO through Crowd.**
These instructions cover how to recover administration rights from the local 'Confluence Internal Directory' only. You won't be able to authenticate as a local Confluence administrator while Crowd SSO is enabled. See [Integrating Crowd with Atlassian Confluence](#) for info on how to configure or disable Crowd SSO.
- **You're using Confluence 3.4 or earlier.**
Please refer to the older documentation if you're still using [OSUser](#) or [AtlassianUser](#).

On this page:

- [Before you Start](#)
- [Step 1. Identify Administrator](#)
- [Step 2. Replace Administrator Password](#)
- [Step 3. Put the Internal Directory in First Position](#)
- [Step 4. Clean Up](#)

Before you Start

The following instructions include example SQL that should work on MySQL and PostgreSQL. You may need to customise the queries for other databases or for your installation.

We **strongly** recommend testing the queries on a test database before modifying your production database.

If you know the admin username with a valid email address, and you have outgoing mail configured, you can reset the password using the forgot password link instead.

We'll send a link to your admin email account to reset your password.

[Log in](#)

[Forgot your password?](#)

Get access to your database

If you're using the embedded H2 database, you can find the files containing your database in `<confluence`

-home-directory>/database. See [Embedded H2 Database](#) for information on how to connect.

If you're using an external production database, connect to the database with your normal tools. You'll need to have permission to run queries and update data in the database.

Step 1. Identify Administrator

To find out which usernames have admin privileges, connect to your database using a database admin tool such as [DBVisualiser](#). Download a database admin tool now if you don't have one installed already. Then connect to your database and retrieve the list of administrator usernames and IDs with:

```
select u.id, u.user_name, u.active from cwd_user u
join cwd_membership m on u.id=m.child_user_id join cwd_group g on
m.parent_id=g.id join cwd_directory d on d.id=g.directory_id
where g.group_name = 'confluence-administrators' and d.directory_name='Confluence
Internal Directory';
```

If there are multiple results, choose one ID/username combination to use for the following steps.

If there are no results, skip down to [If No Local Administrator Exists](#).

It's important to make sure that the "active" field contains a value of "T". Without this flag, trying to authenticate with this user is a non starter.

To set active to true run the following query replacing "<user_name>" with the username from the previous query

```
UPDATE cwd_user
SET active = 'T'
WHERE user_name = '<user_name>;'
```

If No Local Administrator Exists

There may be no administrators in your Internal Directory. If this is the case, you need to add one:

1. Add a new admin user by running:

```
insert into cwd_user(id, user_name, lower_user_name, active, created_date,
updated_date, first_name, lower_first_name, last_name, lower_last_name,
display_name, lower_display_name, email_address, lower_email_address,
directory_id, credential) values (1212121, 'admin', 'admin', 'T',
'2009-11-26 17:42:08', '2009-11-26 17:42:08', 'A. D.', 'a. d.',
'Ministrator', 'ministrator', 'A. D. Ministrator', 'a. d. ministrator',
'admin@example.com', 'admin@example.com', (select id from cwd_directory
where directory_name='Confluence Internal Directory'),
'x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1
ONxp3Y3pB37A=');

insert into user_mapping values ('2c9681954172cf5600000000000000001',
'admin', 'admin');
```

2. Add new groups by running:

```
insert into cwd_group(id, group_name, lower_group_name, active, local,
created_date, updated_date, description, group_type, directory_id)
values (
'888888','confluence-administrators','confluence-administrators','T','F','2
011-03-21 12:20:29','2011-03-21 12:20:29',NULL,'GROUP',(select id from
cwd_directory where directory_name='Confluence Internal Directory'));
insert into cwd_group(id, group_name, lower_group_name, active, local,
created_date, updated_date, description, group_type, directory_id)
values ( '999999','confluence-users','confluence-users','T','F','2011-03-21
12:20:29','2011-03-21 12:20:29',NULL,'GROUP',(select id from cwd_directory
where directory_name='Confluence Internal Directory'));
```

3. Add group memberships into cwd_membership:

```
insert into cwd_membership (id, parent_id, child_user_id) values (888888,
(select id from cwd_group where group_name='confluence-users' and
directory_id=(select id from cwd_directory where directory_name='Confluence
Internal Directory')), 1212121);
insert into cwd_membership (id, parent_id, child_user_id) values (999999,
(select id from cwd_group where group_name='confluence-administrators' and
directory_id=(select id from cwd_directory where directory_name='Confluence
Internal Directory')), 1212121);
```

If you're using an Oracle database, use **sysdate** instead of a string for the **created_date** column.

Step 2. Replace Administrator Password

Confluence doesn't store passwords in plain text in the database, but uses hashes computed from the original password. You'll need to insert a hash, rather than the plain password, over the existing password in the database. Below is the hash for the password `admin`

```
x61Ey6l2Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1ONxp3Y3
pB37A==
```

To change the password to `admin` for a given username:

1. Shut down Confluence
2. Connect to your database
3. Run the following SQL:

```
update cwd_user set credential =
'x61Ey6l2Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1
ONxp3Y3pB37A=='
where id=<id from Stage 1>;
```

Step 3. Put the Internal Directory in First Position

Start Confluence, and try logging in with the username of the user you updated/created and the password 'admin'. If this works, skip to Step 4; otherwise, your Internal Directory doesn't have high enough priority.

To put your Internal Directory in first position:

1. Find the directory names and their order:

```
select d.id, d.directory_name, m.list_index from cwd_directory d join
cwd_app_dir_mapping m on d.id=m.directory_id;
```

2. Take note of the ID with list_index 0, and the list_index and ID of the Confluence Internal Directory
3. Switch the order of the directories:

```
update cwd_app_dir_mapping set list_index = 0 where directory_id =
<Internal Directory id>;
update cwd_app_dir_mapping set list_index = <Noted Internal Directory
list_index> where directory_id = <Directory id that had list_index 0>;
```

4. Check to see if the directory is active (the 'active' column should be set to 'T'):

```
select id, directory_name, active from cwd_directory where id = <Internal
Directory id>;
```

5. If necessary, activate the directory:

```
update cwd_directory set active = 'T' where id = <Internal Directory id>;
```

Step 4. Clean Up

To tidy up:

1. Start Confluence
2. Log in with your modified/created username and use password `admin`
3. Change your password

Don't leave your password as admin; if you do, your instance won't be secure.

4. If you created a new user in Stage 2, create a new admin via the UI and delete the admin you created in Stage 2
5. If you followed Stage Three, go to



> **General Configuration** > **User Directories** and rearrange your directories so they're correctly configured again.

Edit User Details

You can view and edit the details of Confluence users, including their name, password, email address, group membership, and ability to access Confluence.

Edit a user's details

1. Choose



> **User management**

2. Do either of the following:

- Choose **Show all users** to list everyone in the

- 'confluence-users' or 'users' group
- Enter a username, full name or email address in the **Find User** field and hit **Search**

If you're already viewing someone's profile, choose **Administer User** in the sidebar.

- Select the user you want to manage

On this page:

- [Edit a user's details](#)
- [Reset login count](#)

Related pages:

- [Delete or Disable Users](#)
- [Adding or Removing Users in Groups](#)
- [Add and Invite Users](#)

Now you'll see the person's current details and links allowing you to edit them.

- View Profile** — View the user's profile.
- Edit Groups** — [Add or remove](#) this user from a group.
- Edit Details** — Change details such as the user's name, email address, contact details and team or department information. In some instances you may be able to change usernames as well. See [Change a Username](#) for information.
- Set Password** — Edit the user's password details.
- Delete** — You can [delete](#) a user permanently if the user has not added or edited any content on the site.
- Disable** — You can [disable](#) (i.e. deactivate) access for a user who has already added or edited any content on the site.

View User: juser

« [Back to Users](#)

User **juser**

Full Name **Josh User**

Email josh@user.com

Directory **Confluence Internal Directory**

Created **Feb 24, 2015 10:13**

Last Updated **Feb 24, 2015 10:13**

Login **Last Login: Jun 22, 2015 10:56**
Current Failed Login Count: 0

Groups  **confluence-users**

[View Profile](#) [Edit Groups](#) [Edit Details](#) [Set Password](#) [Delete](#) [Disable](#)

Reset login count

Confluence records the number of failed login attempts made against each user account. When the login attempts exceed a preset number, the user is prompted to authenticate using CAPTCHA until they successfully log in.

If the user you're administering has any failed login attempts, you can manually set the failed login count for a user back to zero by clicking **Reset Failed Login Count**.

Multiple user directories

You can define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you could use the default Confluence **internal directory** and connect to an **LDAP** directory server. In that case, you can define the **directory order** to determine where Confluence looks first when processing users and groups.

Here is a summary of how the directory order affects the processing:

- The order of the directories is the order in which they will be searched for users and groups.
- Changes to users and groups will be made only in the first directory where the application has permission to make changes.

See [Managing Multiple Directories](#).

Change a Username

As a Confluence administrator, you can change a user's username. This could be for any reason, but might happen when someone changes their name, for example.

Each active users must have a unique username, so no two *active* users can have the same username. You can, however, assign the username of a [disabled user](#) to another active user.

The procedure for changing a username depends on where you manage your users. See [Configuring User Directories](#) for more info.


On this page:

- [Confluence-managed users](#)
- [Users managed in an external directory](#)
- [Notes](#)

Confluence-managed users

If you manage your users in the Confluence internal directory, you can rename your user in Confluence. You'll need Confluence Administrator permissions to change a username.

To change a username:

1. Choose 
 - > **User management**
2. Search for the user or choose **Show all users**
3. Select the user you'd like to edit and choose **Edit Details**
4. Enter the new username and choose **Submit**

That person will need to use their new username to log in to Confluence from now on. The new username will also be reflected throughout Confluence, including in [@mentions](#).

Users managed in an external directory

If you don't manage your users in the Confluence internal directory, you may still be able to change someone's username. Confluence can't update external users, but it will detect changes in usernames coming from *some* external directories.

The following table shows the instances where you may be able to change a username in your external directory and have the change detected in Confluence.

User directory	Where to rename the user
Internal directory with LDAP authentication	Rename the user in the LDAP directory, Confluence will detect the renamed user. Note: you must have 'Copy User on Login' enabled. See Copying Users on Login for more information.
JIRA 6.1 or later	Rename the user in JIRA, Confluence will automatically detect the renamed user.
Atlassian Crowd 2.7 or later	Rename the user in Crowd, Confluence will automatically detect the renamed user.
LDAP	Rename the user in your LDAP directory, Confluence will automatically detect the renamed user.

Notes

Some important things to note about changing usernames:

- **Mentions and page history** – Any user mentions in current pages will automatically reflect the user's new username, but any mentions in page versions created prior to Confluence 5.3 will include the user's old username.
- **Personal Spaces** – If a Confluence Administrator renames a user who has a personal space, the space key for that space will remain as the original username. For example, if jsmith's username is changed to jbrown, their personal space key will remain ~jsmith.

Managing Site-Wide Permissions and Groups

Permissions determine what people can do on your Confluence site. Confluence recognises permissions at [site level](#) and at [space level](#), as well as [page-level restrictions](#).

You can create [groups](#) and allocate people to them, so that you can assign permissions to a number of people at once. It's quicker to give a group access to Confluence than giving every member access individually.

You can also set the access levels for [anonymous users](#) or deny access to unlicensed users from linked applications, such as JIRA Service Desk.

Related pages:

- [Confluence Security Overview and Advisories](#)
- [Global Permissions Overview](#)

Confluence Groups for Administrators

[Grouping users](#) in Confluence is a great way to cut down the work required when managing permissions and restrictions. Once you have a group of users, you can assign that group a set of [global permissions](#). For example, if you don't want that group of users to be able to create spaces, you can revoke the 'Create Space(s)' permission.

Other users can also take advantage of Confluence groups. [Space admins](#) can assign a set of [space permissions](#) to a group rather than to each individual user, and other users with the 'Add/Delete Restrictions' space permission can add and remove [page restrictions](#) for groups.

Special groups

There are three special default groups in Confluence:

1. **system-administrators (sys admin)** – The user that sets up a Confluence instance will be the first sys admin; they can also appoint

other sys admins by granting them the 'System Administrator' [global permission](#). They can perform all Confluence administrative functions, including assigning permissions to other users.

2. **confluence-administrators**: This is a group of 'super-users' who can access the Confluence administration screens ('administration console') and perform site-wide administration. Members of this group can also see all spaces in the Confluence site. Any user who is a member of this group has site-wide administration powers, regardless of any other setting. The settings on the [global permissions screen](#) do **not** affect the powers allowed to members of this group.
3. **confluence-users**: This is the default group for all new users. Permissions you assign to this group will be assigned to all newly created users.

On this page:

- [Special groups](#)
- [Anonymous users](#)
- [Add or delete groups](#)
- [Confluence Administrator permission vs confluence-admin group comparison](#)
- [Notes](#)

Related pages:

- [Manage Users](#)
- [Global Permissions Overview](#)

The Confluence Administrator permission and the 'confluence-administrators' group are not related.

Going by the names, you would think the 'confluence-administrators' group and the 'Confluence Administrator' permission are related – but they are not. Granting a user or a group 'Confluence Administrator' permission is *not* the same as granting them membership of the 'confluence-administrators' group. Granting the 'Confluence Administrator' permission enables access to only a subset of the administrative functions. Granting membership to the 'confluence-administrators' group gives complete access.


[View the comparison table.](#)

Anonymous users

All users who don't log in when they access Confluence are considered 'anonymous'. You can [grant anonymous users the 'Use Confluence' permission](#) via the Global Permissions screen if you need to. This will allow non-registered users to access pages and spaces in Confluence. A space administrator can further control anonymous access per space via the [space permissions](#).

Add or delete groups

To add a new group:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration
2. Choose **Groups** in the left-hand panel
3. Choose **Add Group**
4. Enter a name for your group and choose **Save**

You're now ready to start [adding users](#) to the group.

To delete a group:

1. Choose the **cog icon**



- , then choose **General Configuration** under Confluence Administration
2. Choose **Groups** in the left-hand panel
You will see a list of all existing groups along with links to remove them.
 3. Choose **Delete** next to the group you want to remove

Confluence Administrator permission vs confluence-admin group comparison



Granting the Confluence Administrator permission to someone allows them access to many, but not all, options in the administration console (



> **General configuration**). Expand the comparison table to view the options available to people granted the Confluence Administrator permission, and to those in the confluence-admin group.

▼ [Click to view the comparison table](#)

Administration option	Confluence Admin permission	confluence-administrators group
CONFIGURATION		
General Configuration		
Further Configuration		
Manage Referrers		
Languages		
Shortcut Links		
Global Templates and Blueprints		
Import Templates		
Mail Servers		
Recommended Updates Email		
User Macros		
In-app Notifications		
HipChat Integration		
Attachment Storage		
Spam Prevention		
PDF Export Language Support		
Configure Code Macro		
WebDAV Configuration		
Office Connector		
ATLASSIAN MARKETPLACE		
Find new add-ons		
Manage add-ons		

Purchased add-ons		
USERS & SECURITY		
Users		
Groups		
Security Configuration		
Global Permissions		
Space Permissions		
User Directories		
Whitelist		
LOOK AND FEEL		
Themes		
Colour Scheme		
Layouts		
Stylesheet		
Site Logo and Favicon		
PDF Layout		
PDF Stylesheet		
Default Space Logo		
		Custom HTML
ADMINISTRATION		
System Information		
Backup & Restore		
Content Indexing		
Mail Queue		
Cache Management		
Scheduled Jobs		
License Details		
Logging and Profiling		
Thread Dump		
Application Links		
Application Navigator		
Analytics		

Import from Another Wiki		
Atlassian Support Tools		
Clustering		

Notes

- **Multiple user directories:** You can define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you could use the default Confluence **internal directory** and connect to an **LDAP** directory server. In that case, you can define the **directory order** to determine where Confluence looks first when processing users and groups.

Here is a summary of how the directory order affects the processing:

- The order of the directories is the order in which they will be searched for users and groups.
- Changes to users and groups will be made only in the first directory where the application has permission to make changes.

See [Managing Multiple Directories](#).

Adding or Removing Users in Groups

If you are a [Confluence Administrator](#), you can add users and groups, and assign users to groups, in order to determine their permissions.

This page tells you how to add a user to a group or remove a user from a group. For an overview of users and groups, please refer to [Confluence Groups](#) and [Manage Users](#).

You can edit group membership in two places:

- From the group management screen
- From the user management screen for a particular user

On this page:

- [Add and remove members via group management](#)
- [Edit group membership via user management](#)

Related pages:

- [Manage Users](#)
- [Global Permissions Overview](#)

Add and remove members via group management

This is the recommended method. It allows you to manage the group membership for a number of users at the same time.

To add members to a group:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Groups** in the left-hand panel
3. Choose the group to which you want to add users
4. Choose **Add Members**
5. Type the username(s) of the people you want to add to the group
 - If you want to add more than one member, separate the usernames with commas
 - You can also search for and select users by choosing the search icon
6. Choose **Add** to add the member(s) to the group

To remove members from a group:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Groups** in the left-hand panel
3. Choose the group from which you want to remove the user

4. Choose the **Delete user from group** icon next to the user whose group membership you want to remove

Screenshot: Adding members

Group Members: techwriters

[Groups](#) [Cancel](#)

techwriters


Add Members Enter a comma separated list of user names to add users to this group

[Cancel](#)

Edit group membership via user management

You can update a user's group membership from the user management screen. This functionality allows you to update one user at a time.

To add a user to a group or remove a user from a group:

1. Go to the user management screen for the user concerned. There are two ways to do this:
 - Either,
 - Go to the user's [Profile](#) and choose **Administer User** on the user's profile screen.
 - Or, Choose the **cog icon**

, then choose **General Configuration** under Confluence Administration
 - a. Choose **Users** in the left-hand panel
 - b. Choose **Show all users**, or search for a specific user by entering all or part of their username, full name or email address
 - c. Choose the username you want to edit
2. Choose **Edit Groups**
3. Select the group(s) for this user

To remove a user from a group, remove the tick mark in the relevant check box.

Screenshot: Editing a user's groups

Edit User Groups : connie

[« Return to view user: connie](#)

Select group memberships for **connie**

☐ confluence-administrators

☒ confluence-users

☐ developers

☐ techwriters

[Select All Groups](#) · [Deselect All Groups](#)

You can define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you could use the default Confluence **internal directory** and connect to an **LDAP** directory server. In that case, you can define the **directory order** to determine where Confluence looks first when processing users and groups.

Here is a summary of how the directory order affects the processing:

- The order of the directories is the order in which they will be searched for users and groups.
- Changes to users and groups will be made only in the first directory where the application has

permission to make changes.
See [Managing Multiple Directories](#).

Global Permissions Overview

Global Permissions determine the actions which a user is allowed to perform in Confluence at a site level. To assign global permissions to a user or group you need Confluence Administrator or greater permissions.

Note: The first system administrator is defined during initial setup. During the initial configuration of Confluence, the Setup Wizard asks for the username of the System Administrator. This user will have the 'System Administrator' permission and will be a member of the 'confluence-administrators' group.

On this page:

- [Overview of global permissions](#)
- [Comparing the System Administrator permission with the Confluence Administrator permission](#)
- [Comparing the confluence-administrators group with the administrator permissions](#)
- [Updating global permissions](#)
- [Revoking access for unlicensed users from JIRA Service Desk](#)
- [Error messages you may see](#)

Overview of global permissions

The following global permissions can be applied to groups and individuals.

Global Permission	Description
Can Use	This is the most basic permission that allows users to access the site. Users with this permission count towards the number of users allowed by your license.
Attach Files to User Profile	This allows the user to upload files to be stored in their user profile. This feature was made obsolete by the introduction of personal spaces in Confluence 2.2. Hence, this permission is no longer relevant. Attachments can be accessed from a user profile view (for example, an image within the 'About Me' field of a profile view) by attaching these files to a page within that user's personal space and referencing them using appropriate wiki markup code.
Personal Space	This permission allows the user to create a personal space.
Create Space(s)	This permission allows users to create new spaces within your Confluence site. When a space is created, the creator automatically has the 'Admin' permission for that space and can perform space-wide administrative functions.
Confluence Administrator	This permission allows users to access the 'Administration Console' that controls site-wide administrative functions. Users with this permission can perform most, but not all, of the Confluence administrative functions. See the comparison of 'System Administrator' and 'Confluence Administrator' below.
System Administrator	This permission allows users to access the 'Administration Console' that controls site-wide administrative functions. Users with this permission can perform all the Confluence administrative functions, including the ones which the 'Confluence Administrator' permission does not allow. See the comparison of 'System Administrator' and 'Confluence Administrator' below. Refer also to the note about the 'confluence-administrators' group below.

Comparing the System Administrator permission with the Confluence Administrator permission

Confluence recognises two levels of administrator:

- **System Administrator** – Users with this permission can perform all the Confluence administrative functions, including the ones which the 'Confluence Administrator' permission does not allow.
- **Confluence Administrator** – Users with this permission can perform most, but not all, of the Confluence administrative functions.

The two-tier administration is useful when you want to delegate some administrator privileges to project managers or team leaders. You can give 'Confluence Administrator' permission to users who should be able to perform most administrative functions, but should not be able to perform functions that can compromise the security of the Confluence system.

The following functions are granted to the 'System Administrator' permission but excluded from the 'Confluence Administrator' permission:

Administration Screen	Excluded from Confluence Administrator permission
General Configuration	The following functionality is disallowed: <ul style="list-style-type: none"> • Server Base URL • Public Signup • Connection Timeouts
Further Configuration	The following functionality is disallowed: <ul style="list-style-type: none"> • Remote API plugin
Security Configuration	The following functionality is disallowed: <ul style="list-style-type: none"> • External user management • Append wildcards to user and group searches • Anti XSS Mode • Enable Custom Stylesheets for Spaces • Show system information on the 500 page • Maximum RSS Items • XSRF Protection
Plugins	The following functionality is disallowed: <ul style="list-style-type: none"> • Upgrade • Install • Confluence Upgrade Check
Daily Backup Admin	This function is disallowed entirely.
Mail Servers	This function is disallowed entirely.
User Macros	This function is disallowed entirely.
Attachment Storage	This function is disallowed entirely.
Layouts	This function is disallowed entirely.
Custom HTML	This function is disallowed entirely.
Backup & Restore	This function is disallowed entirely.
Logging and Profiling	This function is disallowed entirely.

Cluster Configuration	This function is disallowed entirely.
Scheduled Jobs	This function is disallowed entirely.
Application Links	People with the 'Confluence Administrator' permission can add, modify and remove application links and project links. For example, they can link Confluence to JIRA. However, Confluence administrators can configure only OAuth authentication for application links.
Office Connector configuration	This function is disallowed entirely.

Comparing the confluence-administrators group with the administrator permissions

The 'confluence-administrators' group defines a set of 'super-users' who can access the Confluence administration console and perform site-wide administration. Members of this group can also see the content of all pages and spaces in the Confluence instance, regardless of [space permissions](#). They cannot immediately see the pages that exclude them via [page restrictions](#) without knowing the direct URL to the page. They can remove the page restrictions via the Space Administration screen if need be. For example, they will not see restricted pages displayed by the [children macro](#). But they are able to access restricted pages directly using the page URL.

The settings on the 'Global Permissions' screen do not affect the powers allowed to members of the 'confluence-administrators' group .

Granting the 'System Administrator' or 'Confluence Administrator' permission to a user will *not* automatically grant the user access to all spaces in the site. These permissions will only give access to the administration console.

Be aware, however, that users with 'System Administrator' can add themselves to the 'confluence-administrators' group and become a super-user.

The Confluence Administrator permission and the 'confluence-administrators' group are not related.

Going by the names, you would think the 'confluence-administrators' group and the 'Confluence Administrator' permission are related – but they are not. Granting a user or a group 'Confluence Administrator' permission is *not* the same as granting them membership of the 'confluence-administrators' group. Granting the 'Confluence Administrator' permission enables access to only a subset of the administrative functions. Granting membership to the 'confluence-administrators' group gives complete access.

Updating global permissions

To view the global permissions for a group or user:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Global Permissions** in the left-hand panel. The 'View Global Permissions' screen appears.

Add or edit group and user permissions as follows:

To add permissions for a group:

1. First add the group to Confluence, if you have not already done so.
2. Choose **Edit Permissions**. The 'Edit Global Permissions' screen appears.
3. Enter the group name in the **Grant browse permission to** box in the 'Groups' section. You can search for the group name.
4. Choose **Add**.

- The group will appear in the list and you can now edit its permissions.

To add permissions for a specific user:

(Consider adding the user to a group and then assigning the permissions to the group, as described above, instead of assigning permissions to the specific user.)

- First add the user to Confluence, if you have not already done so.
- Choose **Edit Permissions**. The 'Edit Global Permissions' screen appears.
- Enter the username in the **Grant browse permission to** box in the 'Individual Users' section. You can search for the username.
- Choose **Add**.
- The username will appear in the list and you can now edit its permissions.

To add or edit the permissions for a user or group:

- Select, or clear, the check box under the relevant permission in the row for the relevant user/group. A selected check box indicates that the permission is granted.
- To allow anonymous access to your Confluence site, select the 'Use Confluence' and 'View User Profile' options in the 'Anonymous Access' section.
- Choose **Save All** to save your changes.

Screenshot: Editing global permissions

Edit Global Permissions

Global permissions define what people can do in your site. You can grant permissions to groups and individuals, and choose to open up your site to anonymous users.

See [Global Permissions overview](#) to find out more about managing permissions for your entire Confluence site.

Licensed Users

Groups

Grant permissions for the entire site to all the members of a group.

	Attach Files to User Profile [?]	Personal Space [?]	Create Space(s) [?]	Confluence Administrator [?]	System Administrator [?]
administrators	<input checked="" type="checkbox"/> can use	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
confluence-users	<input checked="" type="checkbox"/> can use	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
site-admins	<input checked="" type="checkbox"/> can use	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
system-administrators	<input checked="" type="checkbox"/> can use	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Grant browse permission to

Individual Users

Grant permissions to individual users, regardless of which groups they are a member of.

No users currently have individual global permissions.

Grant browse permission to

Anonymous Access


Make your Confluence site public. You can choose which spaces anonymous users can access. Anonymous users are not included in your license count.

	Use Confluence [?]	View User Profiles [?]
Anonymous	<input type="checkbox"/> can use	<input type="checkbox"/>

Revoking access for unlicensed users from JIRA Service Desk

If you're using Confluence as a knowledge base for JIRA Service Desk, you can choose to allow all active users and customers (that is logged in users who do not have a Confluence license) to view pages in specific spaces. This permission can only be turned on via JIRA Service Desk.

To revoke access for unlicensed users:

1. Go to  **> General Configuration > Global Permissions.**
2. Choose **Edit Permissions**
3. Deselect the '**Can Use**' permission under **Unlicensed Access.**

Unlicensed users will no longer be able to access pages in your Confluence site. This can only be re-enabled via JIRA Service Desk.

You can also choose to revoke access for individual spaces from the Space Permissions screen in each space.

Error messages you may see

Confluence will let you know if there is a problem with some permissions. In rare situations, you may see the following error messages below a permission:

- 'User/Group not found' - This message may appear if your LDAP repository is unavailable, or if the user/group has been deleted after the permission was created.
- 'Case incorrect. Correct case is: xxxxxx' - This message may appear if the upper/lower case in the permission does not match the case of the username or group name. If you see a number of occurrences of this message, you should consider running the [routine supplied to fix the problem](#).

Setting Up Public Access


You can enable anonymous access (also known as public access) to your Confluence site by granting the 'Use Confluence' permission to 'anonymous' users. An 'anonymous' user is someone who has not logged in to the Confluence site. The 'Use Confluence' permission is also called 'can use'.

This user category gives you an easy way to administer users who have not logged into the site. Permissions assigned to this category apply to all anonymous users of the site.

Enabling anonymous access to the site

If you want to make your site visible to everyone, including people who have not logged in, you must enable anonymous access at site level.

To enable anonymous access to your site:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration
2. Choose **Global Permissions** in the left-hand panel.
3. Choose **Edit Permissions**.
4. In the 'Anonymous Access' section, select the **can use** check box to enable anonymous access to the content on your site.
5. If you want to allow anonymous users to see [user profiles](#), select the check box in the **View User Profiles** section.
Note: You must grant the 'can use' permission as well, if you want to grant the 'View User Profiles' permission.
6. Choose **Save All**.

Disabling anonymous access to the site

On this page:

- [Enabling anonymous access to the site](#)
- [Disabling anonymous access to the site](#)
- [Granting public access to a space](#)
- [Notes](#)

Related pages:

- [Configuring Captcha for Spam Prevention](#)
- [Add and Invite Users](#)
- [Global Permissions Overview](#)

To disable anonymous access to your site, deselect the **can use** check box, then choose **Save All**. People will not be able to see the content on the site until they have logged in.

Granting public access to a space

To enable public access to a Confluence space, you must grant the following permissions to anonymous users:

- The site-wide 'can use' permission, as described above.
- The relevant [space permissions](#). If you want a space to be publicly accessible, the anonymous user must have at least the 'View Space' permission. To set space permissions, choose **Browse > Space Admin > Permissions**.

Notes

- We **severely warn against** giving anonymous users any administrative privileges, either within a space, or especially over the Confluence site. Giving administrative privileges to untrusted users may lead to a serious security compromise of your site.
- You can allow people to sign up for usernames themselves, and choose other options for user signup and invitations. See [Add and Invite Users](#).

Configuring User Directories

A user directory is a place where you store information about users and groups. User information includes the person's full name, username, password, email address and other personal information. Group information includes the name of the group, the users that belong to the group, and possibly groups that belong to other groups.

The **internal** directory stores user and group information in the Confluence database. You can also connect to **external** user directories, and to Atlassian **Crowd** and **JIRA** applications as directory managers.

Configuring User Directories in Confluence

To configure your Confluence user directories:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Click '**User Directories**' in the left-hand panel.

On this page:

- [Configuring User Directories in Confluence](#)
- [Connecting to a Directory](#)
- [Updating Directories](#)

Related pages

- [Configuring the Internal Directory](#)
- [Connecting to an LDAP Directory](#)
- [Connecting to an Internal Directory with LDAP Authentication](#)
- [Connecting to Crowd or JIRA for User Management](#)
- [Managing Multiple Directories](#)
- [Managing Nested Groups](#)
- [Synchronising Data from External Directories](#)
- [Diagrams of Possible Configurations for User Management](#)
- [User Management Limitations and Recommendations](#)
- [Requesting Support for External User Management](#)
- [Disabling the Built-In User Management](#)
- [Add and Invite Users](#)
- [Managing Site-Wide Permissions and Groups](#)

Connecting to a Directory

You can add the following types of directory servers and directory managers:

- Confluence's internal directory. See [Configuring the Internal Directory](#).
- Microsoft Active Directory. See [Connecting to an LDAP Directory](#).
- Various other LDAP directory servers. See [Connecting to an LDAP Directory](#).

- An LDAP directory for delegated authentication. See [Connecting to an Internal Directory with LDAP Authentication](#).
- Atlassian Crowd or JIRA 4.3 or later. See [Connecting to Crowd or JIRA for User Management](#).

You can add as many external user directories as you need. Note that you can define the **order** of the directories. This determines which directory Confluence will search first, when looking for user and group information. See [Managing Multiple Directories](#).

Updating Directories

Limitations when Editing Directories

You cannot edit, disable or remove the directory your user belongs to. This precaution is designed to prevent administrators from locking themselves out of the application by changing the directory configuration in a way that prevents them logging in or removes their administration permissions.

This limitation applies to all directory types. For example:


- You cannot disable the internal directory if your user is an internal user.
- You cannot disable or remove an LDAP or a Crowd directory if your user comes from that directory.

In some situations, reordering the directories will change the directory that the current user comes from, if a user with the same username happens to exist in both. This behaviour can be used in some cases to create a copy of the existing configuration, move it to the top, then remove the old one. Note, however, that duplicate usernames are not a supported configuration.

You cannot remove the internal directory. This precaution aligns with the recommendation below that you always keep an administrator account active in the internal directory.

Recommendations

The recommended way to edit directory configurations is to log in as an internal user when making changes to external directory configuration.

 We recommend that you keep either an administrator or system administrator user active in your internal directory for troubleshooting problems with your user directories.

Enabling, Disabling and Removing Directories







You can enable or disable a directory at any time. If you disable a directory, your configuration details will remain but the application will not recognise the users and groups in that directory.

You have to disable a directory before you can remove it. Removing a directory will remove the details from the database.

User Directories

The table below shows the user directories currently configured for Confluence.

The order of the directories is the order in which they will be searched for users and groups. Changes to users and groups will be made in the first directory where Confluence has permission to make changes. It is recommended that users only exist in a single directory.

Directory Name	Type	Order	Operations
Confluence Internal Directory	Internal	 	
OpenLDAP	OpenLDAP (Read-Write)	 	Disable Edit Synchronise Last synchronised at 14/01/11 3:07 PM (took 65s).
Crowd	Atlassian Crowd	 	Disable Edit Synchronise Last synchronised at 14/01/11 2:39 PM (took 0s).

[Add Directory](#)

Additional Configuration

- [LDAP Connection Pool Configuration](#)

Screenshot above: Configuring user directories

Configuring the Internal Directory

The internal directory stores user and group information in the Confluence database.

Overview

The internal directory is enabled by default at installation. When you create the first administrator during the setup procedure, that administrator's username and other details are stored in the internal directory.

If needed, you can configure one or more additional user directories. This is useful if you want to grant access to users and groups that are stored in a corporate directory or other directory server.

Diagram of Possible Configuration

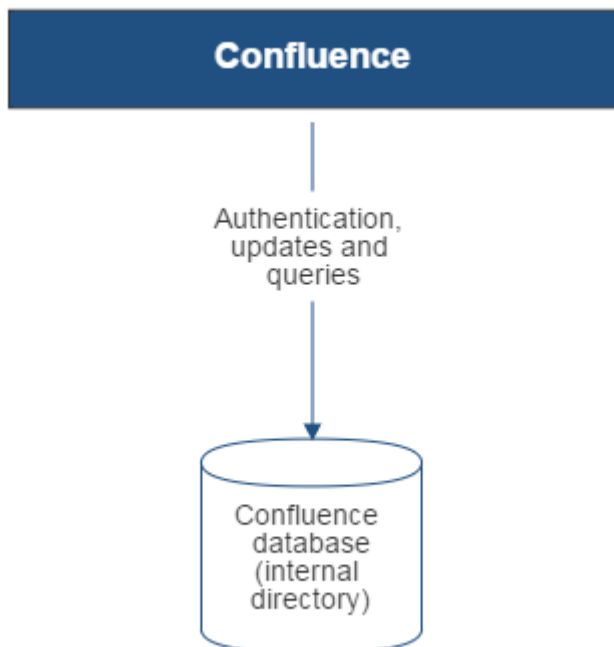


Diagram above: Confluence using its internal directory for user management.

Connecting to an LDAP Directory

You can connect your Confluence application to an LDAP directory for authentication, user and group management.

Overview

An LDAP directory is a collection of data about users and groups. LDAP (Lightweight Directory Access Protocol) is an Internet protocol that web applications can use to look up information about those users and groups from the LDAP server.

We provide built-in connectors for the most popular LDAP directory servers:

On this page:

- [Overview](#)
- [Diagram of Possible Configuration](#)

Related pages:

- [Configuring User Directories](#)
- [How to Reenable the Internal Directory \(Knowledge base article\)](#)

- Microsoft Active Directory
- Apache Directory Server (ApacheDS)
- Apple Open Directory
- Fedora Directory Server
- Novell eDirectory
- OpenDS
- OpenLDAP
- OpenLDAP Using Posix Schema
- Posix Schema for LDAP
- Sun Directory Server Enterprise Edition (DSEE)
- A generic LDAP directory server

When to use this option: Connecting to an LDAP directory server is useful if your users and groups are stored in a corporate directory. When configuring the directory, you can choose to make it read only, read only with local groups, or read/write. If you choose read/write, any changes made to user and group information in the application will also update the LDAP directory.

On this page:

- [Overview](#)
- [Connecting to an LDAP Directory in Confluence](#)
- [Server Settings](#)
- [Schema Settings](#)
- [Permission Settings](#)
 - [Adding Users to Groups Automatically](#)
- [Advanced Settings](#)
- [User Schema Settings](#)
- [Group Schema Settings](#)
- [Membership Schema Settings](#)
- [Diagrams of Some Possible Configurations](#)

Related pages:

- [Configuring User Directories](#)

Connecting to an LDAP Directory in Confluence

To connect Confluence to an LDAP directory:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Click **User Directories** in the left-hand panel.
3. **Add** a directory and select one of these types:
 - **Microsoft Active Directory** – This option provides a quick way to select AD, because it is the most popular LDAP directory type.
 - **LDAP** – You will be able to choose a specific LDAP directory type on the next screen.
4. Enter the values for the settings, as described below.
5. Save the directory settings.
6. Define the **directory order** by clicking the blue up- and down-arrows next to each directory on the 'User Directories' screen. Here is a summary of how the directory order affects the processing:
 - Changes to users and groups will be made only in the first directory where the application has permission to make changes.
 - The order of the directories is the order in which they will be searched for users and groups (by default Confluence aggregates group membership from all directories, so the order does not impact membership itself).

For details see [Managing Multiple Directories](#).

Server Settings

Setting	Description
Name	Enter a meaningful name to help you identify the LDAP directory server. Examples: <ul style="list-style-type: none"> • Example Company Staff Directory • Example Company Corporate LDAP
Directory Type	Select the type of LDAP directory that you will connect to. If you are adding a new LDAP connection, the value you select here will determine the default values for many of the options on the rest of screen. Examples: <ul style="list-style-type: none"> • Microsoft Active Directory • OpenDS • And more.

Hostname	<p>The host name of your directory server. Examples:</p> <ul style="list-style-type: none"> • <code>ad.example.com</code> • <code>ldap.example.com</code> • <code>opens.example.com</code>
Port	<p>The port on which your directory server is listening. Examples:</p> <ul style="list-style-type: none"> • 389 • 10389 • 636 (for example, for SSL)
Use SSL	<p>Check this if the connection to the directory server is an SSL (Secure Sockets Layer) connection. Note that you will need to configure an SSL certificate in order to use this setting.</p>
Username	<p>The distinguished name of the user that the application will use when connecting to the directory server. Examples:</p> <ul style="list-style-type: none"> • <code>cn=administrator,cn=users,dc=ad,dc=example,dc=com</code> • <code>cn=user,dc=domain,dc=name</code> • <code>user@domain.name</code> <p>Ensure that this is an administrator user for the LDAP engine. For example, in Active Directory the user will need to be a member of the built-in Administrators group. The specific privileges for the LDAP user that is used to connect to LDAP are bind and read (user info, group info, group membership, update sequence number, deleted objects). The need for admin privileges is because a normal user can't access uSNChanged attribute and deleted objects container, causing incremental sync to fail silently, this was reported in a bug here CWD-3093.</p>
Password	<p>The password of the user specified above.</p> <p>Note: Connecting to an LDAP server requires that this application log in to the server with the username and password configured here. As a result, this password cannot be one-way hashed - it must be recoverable in the context of this application. The password is currently stored in the database in plain text without obfuscation. To guarantee its security, you need to ensure that other processes do not have OS-level read permissions for this application's database or configuration files.</p>

Schema Settings

Setting	Description
Base DN	<p>The root distinguished name (DN) to use when running queries against the directory server. Examples:</p> <ul style="list-style-type: none"> • <code>o=example,c=com</code> • <code>cn=users,dc=ad,dc=example,dc=com</code> • For Microsoft Active Directory, specify the base DN in the following format: <code>dc=domain1,dc=local</code>. You will need to replace the <code>domain1</code> and <code>local</code> for your specific configuration. Microsoft Server provides a tool called <code>ldp.exe</code> which is useful for finding out and configuring the the LDAP structure of your server.
Additional User DN	<p>This value is used in addition to the base DN when searching and loading users. If no value is supplied, the subtree search will start from the base DN. Example:</p> <ul style="list-style-type: none"> • <code>ou=Users</code>
Additional Group DN	<p>This value is used in addition to the base DN when searching and loading groups. If no value is supplied, the subtree search will start from the base DN. Example:</p> <ul style="list-style-type: none"> • <code>ou=Groups</code>

Permission Settings

Note: You can only assign LDAP users to local groups when 'External Management User Management' is not selected.

Setting	Description
Read Only	LDAP users, groups and memberships are retrieved from your directory server and can only be modified via your directory server. You cannot modify LDAP users, groups or memberships via the application administration screens.
Read Only, with Local Groups	<p>LDAP users, groups and memberships are retrieved from your directory server and can only be modified via your directory server. You cannot modify LDAP users, groups or memberships via the application administration screens. However, you can add groups to the internal directory and add LDAP users to those groups.</p> <p>Note for Confluence users: Users from LDAP are added to groups maintained in Confluence's internal directory the first time they log in. This is only done once per user. There is a known issue with Read Only, with Local Groups in Confluence that may apply to you. See CONF-28621 - User Loses all Local Group Memberships If LDAP Sync is Unable to find the User, but the User appears again in subsequent syncs OPEN</p>
Read/Write	LDAP users, groups and memberships are retrieved from your directory server. When you modify a user, group or membership via the application administration screens, the changes will be applied directly to your LDAP directory server. Please ensure that the LDAP user specified for the application has modification permissions on your LDAP directory server.

Adding Users to Groups Automatically

Setting	Description
Default Group Memberships	<p><i>Option available in Confluence 3.5 and later, and JIRA 4.3.3 and later.</i> This field appears if you select the 'Read Only, with Local Groups' permission. If you would like users to be automatically added to a group or groups, enter the group name(s) here. To specify more than one group, separate the group names with commas.</p> <p><i>In Confluence 3.5 to Confluence 3.5.1:</i> Each time a user logs in, their group memberships will be checked. If the user does not belong to the specified group(s), their username will be added to the group(s). If a group does not yet exist, it will be added locally.</p> <p><i>In Confluence 3.5.2 and later, and JIRA 4.3.3 and later:</i> The first time a user logs in, their group memberships will be checked. If the user does not belong to the specified group(s), their username will be added to the group(s). If a group does not yet exist, it will be added locally. On subsequent logins, the username will <i>not</i> be added automatically to any groups. This change in behaviour allows users to be removed from automatically-added groups. In Confluence 3.5 and 3.5.1, they would be re-added upon next login.</p> <p>Please note that there is no validation of the group names. If you mis-type the group name, authorisation failures will result – users will not be able to access the applications or functionality based on the intended group name.</p> <p>Examples:</p> <ul style="list-style-type: none"> confluence-users confluence-users , jira-administrators , jira-core-users

Advanced Settings

Setting	Description
---------	-------------

Enable Nested Groups	Enable or disable support for nested groups. Some directory servers allow you to define a group as a member of another group. Groups in such a structure are called 'nested groups'. If you are using groups to manage permissions, you can create nested groups to allow inheritance of permissions from one group to its sub-groups.
Manage User Status Locally	If true, you can activate and deactivate users in Crowd independent of their status in the directory server.
Filter out expired users	If true, user accounts marked as expired in ActiveDirectory will be automatically removed. For cached directories, the removal of a user will occur during the first synchronisation after the account's expiration date.
Use Paged Results	Enable or disable the use of the LDAP control extension for simple paging of search results. If paging is enabled, the search will retrieve sets of data rather than all of the search results at once. Enter the desired page size – that is, the maximum number of search results to be returned per page when paged results are enabled. The default is 1000 results.
Follow Referrals	Choose whether to allow the directory server to redirect requests to other servers. This option uses the node referral (JNDI lookup <code>java.naming.referral</code>) configuration setting. It is generally needed for Active Directory servers configured without proper DNS, to prevent a 'javax.naming.PartialResultException: Unprocessed Continuation Reference(s)' error.
Naive DN Matching	<p>If your directory server will always return a consistent string representation of a DN, you can enable naive DN matching. Using naive DN matching will result in a significant performance improvement, so we recommend enabling it where possible.</p> <p>This setting determines how your application will compare DN's to determine if they are equal.</p> <ul style="list-style-type: none"> • If this checkbox is selected, the application will do a direct, case-insensitive, string comparison. This is the default and recommended setting for Active Directory, because Active Directory guarantees the format of DN's. • If this checkbox is not selected, the application will parse the DN and then check the parsed version.
Enable Incremental Synchronisation	<p>Enable incremental synchronisation if you only want changes since the last synchronisation to be queried when synchronising a directory.</p> <p>Please be aware that when using this option, the user account configured for synchronisation must have read access to:</p> <ul style="list-style-type: none"> • The <code>uSNChanged</code> attribute of all users and groups in the directory that need to be synchronised. • The objects and attributes in the Active Directory deleted objects container (see Microsoft's Knowledge Base Article No. 892806 for details). <p>If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application.</p> <p>This setting is only available if the directory type is set to "Microsoft Active Directory".</p>
Synchronisation Interval (minutes)	Synchronisation is the process by which the application updates its internal store of user data to agree with the data on the directory server. The application will send a request to your directory server every x minutes, where 'x' is the number specified here. The default value is 60 minutes.
Read Timeout (seconds)	The time, in seconds, to wait for a response to be received. If there is no response within the specified time period, the read attempt will be aborted. A value of 0 (zero) means there is no limit. The default value is 120 seconds.

Search Timeout (seconds)	The time, in seconds, to wait for a response from a search operation. A value of 0 (zero) means there is no limit. The default value is 60 seconds.
Connection Timeout (seconds)	<p>This setting affects two actions. The default value is 0.</p> <ul style="list-style-type: none"> The time to wait when getting a connection from the connection pool. A value of 0 (zero) means there is no limit, so wait indefinitely. The time, in seconds, to wait when opening new server connections. A value of 0 (zero) means that the TCP network timeout will be used, which may be several minutes.

User Schema Settings

Setting	Description
User Object Class	<p>This is the name of the class used for the LDAP user object. Example:</p> <ul style="list-style-type: none"> <code>user</code>
User Object Filter	<p>The filter to use when searching user objects. Example:</p> <ul style="list-style-type: none"> <code>(&(objectCategory=Person)(sAMAccountName=*))</code> <p>More examples can be found here and here.</p>
User Name Attribute	<p>The attribute field to use when loading the username. Examples:</p> <ul style="list-style-type: none"> <code>cn</code> <code>sAMAccountName</code> <p>NB: In Active Directory, the 'sAMAccountName' is the 'User Logon Name (pre-Windows 2000)' field. The User Logon Name field is referenced by 'cn'.</p>
User Name RDN Attribute	<p>The RDN (relative distinguished name) to use when loading the username. The DN for each LDAP entry is composed of two parts: the RDN and the location within the LDAP directory where the record resides. The RDN is the portion of your DN that is not related to the directory tree structure. Example:</p> <ul style="list-style-type: none"> <code>cn</code>
User First Name Attribute	<p>The attribute field to use when loading the user's first name. Example:</p> <ul style="list-style-type: none"> <code>givenName</code>
User Last Name Attribute	<p>The attribute field to use when loading the user's last name. Example:</p> <ul style="list-style-type: none"> <code>sn</code>
User Display Name Attribute	<p>The attribute field to use when loading the user's full name. Example:</p> <ul style="list-style-type: none"> <code>displayName</code>
User Email Attribute	<p>The attribute field to use when loading the user's email address. Example:</p> <ul style="list-style-type: none"> <code>mail</code>
User Password Attribute	<p>The attribute field to use when loading a user's password. Example:</p> <ul style="list-style-type: none"> <code>unicodePwd</code>

User Unique ID Attribute	<p>The attribute used as a unique immutable identifier for user objects. This is used to track username changes and is optional. If this attribute is not set (or is set to an invalid value), user renames will not be detected — they will be interpreted as a user deletion then a new user addition.</p> <p>This should normally point to a UUID value. Standards-compliant LDAP servers will implement this as 'entryUUID' according to RFC 4530. This setting exists because it is known under different names on some servers, e.g. 'objectGUID' in Microsoft Active Directory.</p>
--------------------------	--

Group Schema Settings

Setting	Description
Group Object Class	<p>This is the name of the class used for the LDAP group object. Examples:</p> <ul style="list-style-type: none"> groupOfUniqueNames group
Group Object Filter	<p>The filter to use when searching group objects. Example:</p> <ul style="list-style-type: none"> (& (objectClass=group) (cn=*))
Group Name Attribute	<p>The attribute field to use when loading the group's name. Example:</p> <ul style="list-style-type: none"> cn
Group Description Attribute	<p>The attribute field to use when loading the group's description. Example:</p> <ul style="list-style-type: none"> description

Membership Schema Settings

Setting	Description
Group Members Attribute	<p>The attribute field to use when loading the group's members. Example:</p> <ul style="list-style-type: none"> member
User Membership Attribute	<p>The attribute field to use when loading the user's groups. Example:</p> <ul style="list-style-type: none"> memberOf
Use the User Membership Attribute, when finding the user's group membership	<p>Check this if your directory server supports the group membership attribute on the user. (By default, this is the 'memberOf' attribute.)</p> <ul style="list-style-type: none"> If this checkbox is selected, your application will use the group membership attribute on the user when retrieving the list of groups to which a given user belongs. This will result in a more efficient retrieval. If this checkbox is not selected, your application will use the members attribute on the group ('member' by default) for the search. If the Enable Nested Groups checkbox is selected, your application will ignore the Use the User Membership Attribute option and will use the members attribute on the group for the search.

Use the User Membership Attribute, when finding the members of a group	<p>Check this if your directory server supports the user membership attribute on the group. (By default, this is the 'member' attribute.)</p> <ul style="list-style-type: none">• If this checkbox is selected, your application will use the group membership attribute on the user when retrieving the members of a given group. This will result in a more efficient search.• If this checkbox is not selected, your application will use the members attribute on the group ('member' by default) for the search.
--	---

Diagrams of Some Possible Configurations

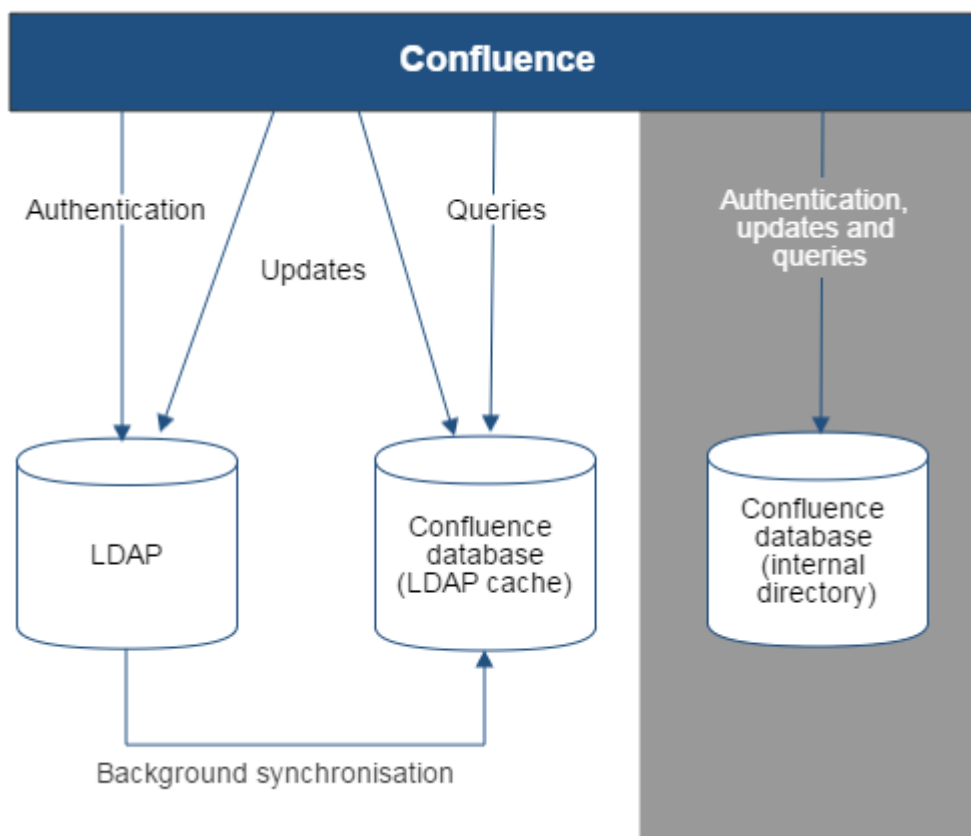


Diagram above: Confluence connecting to an LDAP directory.

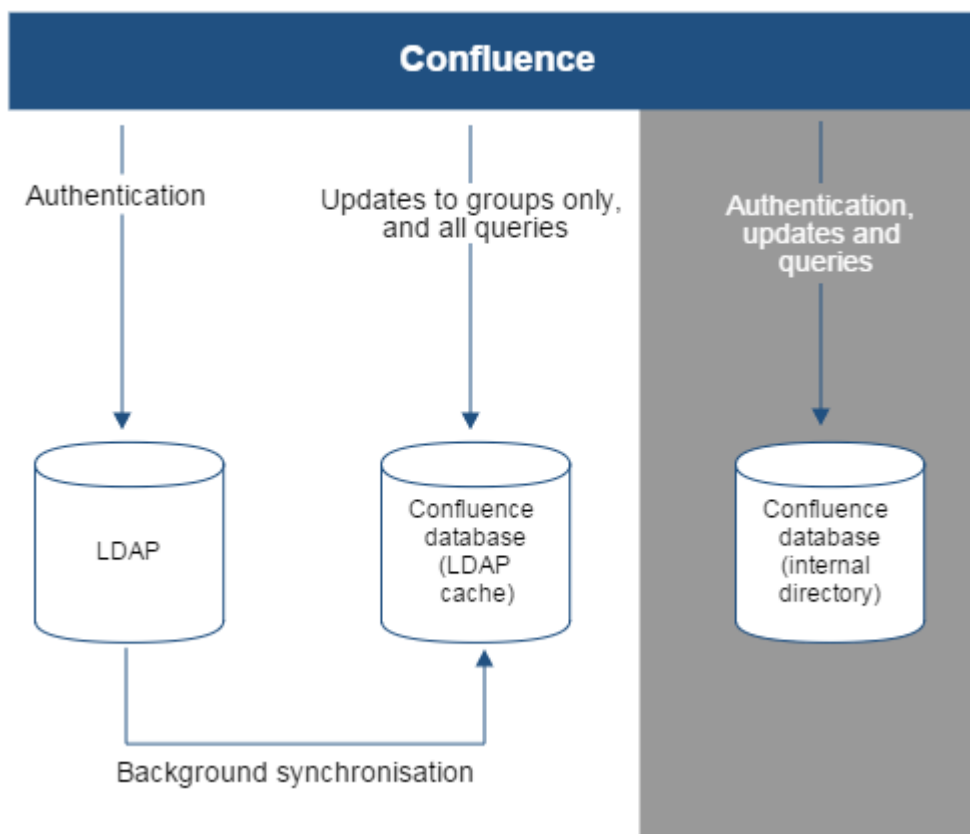


Diagram above: Confluence connecting to an LDAP directory with permissions set to read only and local groups.

Configuring the LDAP Connection Pool

When connection pooling is enabled, the LDAP directory server maintains a pool of connections and assigns them as needed. When a connection is closed, the directory server returns the connection to the pool for future use. This can improve performance significantly.

Related pages:

- [Connecting to an LDAP Directory](#)
- [Configuring User Directories](#)

To configure your LDAP connection pool:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Click '**User Directories**' in the left-hand panel.
3. Click '**LDAP Connection Pool Configuration**' in the 'Additional Configuration' section.

Setting	Description	Default Value
Initial Pool Size	The number of LDAP connections created when initially connecting to the pool.	1
Preferred Pool Size	The optimal pool size. LDAP will remove idle connections when the number of connections grows larger than this value. A value of 0 (zero) means that there is no preferred size, so the number of idle connections is unlimited.	10

Maximum Pool Size	The maximum number of connections. When the number of connections reaches this value, LDAP will refuse further connections. As a result, requests made by an application to the LDAP directory server will be blocked. A value of 0 (zero) means that the number of connections is unlimited.	0
Pool Timeout (seconds)	The length of time, in seconds, that a connection may remain idle before being removed from the pool. When the application is finished with a pooled connection, the connection is marked as idle, waiting to be reused. A value of 0 (zero) means that the idle time is unlimited, so connections will never be timed out.	30
Pool Protocol	Only these protocol types will be allowed to connect to the LDAP directory server. If you want to allow multiple protocols, enter the values separated by a space. Valid values are: <ul style="list-style-type: none"> • plain • ssl 	plain ssl (Both plain and ssl)
Pool Authentication	Only these authentication types will be allowed to connect to the LDAP directory server. If you want to allow multiple authentication types, enter the values separated by a space. See RFC 2829 for details of LDAP authentication methods. Valid values are: <ul style="list-style-type: none"> • none • simple • DIGEST-MD5 	simple

Notes:

- The connection pool settings are system wide and will be used to create a new connection pool for every configured LDAP directory server.
- You must restart your application server for these settings to take effect.

Configuring an SSL Connection to Active Directory

If you want to configure a read/write connection with Microsoft Active Directory, you will need to install an SSL certificate, generated by your Active Directory server, onto your Confluence server and then install the certificate into your JVM keystore.

On this page:

- [Prerequisites](#)
- [Step 1. Install the Active Directory Certificate Services](#)
- [Step 2. Obtain the Server Certificate](#)
- [Step 3. Import the Server Certificate](#)

There's a [Confluence SSL plugin](#) that facilitates this process.

Updating user, group, and membership details in Active Directory requires that your Atlassian application be running in a JVM that trusts the AD server. To do this, we generate a certificate on the Active Directory server, then import it into Java's keystore.

Prerequisites

To generate a certificate, you need the following components installed on the Windows Domain Controller to which you're connecting.

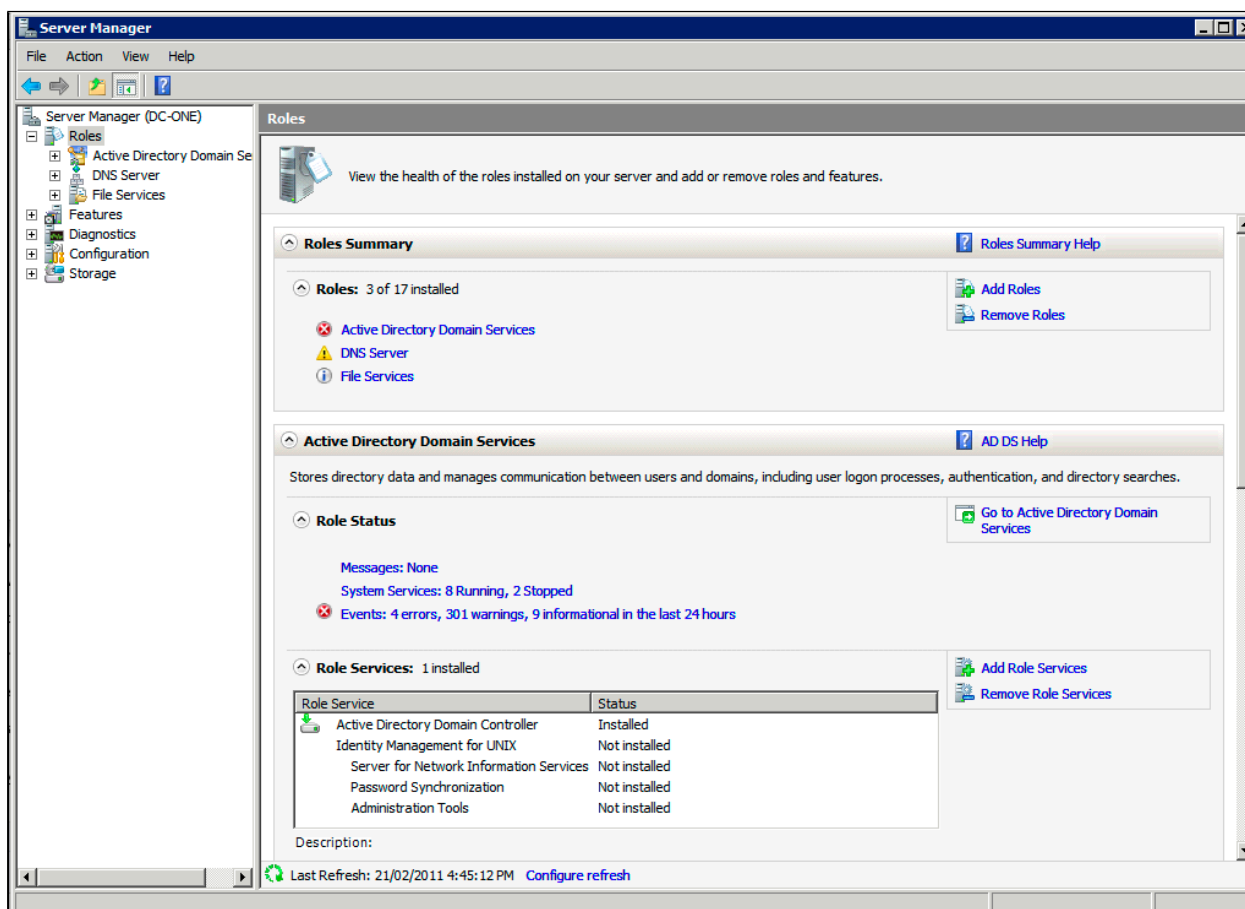
Required Component	Description
Internet Information Services (IIS)	This is required before you can install Windows Certificate Services.

Windows Certificate Services	This installs a certification authority (CA) which is used to issue certificates. Step 1, below, explains this process.
Windows 2000 Service Pack 2	Required if you are using Windows 2000
Windows 2000 High Encryption Pack (128-bit)	Required if you are using Windows 2000. Provides the highest available encryption level (128-bit).

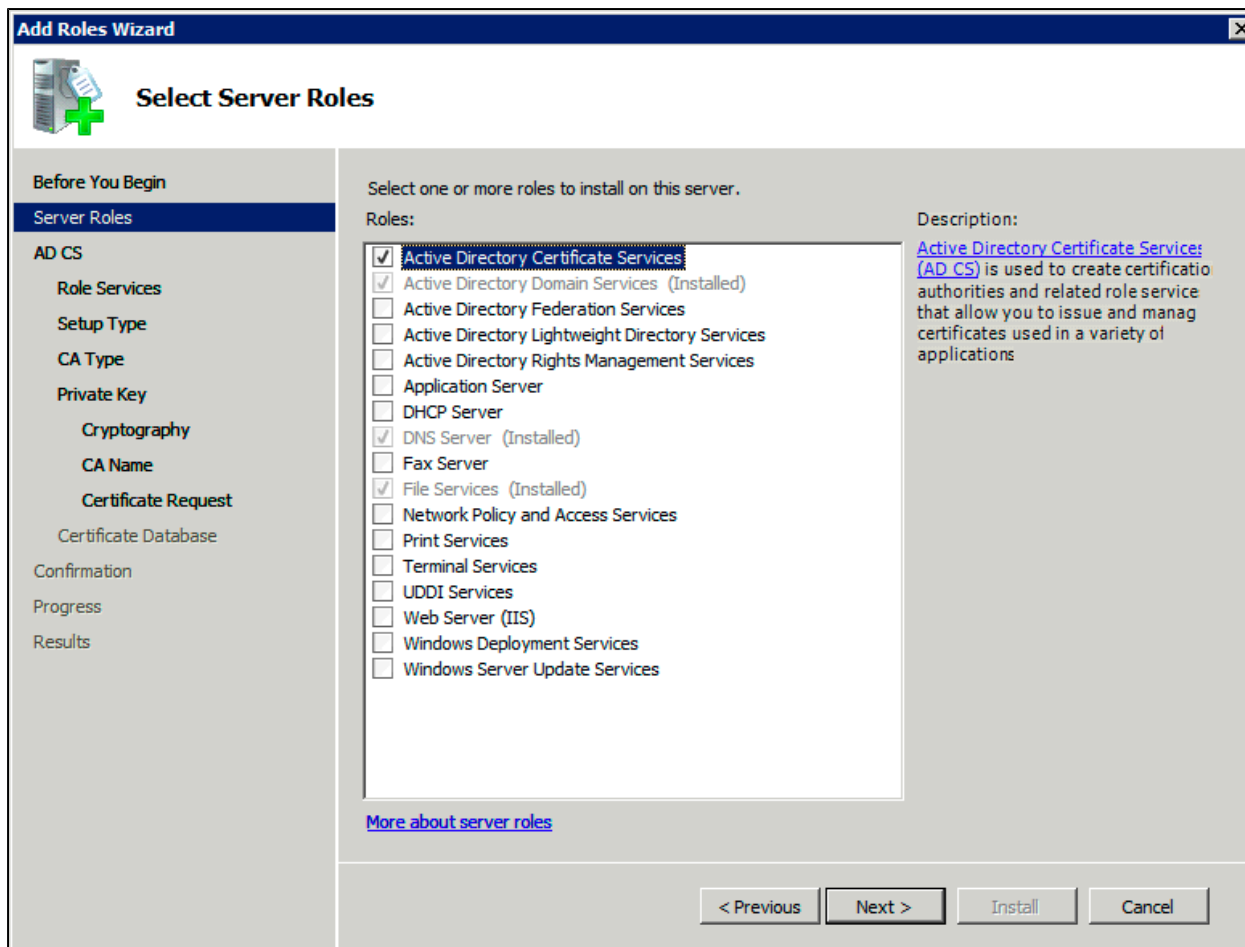
Step 1. Install the Active Directory Certificate Services

If Certificate Services are already installed, skip to step 2, below. The screenshots below are from Server 2008, but the process is similar for Server 2000 and 2003.

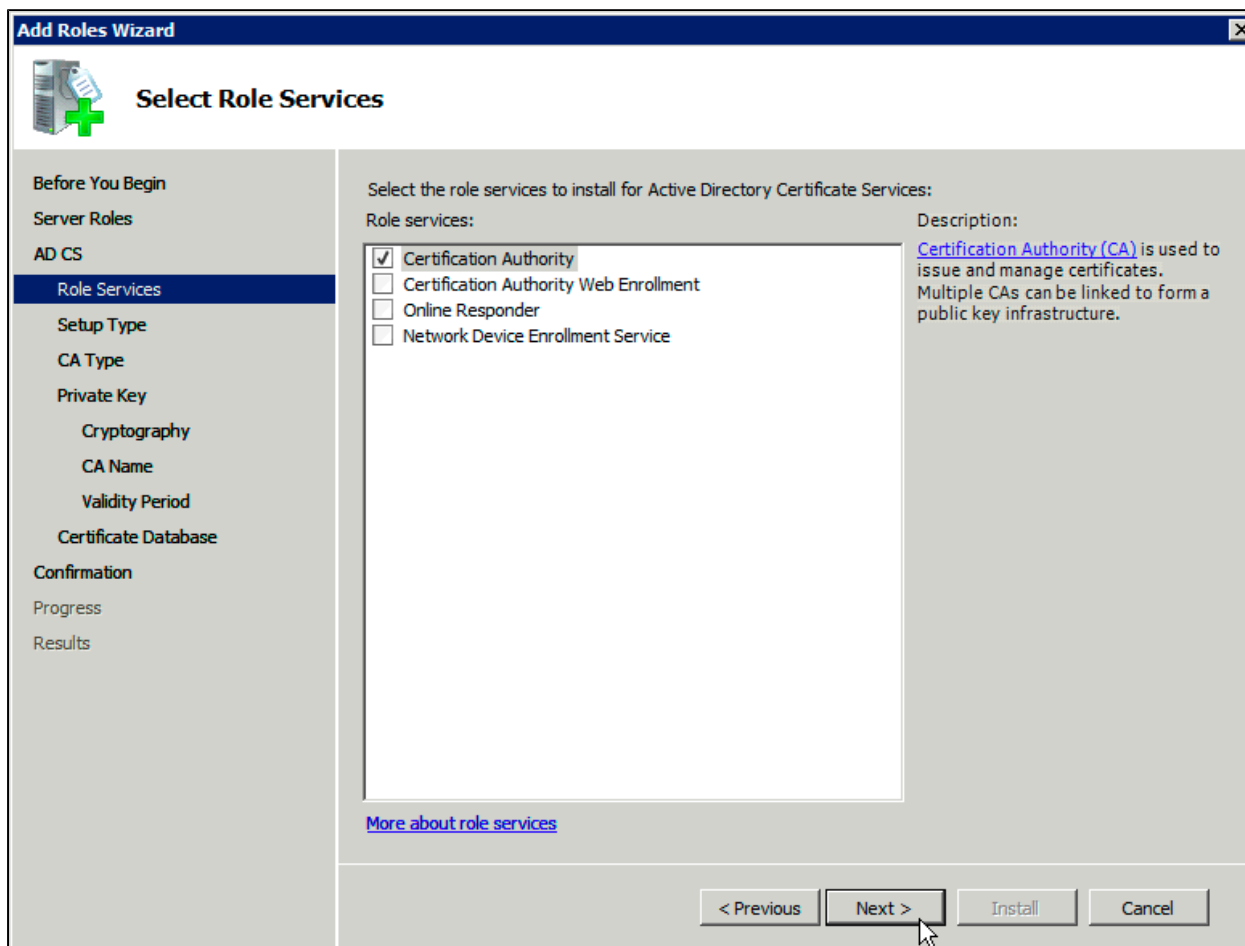
1. Log in to your Active Directory server as an administrator.
2. Click **Start**, point to **Administrative Tools**, and then click **Server Manager**.
3. In the **Roles Summary** section, click **Add Roles**.



4. On the **Select Server Roles** page, select the **Active Directory Certificate Services** check box. Click **Next** twice.



5. On the **Select Role Services** page, select the **Certification Authority** check box, and then click **Next**.



6. On the **Specify Setup Type** page, click **Enterprise**, and then click **Next**.

Add Roles Wizard

Specify Setup Type

Before You Begin
Server Roles
AD CS
Role Services
Setup Type
CA Type
Private Key
Cryptography
CA Name
Validity Period
Certificate Database
Confirmation
Progress
Results

Certification Authorities can use data in Active Directory to simplify the issuance and management of certificates. Specify whether you want to set up an Enterprise or Standalone CA.

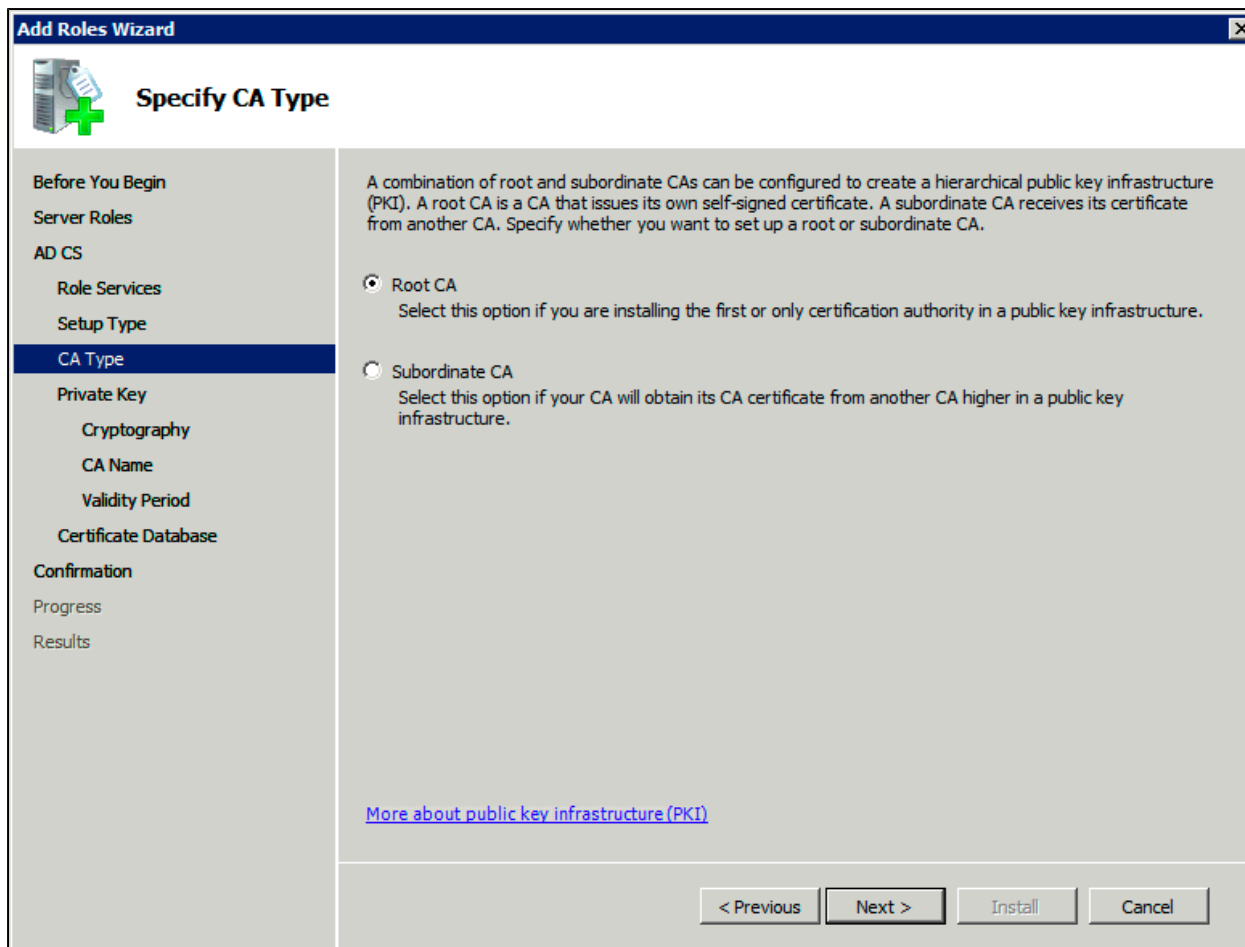
☒ **Enterprise**
Select this option if this CA is a member of a domain and can use Directory Service to issue and manage certificates.

☐ **Standalone**
Select this option if this CA does not use Directory Service data to issue or manage certificates. A standalone CA can be a member of a domain.

[More about the differences between enterprise and standalone setup](#)

< Previous Next > Install Cancel

7. On the **Specify CA Type** page, click **Root CA**, and then click **Next**.



The screenshot shows the 'Add Roles Wizard' window with the title bar 'Add Roles Wizard' and a close button. The main window is titled 'Specify CA Type' and features a green plus icon. On the left is a navigation pane with the following items: 'Before You Begin', 'Server Roles', 'AD CS', 'Role Services', 'Setup Type', 'CA Type' (highlighted), 'Private Key', 'Cryptography', 'CA Name', 'Validity Period', 'Certificate Database', 'Confirmation', 'Progress', and 'Results'. The main content area contains the following text: 'A combination of root and subordinate CAs can be configured to create a hierarchical public key infrastructure (PKI). A root CA is a CA that issues its own self-signed certificate. A subordinate CA receives its certificate from another CA. Specify whether you want to set up a root or subordinate CA.' Below this text are two radio button options: 'Root CA' (selected) with the description 'Select this option if you are installing the first or only certification authority in a public key infrastructure.', and 'Subordinate CA' with the description 'Select this option if your CA will obtain its CA certificate from another CA higher in a public key infrastructure.' At the bottom of the main area is a blue hyperlink: '[More about public key infrastructure \(PKI\)](#)'. At the bottom right are four buttons: '< Previous', 'Next >', 'Install', and 'Cancel'.

8. On the **Set Up Private Key** and **Configure Cryptography for CA** pages, you can configure optional configuration settings, including cryptographic service providers. However, the default values should be fine. Click **Next** twice.

Add Roles Wizard

Set Up Private Key

To generate and issue certificates to clients, a CA must have a private key. Specify whether you want to create a new private key or use an existing one.

- ☒ **Create a new private key**
Use this option if you don't have a private key or wish to create a new private key to enhance security. You will be asked to select a cryptographic service provider and specify a key length for the private key. To issue new certificates, you must also select a hash algorithm.
- ☐ **Use existing private key**
Use this option to ensure continuity with previously issued certificates when reinstalling a CA.
 - ☒ **Select a certificate and use its associated private key**
Select this option if you have an existing certificate on this computer or if you want to import a certificate and use its associated private key.
 - ☐ **Select an existing private key on this computer**
Select this option if you have retained private keys from a previous installation or want to use a private key from an alternate source.

[More about public and private keys](#)

< Previous Next > Install Cancel

9. In the **Common name for this CA** box, type the common name of the CA, and then click **Next**.

Add Roles Wizard

Configure CA Name

Type in a common name to identify this CA. This name is added to all certificates issued by the CA. Distinguished name suffix values are automatically generated but can be modified.

Common name for this CA:
dc-one-DC-ONE-CA

Distinguished name suffix:
DC=dc-one,DC=crowd,DC=sydney,DC=atlassian,DC=com

Preview of distinguished name:
CN=dc-one-DC-ONE-CA,DC=dc-one,DC=crowd,DC=sydney,DC=atlassian,DC=com

[More about configuring a CA name](#)

< Previous Next > Install Cancel

10. On the **Set Validity Period** page, accept the default values or specify other storage locations for the

certificate database and the certificate database log, and then click **Next**.

Add Roles Wizard

Set Validity Period

Before You Begin
Server Roles
AD CS
Role Services
Setup Type
CA Type
Private Key
Cryptography
CA Name
Validity Period
Certificate Database
Confirmation
Progress
Results

A certificate will be issued to this CA to secure communications with other CAs and with clients requesting certificates. The validity period of a CA certificate can be based on a number of factors, including the intended purpose of the CA and security measures that you have taken to secure the CA.

Select validity period for the certificate generated for this CA:
5 Years

CA expiration Date: 21/02/2016 4:57 PM
Note that CA will issue certificates valid only until its expiration date.

[More about setting the certificate validity period](#)

< Previous Next > Install Cancel

Add Roles Wizard

Configure Certificate Database

Before You Begin
Server Roles
AD CS
Role Services
Setup Type
CA Type
Private Key
Cryptography
CA Name
Validity Period
Certificate Database
Confirmation
Progress
Results

The certificate database records all certificate requests, issued certificates, and revoked or expired certificates. The database log can be used to monitor management activity for a CA.

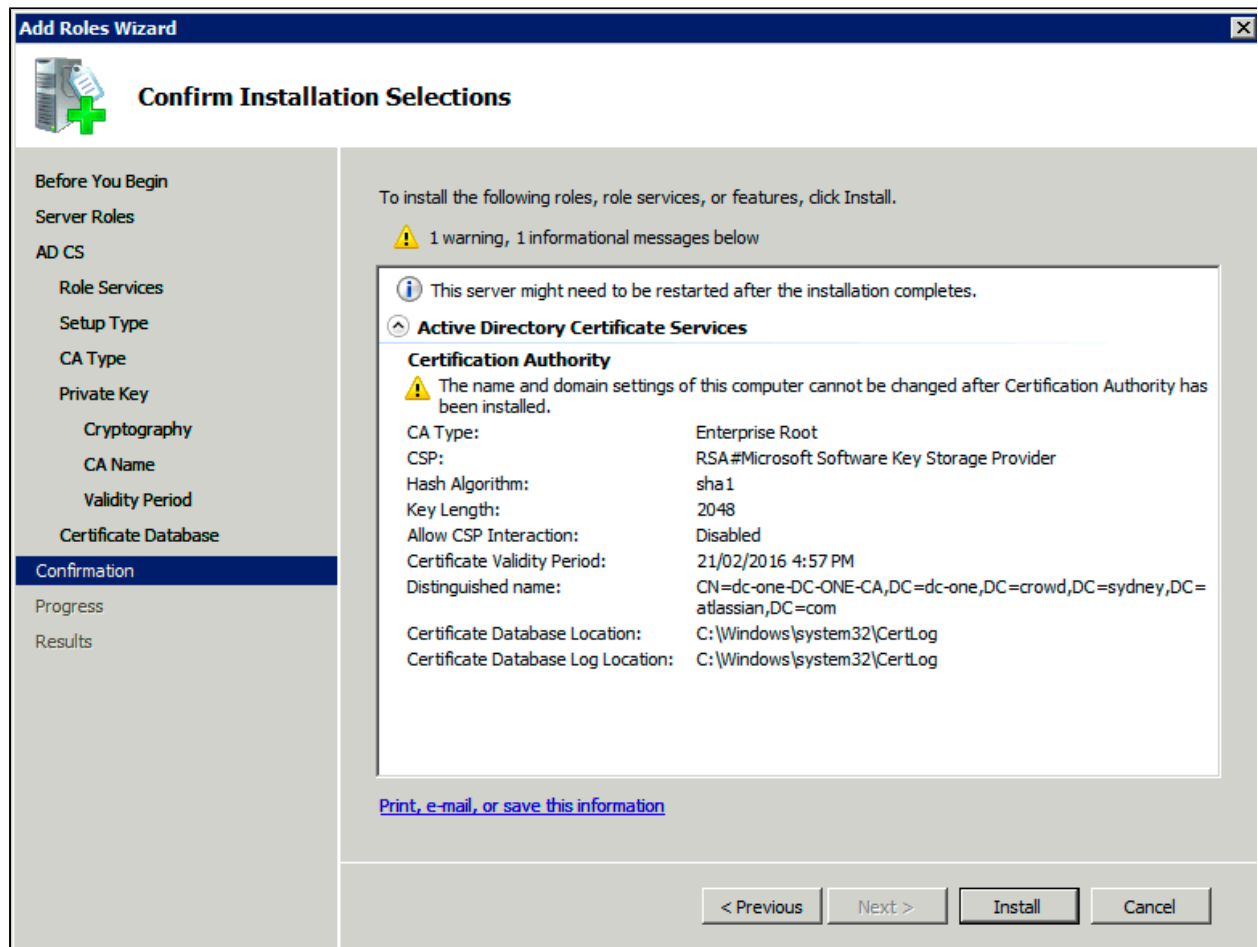
Certificate database location:
C:\Windows\system32\CertLog Browse...

☐ Use existing certificate database from previous installation at this location

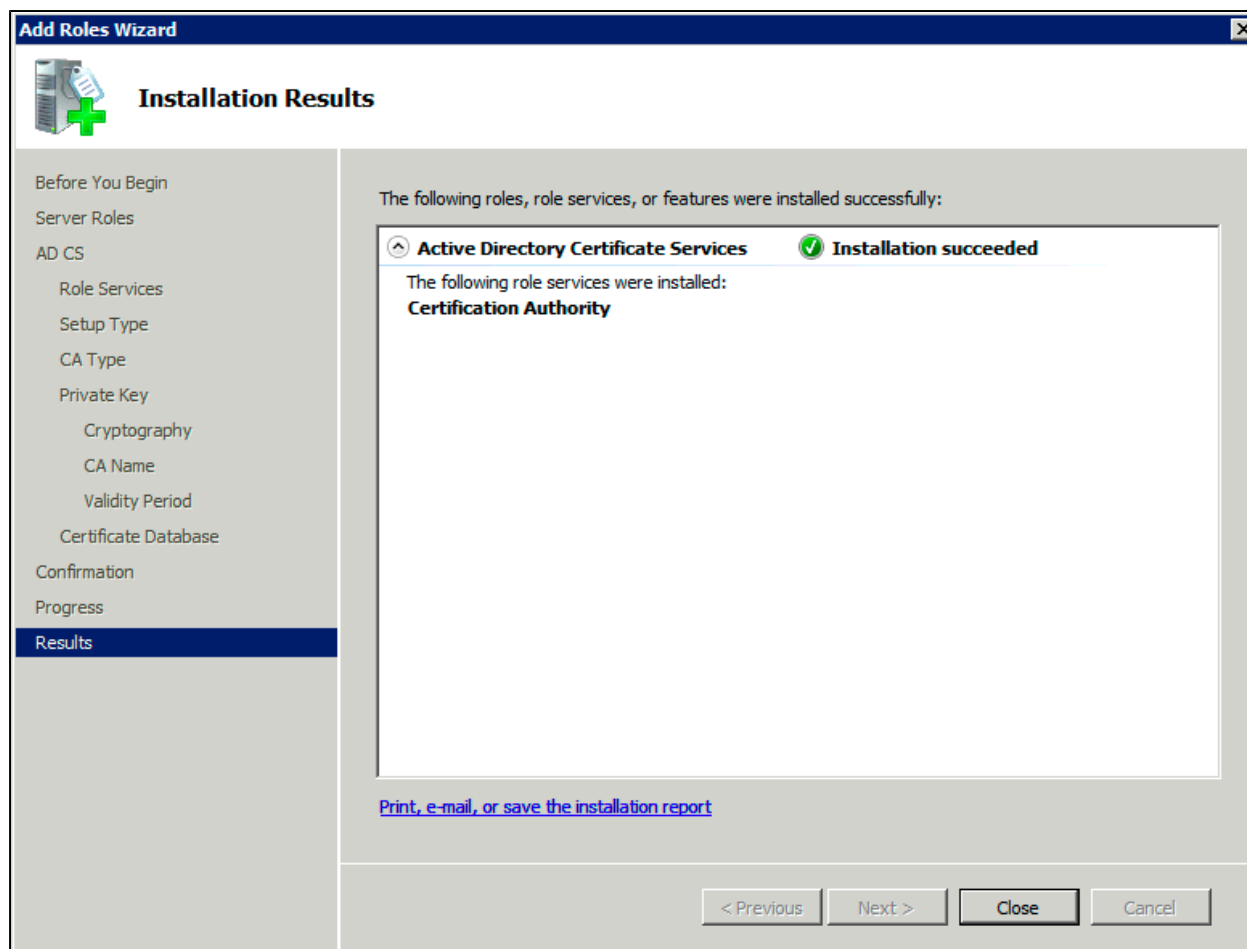
Certificate database log location:
C:\Windows\system32\CertLog Browse...

< Previous Next > Install Cancel

11. After verifying the information on the **Confirm Installation Selections** page, click **Install**.



12. Review the information on the results screen to verify that the installation was successful.



Step 2. Obtain the Server Certificate

The steps above describe how to install the certification authority (CA) on your Microsoft Active Directory server. Next, you will need to add the Microsoft Active Directory server's SSL certificate to the list of accepted certificates used by the JDK that runs your application server.

The Active Directory certificate is automatically generated and placed in root of the C:\ drive, matching a file format similar to the tree structure of your Active Directory server. For example: c:\ad2008.ad01.atlassian.com_ad01.crt.

You can also export the certificate by executing this command on the Active Directory server:

```
certutil -ca.cert client.crt
```

You might still fail in being authenticated with the certificate file above. In this case, Microsoft's [LDAP over SSL \(LDAPS\) Certificate](#) page might help. Be noted that you need:

1. to choose "No, do not export the private key" in step-10 of [Exporting the LDAPS Certificate and Importing for use with AD DS](#) section
2. to choose "DER encoded binary X.509 (.CER)" in step-11 of [Exporting the LDAPS Certificate and Importing for use with AD DS](#) section. This file will be used in the following step.

Step 3. Import the Server Certificate

For an application server to trust your directory's certificate, the certificate must be imported into your Java runtime environment. The JDK stores trusted certificates in a file called a keystore. The default keystore file is called `cacerts` and it lives in the `jre\lib\security` sub-directory of your Java installation.

In the following examples, we use `server-certificate.crt` to represent the certificate file exported by your directory server. You will need to alter the instructions below to match the name actually generated.

Once the certificate has been imported as per the below instructions, you will need to restart the application to

pick up the changes.

Windows

1. Navigate to the directory in which Java is installed. It's probably called something like C:\Program Files\Java\jdk1.5.0_12.

```
cd /d C:\Program Files\Java\jdk1.5.0_12
```

2. Run the command below, where `server-certificate.crt` is the name of the file from your directory server:

```
keytool -importcert -keystore .\jre\lib\security\cacerts -file  
server-certificate.crt
```

3. `keytool` will prompt you for a password. The default keystore password is `changeit`.
4. When prompted `Trust this certificate? [no]:` enter `yes` to confirm the key import:

```
Enter keystore password:  changeit  
Owner: CN=ad01, C=US  
Issuer: CN=ad01, C=US  
Serial number: 15563d6677a4e9e4582d8a84be683f9  
Valid from: Tue Aug 21 01:10:46 ACT 2007 until: Tue Aug 21 01:13:59  
ACT 2012  
Certificate fingerprints:  
    MD5:  D6:56:F0:23:16:E3:62:2C:6F:8A:0A:37:30:A1:84:BE  
    SHA1:  
73:73:4E:A6:A0:D1:4E:F4:F3:CD:CE:BE:96:80:35:D2:B4:7C:79:C1  
Trust this certificate? [no]:  yes  
Certificate was added to keystore
```

You may now change 'URL' to use LDAP over SSL (i.e. `ldaps://<HOSTNAME>:636/`) and use the 'Secure SSL' option when connecting your application to your directory server.

UNIX

1. Navigate to the directory in which the Java used by JIRA is installed. If the default JAVA installation is used, then it would be

```
cd $JAVA_HOME
```

2. Run the command below, where `server-certificate.crt` is the name of the file from your directory server:

```
sudo keytool -importcert -keystore ./jre/lib/security/cacerts -file  
server-certificate.crt
```

3. `keytool` will prompt you for a password. The default keystore password is `changeit`.
4. When prompted `Trust this certificate? [no]:` enter `yes` to confirm the key import:

```
Password:
Enter keystore password:  changeit
Owner: CN=ad01, C=US
Issuer: CN=ad01, C=US
Serial number: 15563d6677a4e9e4582d8a84be683f9
Valid from: Tue Aug 21 01:10:46 ACT 2007 until: Tue Aug 21 01:13:59
ACT 2012
Certificate fingerprints:
    MD5:  D6:56:F0:23:16:E3:62:2C:6F:8A:0A:37:30:A1:84:BE
    SHA1:
73:73:4E:A6:A0:D1:4E:F4:F3:CD:CE:BE:96:80:35:D2:B4:7C:79:C1
Trust this certificate? [no]:  yes
Certificate was added to keystore
```

You may now change 'URL' to use LDAP over SSL (i.e. `ldaps://<HOSTNAME>:636/`) and use the 'Secure SSL' option when connecting your application to your directory server.

Mac OS X

1. Navigate to the directory in which Java is installed. This is usually

```
cd /Library/Java/Home
```

2. Run the command below, where `server-certificate.crt` is the name of the file from your directory server:

```
sudo keytool -importcert -keystore ./jre/lib/security/cacerts -file
server-certificate.crt
```

3. `keytool` will prompt you for a password. The default keystore password is `changeit`.
4. When prompted `Trust this certificate? [no]:` enter `yes` to confirm the key import:

```
Password:
Enter keystore password:  changeit
Owner: CN=ad01, C=US
Issuer: CN=ad01, C=US
Serial number: 15563d6677a4e9e4582d8a84be683f9
Valid from: Tue Aug 21 01:10:46 ACT 2007 until: Tue Aug 21 01:13:59
ACT 2012
Certificate fingerprints:
    MD5:  D6:56:F0:23:16:E3:62:2C:6F:8A:0A:37:30:A1:84:BE
    SHA1:
73:73:4E:A6:A0:D1:4E:F4:F3:CD:CE:BE:96:80:35:D2:B4:7C:79:C1
Trust this certificate? [no]:  yes
Certificate was added to keystore
```

You may now change 'URL' to use LDAP over SSL (i.e. `ldaps://<HOSTNAME>:636/`) and use the 'Secure SSL' option when connecting your application to your directory server.

RELATED TOPICS

[Connecting to an LDAP Directory](#)
[Configuring User Directories](#)

Connecting to an Internal Directory with LDAP Authentication

You can connect your Confluence application to an LDAP directory for delegated authentication. This means that Confluence will have an internal directory that uses LDAP for authentication only. There is an option to create users in the internal directory automatically when they attempt to log in, as described in the settings section.

Overview

An internal directory with LDAP authentication offers the features of an internal directory while allowing you to store and check users' passwords in LDAP only. Note that the 'internal directory with LDAP authentication' is separate from the default 'internal directory'. On LDAP, all that the application does is to check the password. The LDAP connection is read only. Every user in the internal directory with LDAP authentication must map to a user on LDAP, otherwise they cannot log in.

When to use this option: Choose this option if you want to set up a user and group configuration within your application that suits your needs, while checking your users' passwords against the corporate LDAP directory. This option also helps to avoid the performance issues that may result from downloading large numbers of groups from LDAP.

On this page:

- [Overview](#)
- [Connecting Confluence to an Internal Directory with LDAP Authentication](#)
- [Server Settings](#)
 - [Copying Users on Login](#)
- [Schema Settings](#)
- [Advanced Settings](#)
- [User Schema Settings](#)
- [Group Schema Settings](#)
- [Membership Schema Settings](#)
- [Diagrams of Possible Configurations](#)

Related pages:

[Configuring User Directories](#)

Connecting Confluence to an Internal Directory with LDAP Authentication

To connect to an internal directory but check logins via LDAP:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Click '**User Directories**' in the left-hand panel.
3. **Add** a directory and select type '**Internal with LDAP Authentication**'.
4. Enter the values for the settings, as described below.
5. Save the directory settings.
6. If you want LDAP users to be used in place of existing internal users, move the 'Internal with LDAP Authentication' directory to the top of the list. You can define the **directory order** by clicking the blue up- and down-arrows next to each directory on the '**User Directories**' screen. Here is a summary of how the directory order affects the processing:
 - Changes to users and groups will be made only in the first directory where the application has permission to make changes.
 - The order of the directories is the order in which they will be searched for users and groups (by default Confluence aggregates group membership from all directories, so the order does not impact membership itself).

For details see [Managing Multiple Directories](#).
7. Add your users and groups in Confluence. See [Add and Invite Users](#) and [Managing Site-Wide Permissions and Groups](#).

Server Settings

Setting	Description
---------	-------------

Name	A descriptive name that will help you to identify the directory. Examples: <ul style="list-style-type: none"> • Internal directory with LDAP Authentication • Corporate LDAP for Authentication Only
Directory Type	Select the type of LDAP directory that you will connect to. If you are adding a new LDAP connection, the value you select here will determine the default values for some of the options on the rest of screen. Examples: <ul style="list-style-type: none"> • Microsoft Active Directory • OpenDS • And more.
Hostname	The host name of your directory server. Examples: <ul style="list-style-type: none"> • ad.example.com • ldap.example.com • opens.example.com
Port	The port on which your directory server is listening. Examples: <ul style="list-style-type: none"> • 389 • 10389 • 636 (for example, for SSL)
Use SSL	Check this box if the connection to the directory server is an SSL (Secure Sockets Layer) connection. Note that you will need to configure an SSL certificate in order to use this setting.
Username	The distinguished name of the user that the application will use when connecting to the directory server. Examples: <ul style="list-style-type: none"> • cn=administrator,cn=users,dc=ad,dc=example,dc=com • cn=user,dc=domain,dc=name • user@domain.name
Password	The password of the user specified above.

Copying Users on Login

Setting	Description
Copy User on Login	<p>This option affects what will happen when a user attempts to log in. If this box is checked, the user will be created automatically in the internal directory that is using LDAP for authentication when the user first logs in and their details will be synchronised on each subsequent log in. If this box is not checked, the user's login will fail if the user wasn't already manually created in the directory.</p> <p>If you check this box the following additional fields will appear on the screen, which are described in more detail below:</p> <ul style="list-style-type: none"> • Default Group Memberships • Synchronise Group Memberships • User Schema Settings (described in a separate section below)

Default Group Memberships	<p>This field appears if you check the Copy User on Login box. If you would like users to be automatically added to a group or groups, enter the group name(s) here. To specify more than one group, separate the group names with commas. Each time a user logs in, their group memberships will be checked. If the user does not belong to the specified group(s), their username will be added to the group(s). If a group does not yet exist, it will be added to the internal directory that is using LDAP for authentication.</p> <p>Please note that there is no validation of the group names. If you mis-type the group name, authorisation failures will result – users will not be able to access the applications or functionality based on the intended group name.</p> <p>Examples:</p> <ul style="list-style-type: none"> • <code>confluence-users</code> • <code>bamboo-users,jira-administrators,jira-core-users</code>
Synchronise Group Memberships	<p>This field appears if you select the Copy User on Login checkbox. If this box is checked, group memberships specified on your LDAP server will be synchronised with the internal directory each time the user logs in.</p> <p>If you check this box the following additional fields will appear on the screen, both described in more detail below:</p> <ul style="list-style-type: none"> • Group Schema Settings (described in a separate section below) • Membership Schema Settings (described in a separate section below)

Note: 'Copy Users on Login' must be enabled if you want to be able to [change usernames](#).

Schema Settings

Setting	Description
Base DN	<p>The root distinguished name (DN) to use when running queries against the directory server. Examples:</p> <ul style="list-style-type: none"> • <code>o=example,c=com</code> • <code>cn=users,dc=ad,dc=example,dc=com</code> • For Microsoft Active Directory, specify the base DN in the following format: <code>dc=domain1,dc=local</code>. You will need to replace the <code>domain1</code> and <code>local</code> for your specific configuration. Microsoft Server provides a tool called <code>ldp.exe</code> which is useful for finding out and configuring the the LDAP structure of your server.
User Name Attribute	<p>The attribute field to use when loading the username. Examples:</p> <ul style="list-style-type: none"> • <code>cn</code> • <code>sAMAccountName</code>

Advanced Settings

Setting	Description
Enable Nested Groups	<p>Enable or disable support for nested groups. Some directory servers allow you to define a group as a member of another group. Groups in such a structure are called 'nested groups'. If you are using groups to manage permissions, you can create nested groups to allow inheritance of permissions from one group to its sub-groups.</p>
Use Paged Results	<p>Enable or disable the use of the LDAP control extension for simple paging of search results. If paging is enabled, the search will retrieve sets of data rather than all of the search results at once. Enter the desired page size – that is, the maximum number of search results to be returned per page when paged results are enabled. The default is 1000 results.</p>

Follow Referrals	Choose whether to allow the directory server to redirect requests to other servers. This option uses the node referral (JNDI lookup <code>java.naming.referral</code>) configuration setting. It is generally needed for Active Directory servers configured without proper DNS, to prevent a 'javax.naming.PartialResultException: Unprocessed Continuation Reference(s)' error.
------------------	--

User Schema Settings

Note: this section is only visible when **Copy User on Login** is enabled.

Setting	Description
Additional User DN	This value is used in addition to the base DN when searching and loading users. If no value is supplied, the subtree search will start from the base DN. Example: <ul style="list-style-type: none"> <code>ou=Users</code>
User Object Class	This is the name of the class used for the LDAP user object. Example: <ul style="list-style-type: none"> <code>user</code>
User Object Filter	The filter to use when searching user objects. Example: <ul style="list-style-type: none"> <code>(&(objectCategory=Person)(sAMAccountName=*))</code>
User Name RDN Attribute	The RDN (relative distinguished name) to use when loading the username. The DN for each LDAP entry is composed of two parts: the RDN and the location within the LDAP directory where the record resides. The RDN is the portion of your DN that is not related to the directory tree structure. Example: <ul style="list-style-type: none"> <code>cn</code>
User First Name Attribute	The attribute field to use when loading the user's first name. Example: <ul style="list-style-type: none"> <code>givenName</code>
User Last Name Attribute	The attribute field to use when loading the user's last name. Example: <ul style="list-style-type: none"> <code>sn</code>
User Display Name Attribute	The attribute field to use when loading the user's full name. Example: <ul style="list-style-type: none"> <code>displayName</code>
User Email Attribute	The attribute field to use when loading the user's email address. Example: <ul style="list-style-type: none"> <code>mail</code>

Group Schema Settings

Note: this section is only visible when both **Copy User on Login** and **Synchronise Group Memberships** are enabled.

Setting	Description
Additional Group DN	This value is used in addition to the base DN when searching and loading groups. If no value is supplied, the subtree search will start from the base DN. Example: <ul style="list-style-type: none"> <code>ou=Groups</code>

Group Object Class	This is the name of the class used for the LDAP group object. Examples: <ul style="list-style-type: none"> • <code>groupOfUniqueNames</code> • <code>group</code>
Group Object Filter	The filter to use when searching group objects. Example: <ul style="list-style-type: none"> • <code>(objectCategory=Group)</code>
Group Name Attribute	The attribute field to use when loading the group's name. Example: <ul style="list-style-type: none"> • <code>cn</code>
Group Description Attribute	The attribute field to use when loading the group's description. Example: <ul style="list-style-type: none"> • <code>description</code>

Membership Schema Settings

Note: this section is only visible when both **Copy User on Login** and **Synchronise Group Memberships** are enabled.

Setting	Description
Group Members Attribute	The attribute field to use when loading the group's members. Example: <ul style="list-style-type: none"> • <code>member</code>
User Membership Attribute	The attribute field to use when loading the user's groups. Example: <ul style="list-style-type: none"> • <code>memberOf</code>
Use the User Membership Attribute, when finding the user's group membership	Check this box if your directory server supports the group membership attribute on the user. (By default, this is the 'memberOf' attribute.) <ul style="list-style-type: none"> • If this box is checked, your application will use the group membership attribute on the user when retrieving the members of a given group. This will result in a more efficient retrieval. • If this box is not checked, your application will use the members attribute on the group ('member' by default) for the search.

Diagrams of Possible Configurations

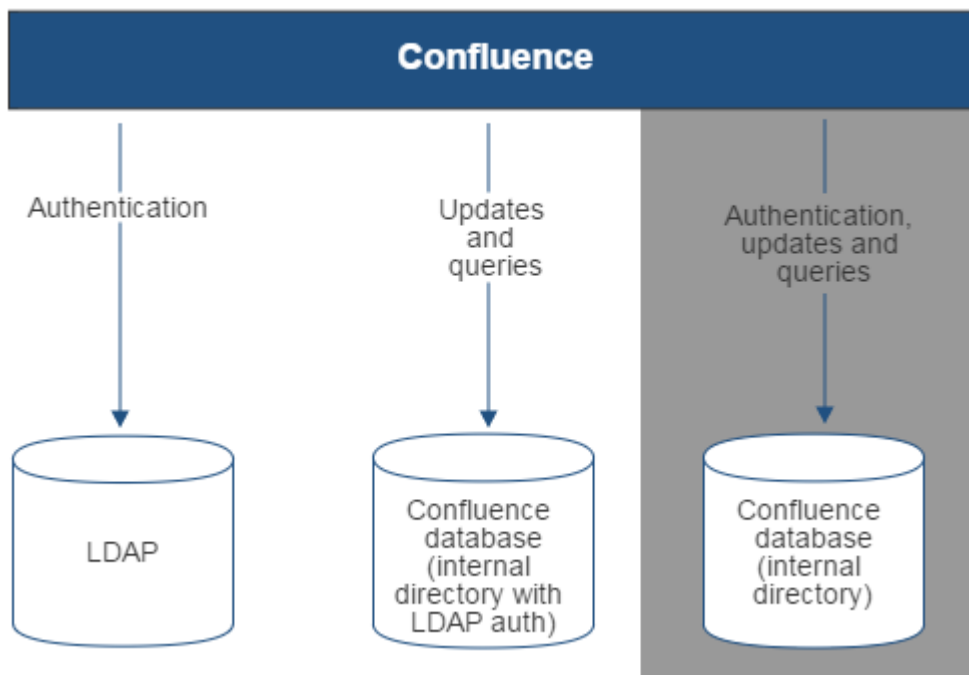


Diagram above: Confluence connecting to an LDAP directory for authentication only.

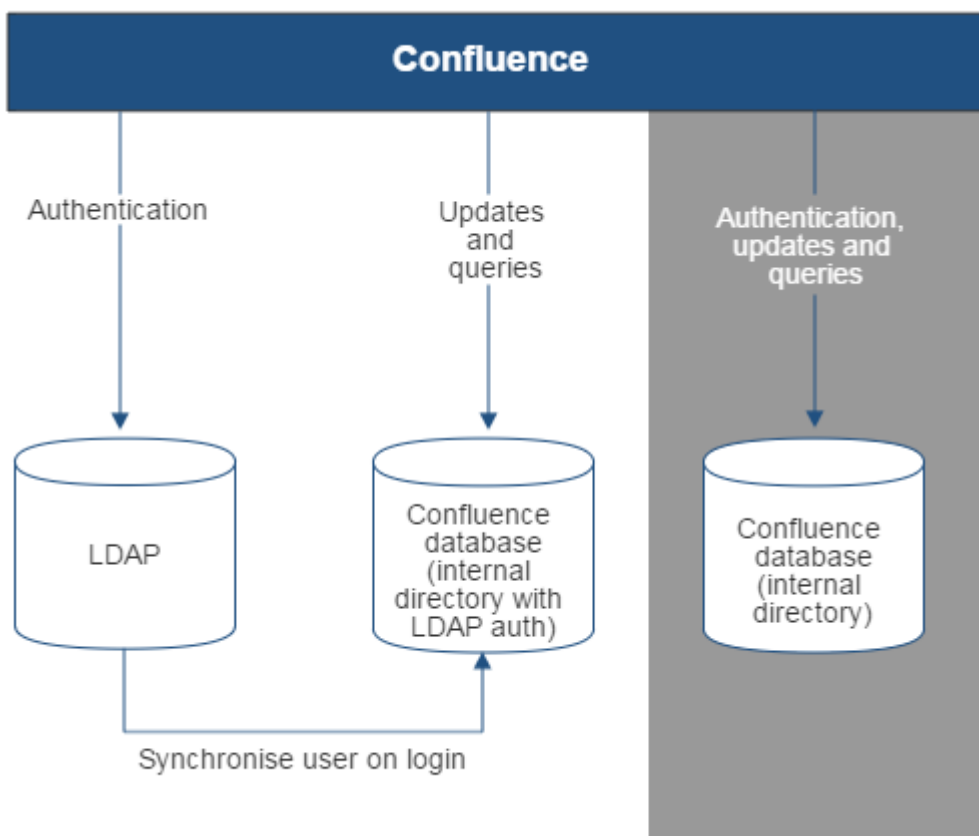


Diagram above: Confluence connecting to an LDAP directory for authentication only, with each user synchronised with the internal directory that is using LDAP authentication when they log in to Confluence.

Connecting to Crowd or JIRA for User Management

You can connect your Confluence application to Atlassian Crowd or to a JIRA application (version 4.3 or later) for management of users and groups, and for authentication.

Connecting Confluence to Crowd for User Management

Atlassian Crowd is an application security framework that handles authentication and authorization for your web-based applications. With Crowd you can integrate multiple web applications and user directories, with support for single sign-on (SSO) and centralized identity management. The Crowd Administration Console provides a web interface for managing directories, users and their permissions. See the [Crowd Administration Guide](#).

When to use this option: Connect to Crowd if you want to use the full Crowd functionality to manage your directories, users and groups. You can connect your Crowd server to a number of directories of all types that Crowd supports, including custom directory connectors.


On this page:

- [Connecting Confluence to Crowd for User Management](#)
- [Connecting Confluence to JIRA applications for User Management](#)
- [Diagrams of Some Possible Configurations](#)
- [Troubleshooting](#)

Related pages:

- [Configuring User Directories](#)

To connect Confluence to Crowd:

1. Go to your **Crowd Administration Console** and define the Confluence application to Crowd. See the Crowd documentation: [Adding an Application](#).
2. Choose the **cog icon**  , then choose **General Configuration** under Confluence Administration
3. Click '**User Directories**' in the left-hand panel.
4. **Add** a directory and select type '**Atlassian Crowd**'. Enter the settings as described below.
5. Save the directory settings.
6. Define the **directory order** by clicking the blue up- and down-arrows next to each directory on the '**User Directories**' screen. Here is a summary of how the directory order affects the processing:
 - Changes to users and groups will be made only in the first directory where the application has permission to make changes.
 - The order of the directories is the order in which they will be searched for users and groups (by default Confluence aggregates group membership from all directories, so the order does not impact membership itself).
 For details see [Managing Multiple Directories](#).
7. If required, configure Confluence to use Crowd for single sign-on (SSO) too. See the Crowd documentation: [Integrating Crowd with Atlassian Confluence](#).

Crowd Settings in Confluence

Setting	Description
Name	A meaningful name that will help you to identify this Crowd server amongst your list of directory servers. Examples: <ul style="list-style-type: none"> • Crowd Server • Example Company Crowd
Server URL	The web address of your Crowd console server. Examples: <ul style="list-style-type: none"> • http://www.example.com:8095/crowd/ • http://crowd.example.com
Application Name	The name of your application, as recognized by your Crowd server. Note that you will need to define the application in Crowd too, using the Crowd administration Console. See the Crowd documentation on adding an application .

Application Password	The password which the application will use when it authenticates against the Crowd framework as a client. This must be the same as the password you have registered in Crowd for this application. See the Crowd documentation on adding an application .
----------------------	--

Note: There is a known issue where the password is not saved in some instances

CONF-33979 - New JIRA/Crowd password not saved after test

[OPEN](#)

when configuring Confluence to use JIRA/Crowd as a external user directory.

Crowd Permissions

Setting	Description
Read Only	The users, groups and memberships in this directory are retrieved from Crowd and can only be modified via Crowd. You cannot modify Crowd users, groups or memberships via the application administration screens.
Read/Write	The users, groups and memberships in this directory are retrieved from Crowd. When you modify a user, group or membership via the application administration screens, the changes will be applied directly to Crowd. Please ensure that the application has modification permissions for the relevant directories in Crowd. See the Crowd documentation: Specifying an Application's Directory Permissions .

Advanced Crowd Settings

Setting	Description
Enable Nested Groups	Enable or disable support for nested groups. Before enabling nested groups, please check to see if the user directory or directories in Crowd support nested groups. When nested groups are enabled, you can define a group as a member of another group. If you are using groups to manage permissions, you can create nested groups to allow inheritance of permissions from one group to its sub-groups.
Synchronisation Interval (minutes)	Synchronisation is the process by which the application updates its internal store of user data to agree with the data on the directory server. The application will send a request to your directory server every x minutes, where 'x' is the number specified here. The default value is 60 minutes.

Connecting Confluence to JIRA applications for User Management

Note that the license tiers for your JIRA application and Confluence do not need to match to use this feature. For example, you can manage a Confluence 50 user license with JIRA Software, even if JIRA Software only has a 25 user license.

Subject to certain limitations, you can connect a number of Atlassian applications to a single JIRA application for centralised user management.

When to use this option: You can connect to a server running **JIRA 4.3** or later, **JIRA Software 7.0** or later, **JIRA Core 7.0** or later, or **JIRA Service Desk 3.0** or later. Choose this option as an alternative to Atlassian Crowd, for simple configurations with a limited number of users.

To connect Confluence to a JIRA application:

1. In your JIRA application go to



> **User Management > JIRA User Server.**

(For JIRA 6.4 and earlier go to your JIRA administration screen then **Users > JIRA User Server**)

- Click **Add Application**.
 - Enter the **application name** and **password** that Confluence will use when accessing JIRA.
 - Enter the **IP address** or addresses of your Confluence server. Valid values are:
 - A full IP address, e.g. 192.168.10.12.
 - A wildcard IP range, using CIDR notation, e.g. 192.168.10.1/16. For more information, see the introduction to [CIDR notation on Wikipedia](#) and [RFC 4632](#).
 - **Save** the new application.
2. Set up the JIRA user directory in Confluence:
 - Choose the **cog icon**

, then choose **General Configuration** under Confluence Administration
 - Click '**User Directories**' in the left-hand panel.
 - **Add** a directory and select type '**Atlassian JIRA**'.
 - Enter the settings as described below. When asked for the **application name** and **password**, enter the values that you defined for your Confluence application in the settings on JIRA.
 - Save the directory settings.
 - Define the **directory order** by clicking the blue up- and down-arrows next to each directory on the '**User Directories**' screen. Here is a summary of how the directory order affects the processing:
 - The order of the directories is the order in which they will be searched for users and groups.
 - Changes to users and groups will be made only in the first directory where the application has permission to make changes.

For details see [Managing Multiple Directories](#).
 3. In order to use Confluence, users must be a member of the `confluence-users` group or have Confluence 'can use' permission. Follow these steps to configure your Confluence groups in your JIRA application:
 - a. Add the `confluence-users` and `confluence-administrators` groups in your JIRA application.
 - b. Add your own username as a member of both of the above groups.
 - c. Choose one of the following methods to give your existing JIRA users access to Confluence:
 - Option 1: In your JIRA application, find the groups that the relevant users belong to. Add the groups as members of one or both of the above Confluence groups.
 - Option 2: Log in to Confluence using your JIRA account and go to the Confluence **Administration Console**. Click '**Global Permissions**' and assign the 'can use' permission to the relevant JIRA groups.

Ensure that you have added Confluence URL into JIRA Whitelist in JIRA Administration >> System >> Security >> Whitelist. For example: <https://confluence.atlassian.com/> or refer to this guide: [Configuring the whitelist](#).

JIRA Settings in Confluence

Setting	Description
Name	<p>A meaningful name that will help you to identify this JIRA server in the list of directory servers. Examples:</p> <ul style="list-style-type: none"> • JIRA Service Desk Server • My Company JIRA

Server URL	The web address of your JIRA server. Examples: <ul style="list-style-type: none">• <code>http://www.example.com:8080</code>• <code>http://jira.example.com</code>
Application Name	The name used by your application when accessing the JIRA server that acts as user manager. Note that you will also need to define your application to that JIRA server, via the ' Other Applications ' option in the 'Users, Groups & Roles' section of the 'Administration' menu.
Application Password	The password used by your application when accessing the JIRA server that acts as user manager.

JIRA Permissions

Setting	Description
Read Only	The users, groups and memberships in this directory are retrieved from the JIRA server that is acting as user manager. They can only be modified via that JIRA server.

Advanced JIRA Settings

Setting	Description
Enable Nested Groups	Enable or disable support for nested groups. Before enabling nested groups, please check to see if nested groups are enabled on the JIRA server that is acting as user manager. When nested groups are enabled, you can define a group as a member of another group. If you are using groups to manage permissions, you can create nested groups to allow inheritance of permissions from one group to its sub-groups.
Synchronisation Interval (minutes)	Synchronisation is the process by which the application updates its internal store of user data to agree with the data on the directory server. The application will send a request to your directory server every x minutes, where 'x' is the number specified here. The default value is 60 minutes.

Diagrams of Some Possible Configurations

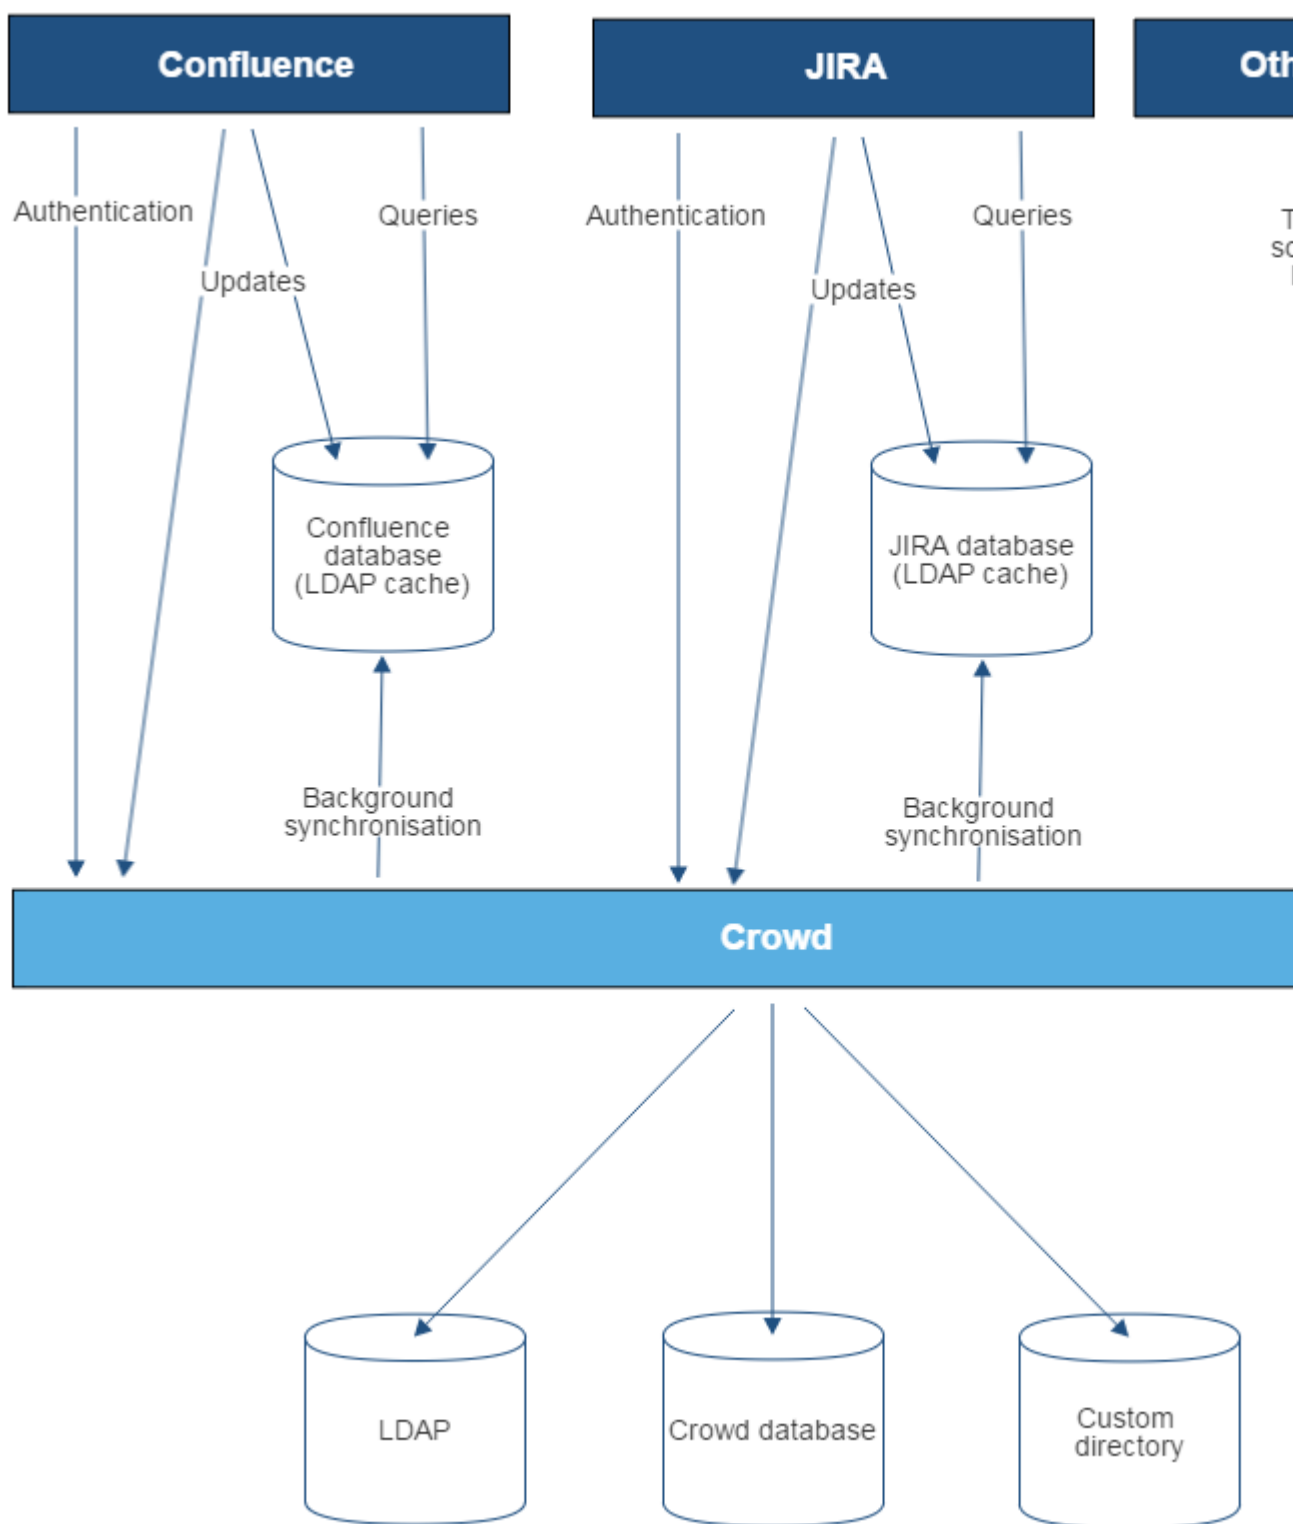


Diagram above: Confluence, JIRA and other applications connecting to Crowd for user management.

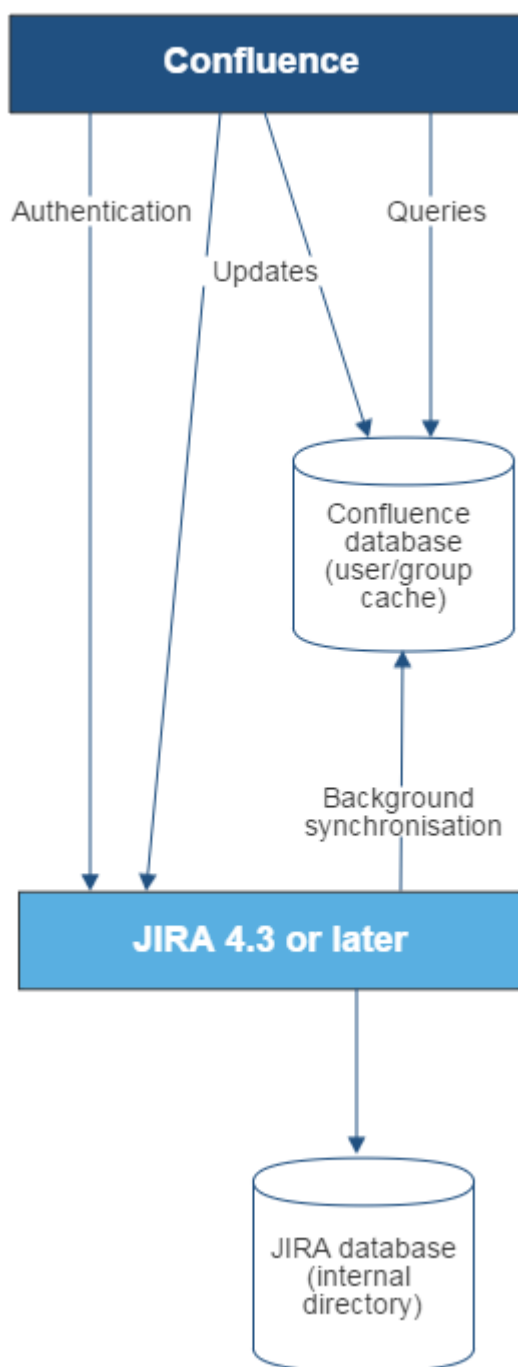


Diagram above: Confluence connecting to JIRA for user management.

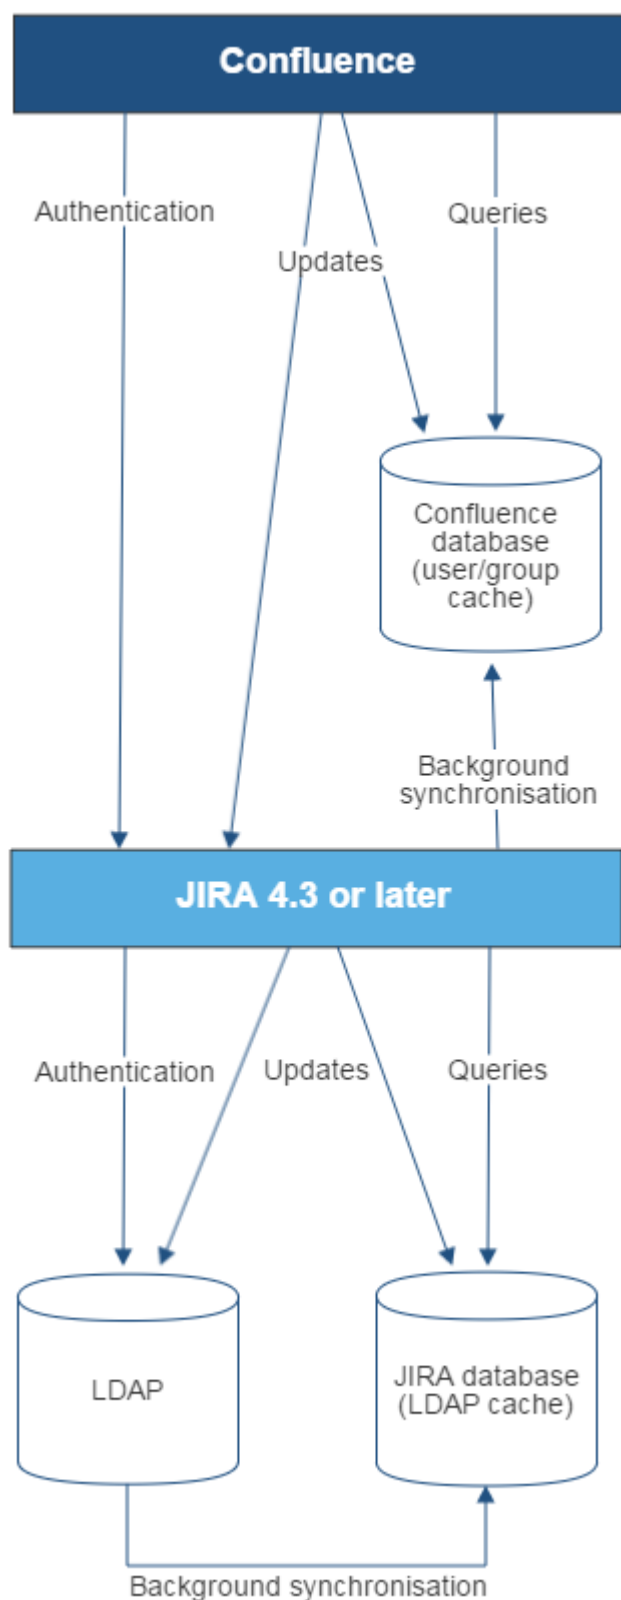


Diagram above: Confluence connecting to JIRA for user management, with JIRA in turn connecting to LDAP.

Troubleshooting

Below are some error messages you may encounter. If you run into problems, you should turn on WARN logging for the relevant class. See [Configuring Logging](#).

Error	Message	Cause
-------	---------	-------

error.jirabaseurl.connection.refused	Connection refused. Check if an instance of JIRA is running on the given url	This may be because: <ul style="list-style-type: none"> • JIRA url is incorrect • JIRA instance is not running on the specified url. • JIRA instance running on the specified url is not 4.3 or later.
error.applicationlink.connection.refused	Failed to establish application link between JIRA server and Confluence server.	Unable to create an application link between JIRA and Confluence. This may be because: <ul style="list-style-type: none"> • Confluence or JIRA url is incorrect • the instance is not running on the specified url • credentials are incorrect. Refer to the Confluence log files for further troubleshooting information.
error.jirabaseurl.not.valid	This is not a valid url for a JIRA application.	A runtime exception has occurred. Refer to the Confluence log files for further troubleshooting information.

Reverting from Crowd or JIRA applications to Internal User Management

If your Confluence site currently uses Crowd or a JIRA application for user management, you can revert to internal user management as described below. If your Confluence instance has only a few users, it is easier to recreate the users and groups in Confluence manually. If you have a large number of users and groups, it is more efficient to migrate the relevant users and groups into the Confluence Internal directory.

On this page:

- [Option 1 – Manually Recreate Users and Groups in Confluence](#)
- [Option 2 – Transfer Crowd/JIRA application Users and Groups to the Confluence Database](#)

Both options provided below will reset the affected users' passwords. When done, be sure to notify them to use the 'Reset My Password' link on the Confluence log in page before they attempt to log in.

Option 1 – Manually Recreate Users and Groups in Confluence

Use this option if you have only a few users and groups.

1. Log in to Confluence as a Confluence system administrator.
2. Go to the user directories administration screen and move the **internal** directory to the top of the list of directories, by clicking the arrows in the '**Order**' column.
3. Make sure that you have at least one user from the **internal** directory in each of the `confluence-users` and `confluence-administrators` groups.
4. Make sure that you have a username in the **internal** directory with Confluence system administrator permissions.
 - If you do not have such a user, add a new one now, and log out of Confluence.
 - Log back in as the user you just added, and go back to the user directories administration screen.
5. Disable the '**Atlassian Crowd**' directory.
6. Manually add the required users and groups in Confluence. They will be added to the internal directory, because you have moved it to the top of the list of directories.
 - If you have assigned Confluence permissions to a group which exists in your JIRA application,

- you must create a group in Confluence with the same name.
 - If a user who exists in your JIRA application has created content or has had permissions assigned to them in Confluence, you must also create that user in Confluence.
7. Add the users to the required groups.

Option 2 – Transfer Crowd/JIRA application Users and Groups to the Confluence Database

Use this option to migrate External Application (Crowd or JIRA applications) users into the Confluence database. You need a knowledge of SQL to perform this task.

The SQL commands given below are tailored for **MySQL**. If you are using a database other than MySQL, you will need to modify the SQL to work in your database.

Step 1. Create Backups

Creating backups is the only way to restore your data if something goes wrong.

1. From Confluence, create a full XML site backup including attachments.
2. Stop Confluence.
3. Make a backup copy of the Confluence home and installation directories.
4. Repeat the above steps for your External Application.
5. From your MySQL administration tool, create a database backup for the Crowd/JIRA application and Confluence databases.

Step 2. Replace Confluence User Management

Use the SQL below to move groups and users from your External Application to Confluence by transferring table content. The SQL provided is specific to MySQL and must be modified for other databases. Find the IDs for your Directories

1. Run the following command and take note of the resulting number. It will be referenced throughout the following instructions as <Confluence Internal ID>.

```
select id from cwd_directory where directory_name='Confluence Internal Directory';
```

2. From the User Directories administration page, find the name of the directory who's users/groups you want to move. Run the following command and take note of the resulting number. It will be referenced throughout the following instructions as <External Application ID>.

```
select id from cwd_directory where directory_name='<External Directory Name>';
```

Move Groups to Confluence

1. It is possible that you have several groups in your Internal Directory that have the same name as groups in your External Application. To find these, run:

```
select distinct a.id, a.directory_id, a.group_name, d.directory_name from cwd_group a join cwd_group b on a.group_name=b.group_name join cwd_directory d on d.id=a.directory_id where a.directory_id != b.directory_id;
```

- a. If you have results from the previous query, for each of the group names that have duplicates, find the id for the group in the Confluence Internal Directory (<internal group id>) and the External Application (<external group id>). Run the following:

```
update cwd_group_attribute set group_id=<internal group id>,
directory_id=<Confluence Internal Id> where group_id=<external group
id>;
update cwd_membership set child_group_id=<internal group id> where
child_group_id=<external group id>;
update cwd_membership set parent_id=<internal group id> where
parent_id=<external group id>;
delete from cwd_group where id=<external group id>;
```

2. Move all the groups in the External Application to the Confluence Internal Directory.

```
update cwd_group set directory_id=<Confluence Internal ID> where
directory_id=<External Application ID>;
```

Move Users to Confluence

1. It is possible that you have several users in your Internal Directory that have the same name as users in your External Application. To find these, run:

```
select distinct a.id, a.directory_id, a.user_name, d.directory_name from
cwd_user a join cwd_user b on a.user_name=b.user_name join cwd_directory d
on d.id=a.directory_id where a.directory_id != b.directory_id;
```

- a. If you have results from the previous query, for each of the user names that have duplicates, find the id for the user in the Confluence Internal Directory (<internal user id>) and the External Application (<external user id>). Run the following:

```
update cwd_membership set child_user_id=<internal user id> where
child_user_id=<external user id>;
update cwd_user_credential_record set user_id=<internal user id>
where user_id=<external user id>;
update cwd_user_attribute set user_id=<internal user id>,
directory_id=<Confluence Internal ID> where user_id=<external user
id>;
delete from cwd_user where id=<external user id>;
```

2. Move all the users in the External Application to the Confluence Internal Directory.

```
update cwd_user set directory_id=<Confluence Internal ID> where
directory_id=<External Application ID>;
```

Delete the External Application directory

1. You need to change the order of your directories so that the Internal directory is at the top, and active.
 - a. If you have only two directories - the Internal and the External Application directory you are deleting, then do the following:

```
update cwd_app_dir_mapping set list_index = 0 where directory_id =
<Confluence Internal ID>;
```

- b. If you have more than two directories, you need to rearrange them so the Internal Directory is at

the top (list_index 0) and the External Application directory you are deleting is at the bottom.

- List the directories and their order using

```
select d.id, d.directory_name, m.list_index from cwd_directory d
join cwd_app_dir_mapping m on d.id=m.directory_id order by
m.list_index;
```

- Change the list indexes so that they are in the order you want. Directory order can be rearranged using

```
update cwd_app_dir_mapping set list_index = <position> where
directory_id = <directory id>;
```

c. Check that the internal directory is enabled.

- List the internal directory. An enabled directory will have its 'active' column set to 'T'

```
select id, directory_name, active from cwd_directory where id =
<Internal Directory id>;
```

- If the internal directory is not active, activate it by

```
update cwd_directory set active = 'T' where id = <Internal
Directory id>;
```

2. When the directories are ordered correctly, delete the External Application directory from the directory order:

```
delete from cwd_app_dir_operation where app_dir_mapping_id = (select id
from cwd_app_dir_mapping where directory_id = <External Application ID>);
delete from cwd_app_dir_mapping where directory_id = <External Application
ID>;
```

3. The External Application directory is referenced in several other tables in the database. You need to remove the remaining references to it:

```
delete from cwd_directory_attribute where directory_id=<External
Application ID>;
delete from cwd_directory_operation where directory_id=<External
Application ID>;
```

4. All references to the External Directory should now have been removed. Delete the directory using:

```
delete from cwd_directory where id = <External Application ID>;
```

Reset passwords

1. All users who were in the External Directory you deleted, including admins, will be unable to log in. Their passwords need to be reset by choosing the 'Forgot your password?' link on the login page. Alternatively, use the instructions at [Restore Passwords To Recover Admin User Rights](#) to reset the administrator password, then set the users' passwords for them via the Manage Users page in the administration screen.

Managing Multiple Directories

This page describes what happens when you have defined more than one user directory in Confluence. For example, you may have an internal directory and you may also connect to an LDAP directory server and/or other types of user directories. When you connect to a new directory server, you also need to define the **directory order**.

Avoid duplicate usernames across directories. If you are connecting to more than one user directory, we recommend that you ensure the usernames are unique to one directory. For example, we do not recommend that you have a user `jsmith` in both 'Directory1' and 'Directory2'. The reason is the potential for confusion, especially if you swap the order of the directories. Changing the directory order can change the user that a given username refers to.





Overview

Here is a summary of how the directory order affects the processing:

- The order of the directories is the order in which they will be searched for users and groups.
- Changes to users and groups will be made only in the first directory where the application has permission to make changes.

Configuring the Directory Order

You can change the order of your directories as defined to Confluence. Select '**User Directories**' from the Confluence Administration Console and click the blue up- and down-arrows next to each directory.

Directory Name	Type	Order
Confluence Internal Directory	Internal	 
OpenLDAP	OpenLDAP (Read-Write)	 

Notes:

- Please read the rest of this page to understand what effect the directory order will have on authentication (login) and permissions in Confluence, and what happens when you update users and groups in Confluence.

Effect of Directory Order

This section summarises the effect the order of the directories will have on login and permissions, and on the updating of users and groups.

Login

The directory order is significant during the authentication of the user, in cases where the same user exists in multiple directories. When a user attempts to log in, the application will search the directories in the order specified, and will use the credentials (password) of the *first occurrence of the user* to validate the login attempt.

Permissions

Aggregating membership (default)

The directory order **is not** significant when granting the user permissions based on group membership as Confluence uses an aggregating membership scheme by default. If the same username exists in more than one directory, the application will aggregate (combine) group membership from all directories where the

On this page:

- [Overview](#)
- [Configuring the Directory Order](#)
- [Effect of Directory Order](#)
 - [Login](#)
 - [Permissions](#)
 - [Updating Users and groups](#)

username appears.

Example:

- You have connected two directories: The Customers directory and the Partners directory.
- The Customers directory is first in the directory order.
- A username `jsmith` exists in both the Customers directory and the Partners directory.
- The user `jsmith` is a member of group `G1` in the Customers directory and group `G2` in the Partners directory.
- The user `jsmith` will have permissions based on membership of both `G1` and `G2` regardless of the directory order.

For administrators upgrading to Confluence 5.7 or later:

How group memberships are determined for users that belong to multiple user directories (such as LDAP, Active Directory, Crowd) changed in Confluence 5.7. Group memberships are now aggregated from *all* directories, not the first one the user appears in. In most cases, this change will have no impact as users generally only exist in one directory, or their memberships are correctly synchronised between user directories. In some rare cases, where group memberships are out of synch, the change may lead to users gaining permissions to view spaces and pages (if they are a member a group in a user directory that was previously being ignored by Confluence).

▼ [Here's an example scenario...](#)



This is Issac. Something went wrong a while ago, so he's got the same username in two user directories, but belongs to different groups.

Right now, the user directories in his organisation's Confluence site look like this:

Confluence User Directories

Active Directory	1
Confluence internal directory	2

and Issac's group memberships in each directory looks like this:

1 Active Directory confluence-users sydney		Confluence internal directory 2 confluence-users developers sydney
---	--	---

The 'Dev Team' page is restricted to the developers group.

- In Confluence 5.6 and earlier, Issac couldn't see this page as we determined his group membership from Active Directory - because it's the first directory in the list it had the highest priority.
- In Confluence 5.7 and beyond, Issac will see the page because we determine his group membership from *all* directories, not just the highest one.

To Confluence his group membership looks like this:



Confluence group membership
 confluence-users
 developers
 sydney

This means after the 5.7 upgrade he can see any pages and spaces that are restricted to the 'developers'

group.

Non-aggregating membership

It is possible to use the REST API to tell Confluence to use a non-aggregating membership scheme as follows:

Turning on non-aggregating membership...

The REST resource supported JSON and XML. You'll need to be a system administrator and logged in to do this.

```
# To GET the current setting
curl -H 'Accept: application/json' -u <username>
<base-url>/rest/crowd/latest/application

# To PUT the setting
curl -H 'Content-type: application/json' -X PUT -d
'{"membershipAggregationEnabled":true}' -u <username>
<base-url>/rest/crowd/latest/application
```

If you've chosen non-aggregating membership, the directory order is significant. If the same username exists in more than one directory, the application will look for group membership only in the first directory where the username appears, based on the directory order.

Example:

- You have connected two directories: The Customers directory and the Partners directory.
- The Customers directory is first in the directory order.
- A username `jsmith` exists in both the Customers directory and the Partners directory.
- The user `jsmith` is a member of group `G1` in the Customers directory and group `G2` in the Partners directory.
- The user `jsmith` will have permissions based on membership of `G1` only, not `G2`.

Updating Users and groups

If you update a user or group via the application's administration screens, the update will be made in the first directory where the application has write permissions.

Example 1:

- You have connected two directories: The Customers directory and the Partners directory.
- The application has permission to update both directories.
- The Customers directory is first in the directory order.
- A username `jsmith` exists in both the Customers directory and the Partners directory.
- You update the email address of user `jsmith` via the application's administration screens.
- The email address will be updated in the Customers directory only, not the Partners directory.

Example 2:

- You have connected two directories: A read/write LDAP directory and the internal directory.
- The LDAP directory is first in the directory order.
- All new users will be added to the LDAP directory. It is not possible to add a new user to the internal directory.

RELATED TOPICS

Configuring User Directories

- [Configuring the Internal Directory](#)
- [Connecting to an LDAP Directory](#)
- [Connecting to an Internal Directory with LDAP Authentication](#)
- [Connecting to Crowd or JIRA for User Management](#)
- [Managing Multiple Directories](#)

- [Managing Nested Groups](#)
- [Synchronising Data from External Directories](#)
- [Diagrams of Possible Configurations for User Management](#)
- [User Management Limitations and Recommendations](#)
- [Requesting Support for External User Management](#)
- [Disabling the Built-In User Management](#)

Managing Nested Groups

Some directory servers allow you to define a group as a member of another group. Groups in such a structure are called 'nested groups'. If you are using groups to manage permissions, you can create nested groups to allow inheritance of permissions from one group to its sub-groups.

This page describes how Confluence handles nested groups that exist in one or more of your directory servers.

Enabling Nested Groups

You can enable or disable support for nested groups on each directory individually. Go to the '**User Directories**' section of the Confluence Administration Console, **edit** the directory and select '**Enable Nested Groups**'. See [Configuring User Directories](#).

Notes:

- Before enabling nested groups for a specific directory type in Confluence, please make sure that your directory server supports nested groups.
- Please read the rest of this page to understand what effect nested groups will have on authentication (login) and permissions in Confluence, and what happens when you update users and groups in Confluence.

Effect of Nested Groups

This section summarises the effect nested groups will have on login and permissions, and on the viewing and updating of users and groups.

Login

When a user logs in, they will be allowed access to the application if they belong to an authorised group or any of its sub-groups.

Permissions

The user will be allowed access to a function if they belong to a group that has the necessary permissions, or if they belong to any of its sub-groups.

Viewing Lists of Group Members

If you ask to view the members of a group, you will see all users who are members of the group and all users belonging its sub-groups, consolidated into one list. We call this a '**flattened**' list.

You cannot view or edit the nested groups themselves. You will not be able to see that one group is a

On this page:

- [Enabling Nested Groups](#)
- [Effect of Nested Groups](#)
 - [Login](#)
 - [Permissions](#)
 - [Viewing Lists of Group Members](#)
 - [Adding and Updating Group Memberships](#)
- [Examples](#)
 - [Example 1: User is Member of Sub-Group](#)
 - [Example 2: Sub-Groups as Members of the 'jira-developers' group](#)
 - [Example 3: Sub-Groups as Members of the 'confluence-users' group](#)
- [Notes](#)

Related pages:

- [Configuring User Directories](#)

member of another group.

Adding and Updating Group Memberships

If you add a user to a group, the user is added to the named group and not to any other groups.

If you try to remove a user from a [flattened](#) list, the following will happen:

- If the user is a member of the top group in the hierarchy (tree) of groups contained in the flattened list, the user will be removed from the group.
- Otherwise, you will see an error message stating that the user is not a direct member of the group.

Examples

Example 1: User is Member of Sub-Group

Let's assume that the following two groups exist in your directory server:

- `staff`
- `marketing`

Memberships:

- The `marketing` group is a member of the `staff` group.
- User `jsmith` is a member of `marketing`.

You will see that `jsmith` is a member of both `marketing` and `staff`. You will not see that the two groups are nested. If you assign permissions to the `staff` group, then `jsmith` will get those permissions.

Example 2: Sub-Groups as Members of the 'jira-developers' group

In an LDAP directory server, we have groups '**engineering-group**' and '**techwriters-group**'. We want to grant both groups developer-level access to our JIRA site. We will assume you have a group in your JIRA site called '**jira-developers**' that have developer-level access.

- Add a group called '**jira-developers**'.
- Add the '**engineering-group**' as a sub-group of '**jira-developers**'.
- Add the '**techwriters-group**' as a sub-group of '**jira-developers**'.

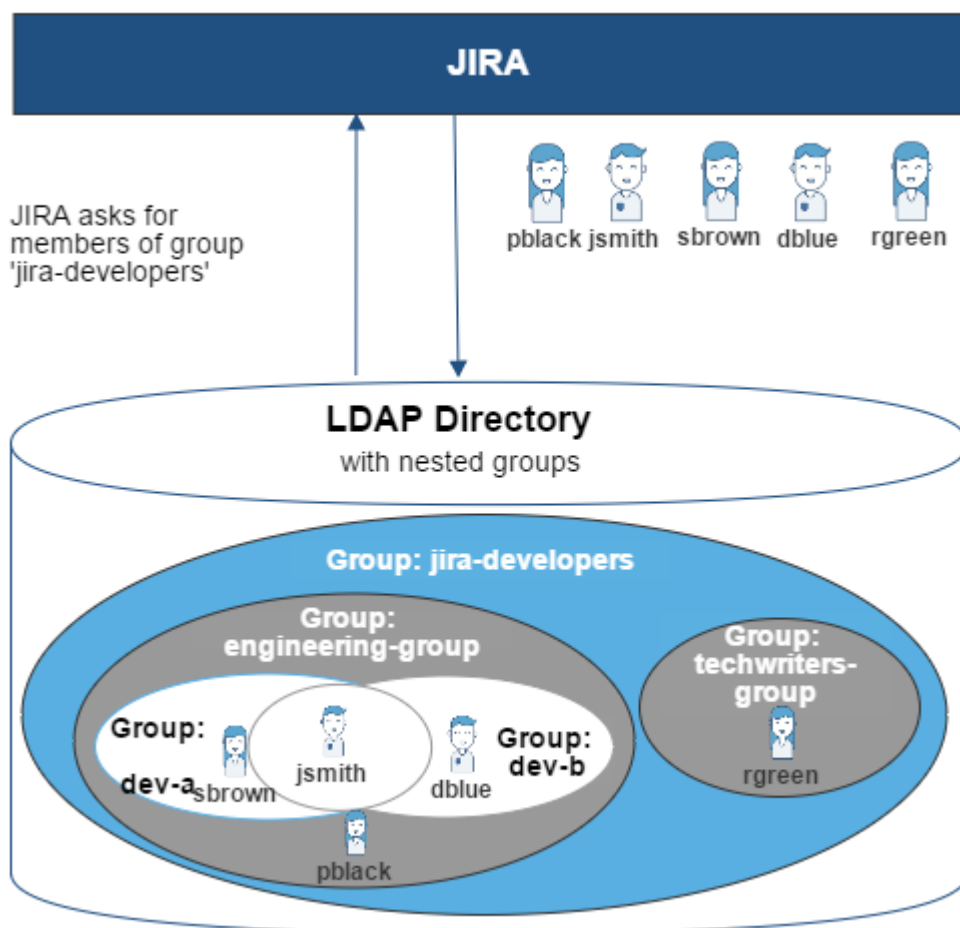
Group memberships are now:

- **jira-developers** — sub-groups: **engineering-group**, **techwriters-group**
- **engineering-group** — sub-groups: **dev-a**, **dev-b**; users: **pblack**
- **dev-a** — users: **jsmith**, **sbrown**
- **dev-b** — users: **jsmith**, **dblue**
- **techwriters-group** — users: **rgreen**

When the JIRA application requests a list of users in the '**jira-developers**' group, it will receive the following list:

- **pblack**
- **jsmith**
- **sbrown**
- **dblue**
- **rgreen**

Diagram: Sub-groups as members of the 'jira-developers' group



Example 3: Sub-Groups as Members of the 'confluence-users' group

In an LDAP directory server, we have groups '**engineering-group**' and '**payroll-group**'. We want to grant both groups access to our Confluence site.

- Add a group called '**confluence-users**'.
- Add the '**engineering-group**' as a sub-group of '**confluence-users**'.
- Add the '**payroll-group**' as a sub-group of '**confluence-users**'.

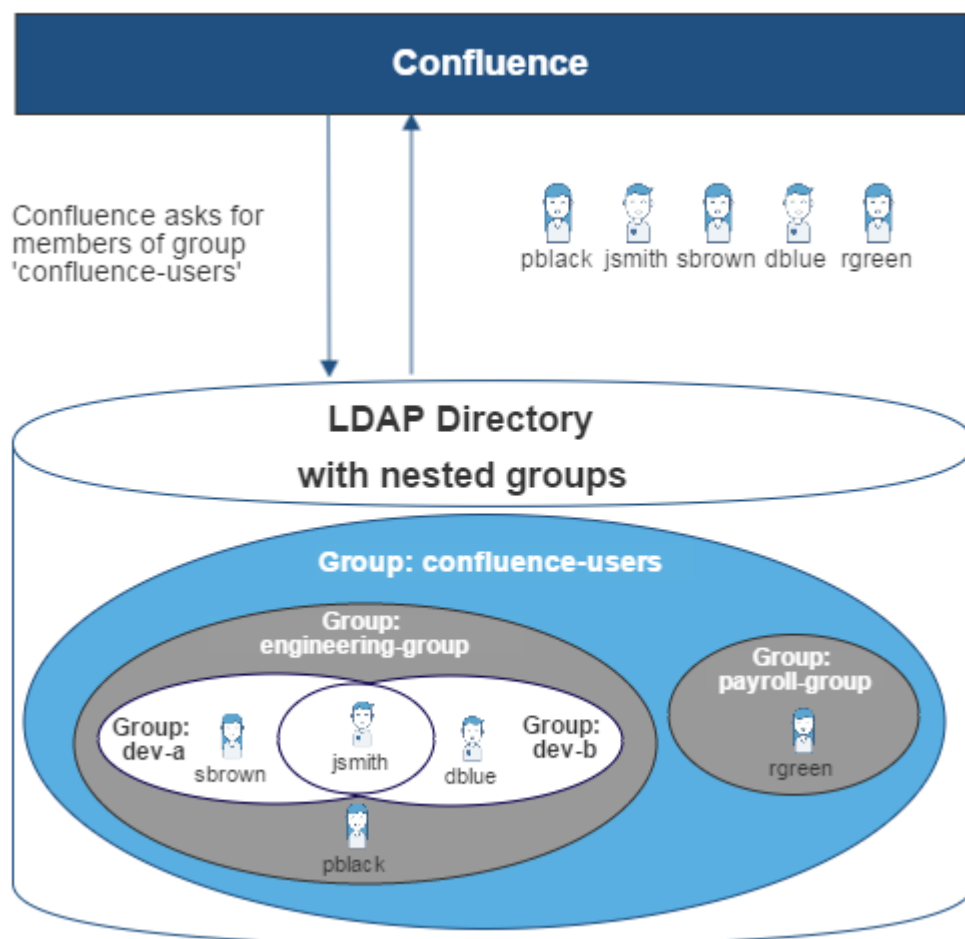
Group memberships are now:

- **confluence-users** — sub-groups: **engineering-group**, **payroll-group**
- **engineering-group** — sub-groups: **dev-a**, **dev-b**; users: **pblack**
- **dev-a** — users: **jsmith**, **sbrown**
- **dev-b** — users: **jsmith**, **dblue**
- **payroll-group** — users: **rgreen**

When Confluence requests a list of users in the '**confluence-users**' group, it will receive the following list:

- **pblack**
- **jsmith**
- **sbrown**
- **dblue**
- **rgreen**

Diagram: Sub-groups as members of the 'confluence-users' group



Notes

- **Possible impact on performance.** Enabling nested groups may result in slower user searches.
- **Definition of nested groups in LDAP.** In an LDAP directory, a nested group is defined as a child group entry whose DN (Distinguished Name) is referenced by an attribute contained within a parent group entry. For example, a parent group '**Group One**' might have an `objectClass=group` attribute and one or more `member=DN` attributes, where the DN can be that of a user *or* that of a group elsewhere in the LDAP tree:

```
member=CN=John Smith,OU=Users,OU=OrgUnitA,DC=sub,DC=domain
member=CN=Group Two,OU=OrgUnitBGroups,OU=OrgUnitB,DC=sub,DC=domain
```

Synchronising Data from External Directories

For certain directory types, Confluence stores a cache of directory information (users and groups) in the application database, to ensure fast recurrent access to user and group data. A synchronisation task runs periodically to update the internal cache with changes from the external directory.

On this page:

- [Affected Directory Types](#)
- [How it Works](#)
- [Finding the Time Taken to Synchronise](#)
- [Manually Synchronising the Cache](#)
- [Configuring the Synchronisation Interval](#)

Related pages:

- [Configuring User Directories](#)

Affected Directory Types

Data caching and synchronisation apply to the following user directory types:

- **LDAP** (Microsoft Active Directory and all supported LDAP directories) where permissions are set to **read only**.
- **LDAP** (Microsoft Active Directory and all supported LDAP directories) where permissions are set to **read only, with local groups**.
- **LDAP** (Microsoft Active Directory and all supported LDAP directories) where permissions are set to **read/write**.
- **Atlassian Crowd**.
- **Atlassian JIRA**.

Data caching and synchronisation do not occur for the following user directory types:

- **LDAP** (Microsoft Active Directory and all supported LDAP directories) where permissions are set to **authentication only, with local groups**.
- **Internal Directory with LDAP Authentication**.
- **Internal Directory**.

How it Works

Here is a summary of the caching functionality:

- The caches are held in the application database.
- When you connect a new external user directory to the application, a synchronisation task will start running in the background to copy all the required users, groups and membership information from the external directory to the application database. This task may take a while to complete, depending on the size and complexity of your user base.
- Note that a user will not be able to log in until the synchronisation task has copied that user's details into the cache.
- A periodic synchronisation task will run to update the database with any changes made to the external directory. The default synchronisation interval, or polling interval, is one hour (60 minutes). You can change the synchronisation interval on the directory configuration screen.
- You can manually synchronise the cache if necessary.
- If the external directory permissions are set to read/write: Whenever an update is made to the users, groups or membership information via the application, the update will also be applied to the cache and the external directory immediately.
- All authentication happens via calls to the external directory. When caching information from an external directory, the application database does not store user passwords.
- All other queries run against the internal cache.




Finding the Time Taken to Synchronise

The '**User Directories**' screen shows information about the last synchronisation operation, including the length of time it took.

Manually Synchronising the Cache

You can manually synchronise the cache by clicking '**Synchronise**' on the '**User Directories**' screen. If a synchronisation operation is already in progress, you cannot start another until the first has finished.

Screen snippet: User directories, showing information about synchronisation

OpenLDAP	OpenLDAP (Read-Write)	 	Disable Edit Synchronise Last synchronised at 14/01/11 3:07 PM (took 65s).
Crowd	Atlassian Crowd	 	Disable Edit Synchronise Last synchronised at 14/01/11 2:39 PM (took 0s).

Configuring the Synchronisation Interval

Note: The option to configure the synchronisation interval for Crowd and JIRA directories is available in **Confluence 3.5.3 and later**. Earlier versions of Confluence allow you to configure the interval for LDAP directories only.

You can set the '**Synchronisation Interval**' on the directory configuration screen. The synchronisation interval is the period of time to wait between requests for updates from the directory server.

The length you choose for your synchronisation interval depends on:

- The length of time you can tolerate stale data.
- The amount of load you want to put on the application and the directory server.
- The size of your user base.

If you synchronise more frequently, then your data will be more up to date. The downside of synchronising more frequently is that you may overload your server with requests.

If you are not sure what to do, we recommend that you start with an interval of 60 minutes (this is the default setting) and reduce the value incrementally. You will need to experiment with your setup.

Diagrams of Possible Configurations for User Management

The aim of these diagrams is to help people understand each directory type at a glance. We have kept the diagrams simple and conceptual, with just enough information to be correct.

Some things that we do **not** attempt to show:

- In most cases, we do not attempt to show that you can have multiple directory types mapped to Confluence at the same time. We illustrate that fact in just the first two LDAP diagrams.
- We have not included a diagram for Confluence's legacy connection to JIRA database.
- We do not attempt to show all of the possible configurations and layered connections that are available now that you can use JIRA as a directory manager.

On this page:

- [Confluence Internal Directory](#)
- [Confluence with Read/Write Connection to LDAP](#)
- [Confluence with Read-Only Connection to LDAP, with Local Groups](#)
- [Confluence Internal Directory with LDAP Authentication](#)
- [Confluence with LDAP Authentication, Copy Users on First Login](#)
- [Confluence Connecting to JIRA](#)
- [Confluence Connecting to JIRA and JIRA Connecting to LDAP](#)
- [Confluence and JIRA Connecting to Crowd](#)

Related pages:

- [Configuring User Directories](#)

Confluence Internal Directory

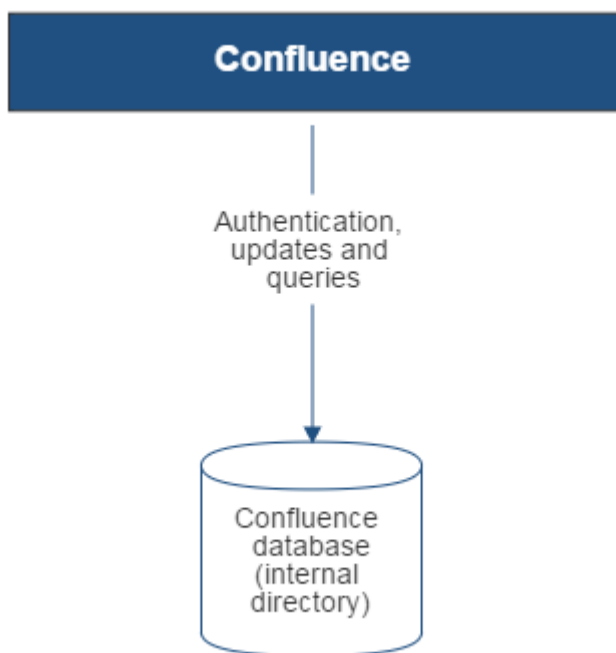


Diagram above: Confluence using its internal directory for user management.

Confluence with Read/Write Connection to LDAP

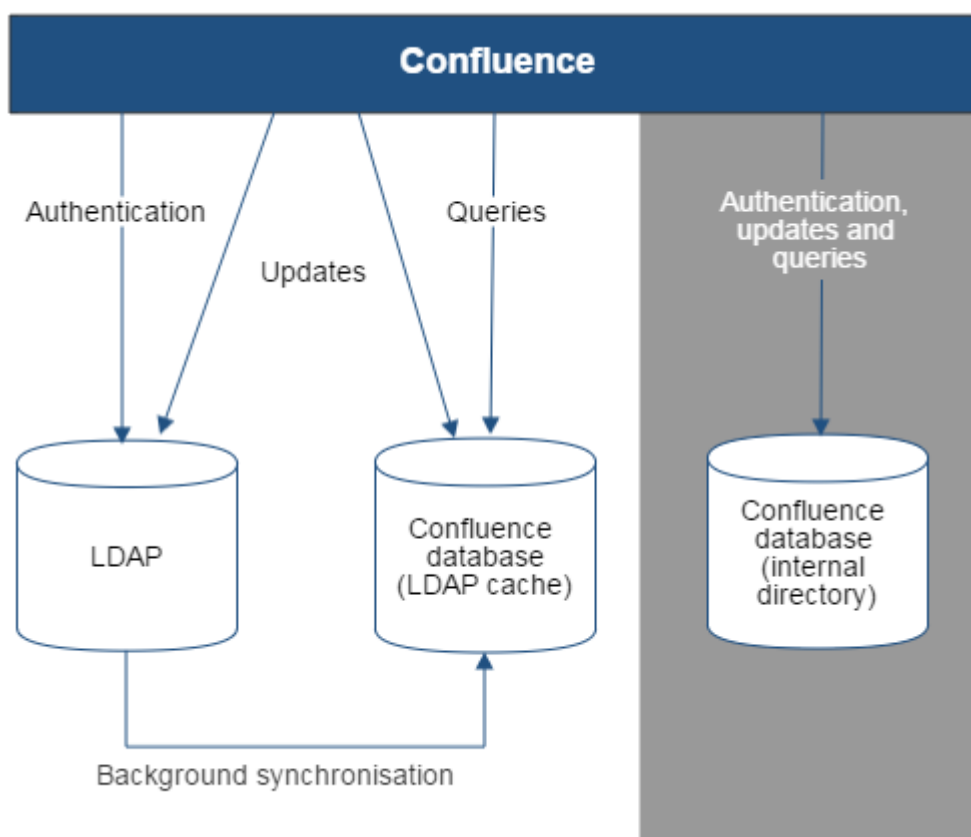


Diagram above: Confluence connecting to an LDAP directory.

Confluence with Read-Only Connection to LDAP, with Local Groups

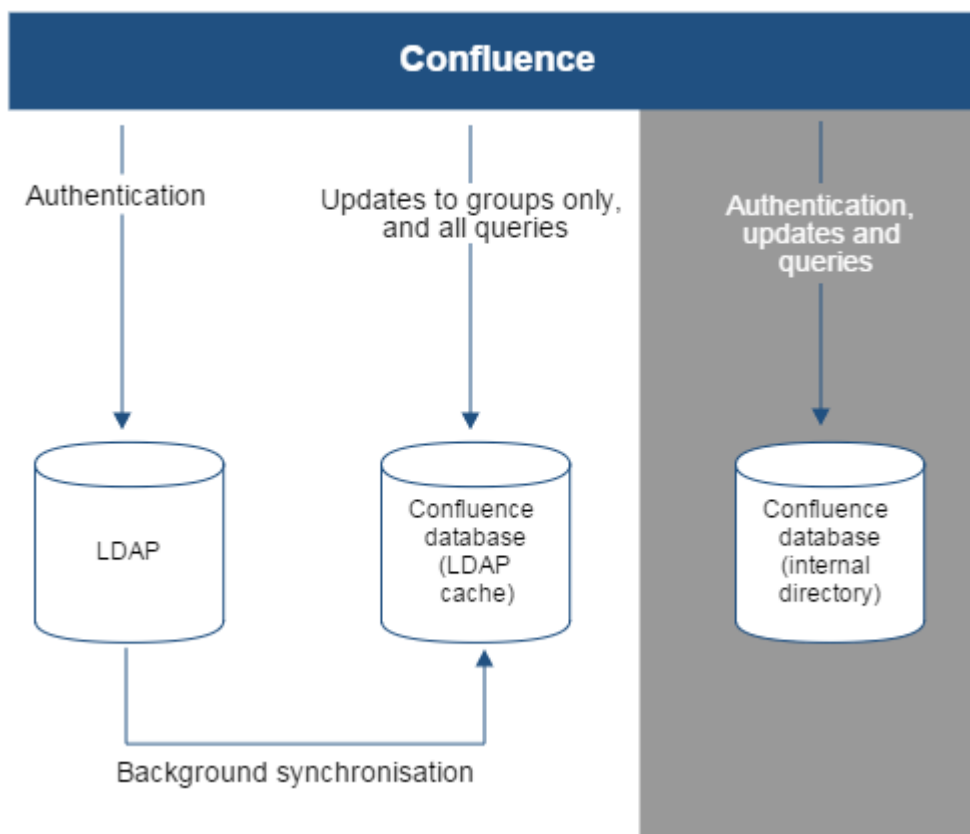


Diagram above: Confluence connecting to an LDAP directory with permissions set to read only and local groups.

Confluence Internal Directory with LDAP Authentication

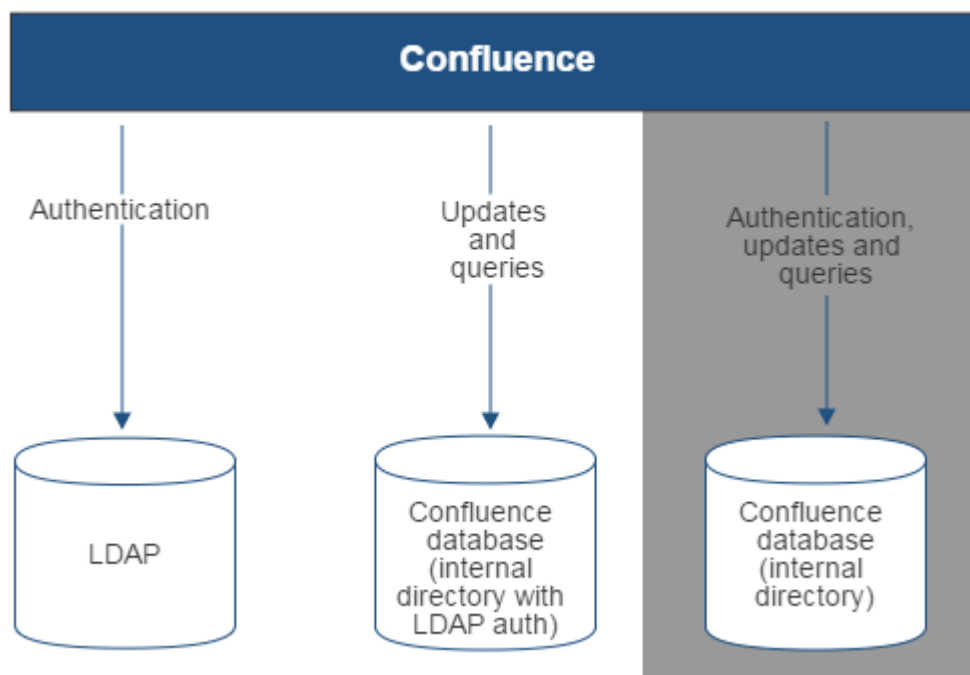


Diagram above: Confluence connecting to an LDAP directory for authentication only.

Confluence with LDAP Authentication, Copy Users on First Login

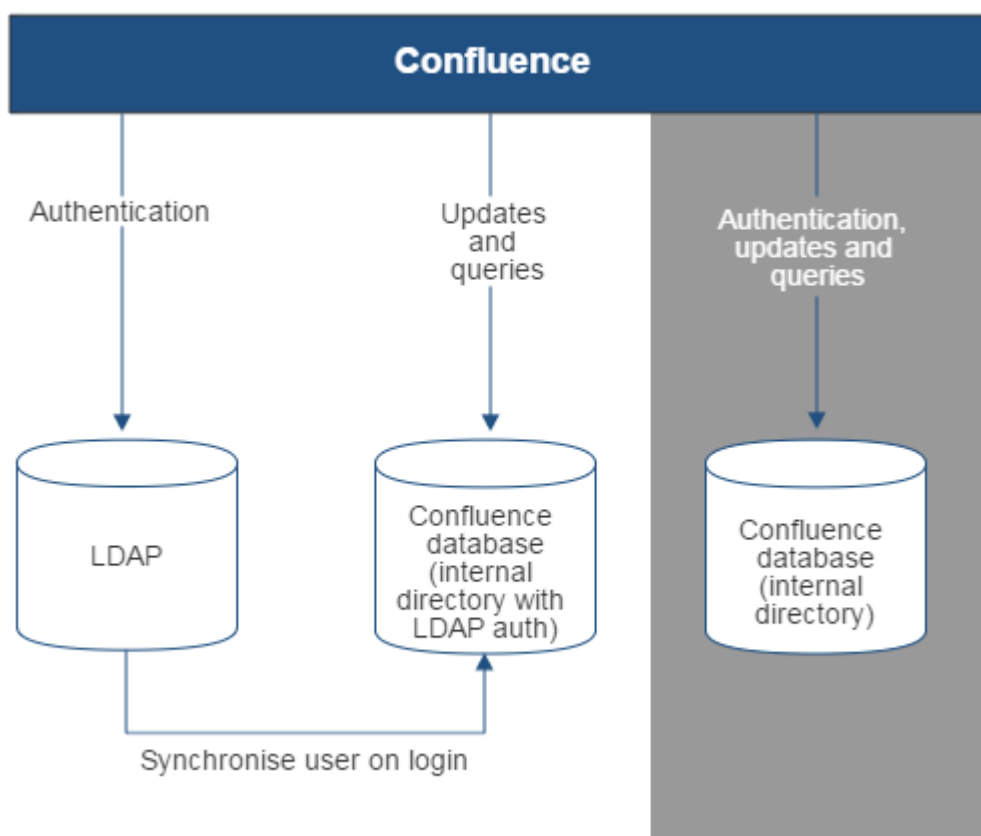


Diagram above: Confluence connecting to an LDAP directory for authentication only, with each user synchronised with the internal directory that is using LDAP authentication when they log in to Confluence.

Confluence Connecting to JIRA

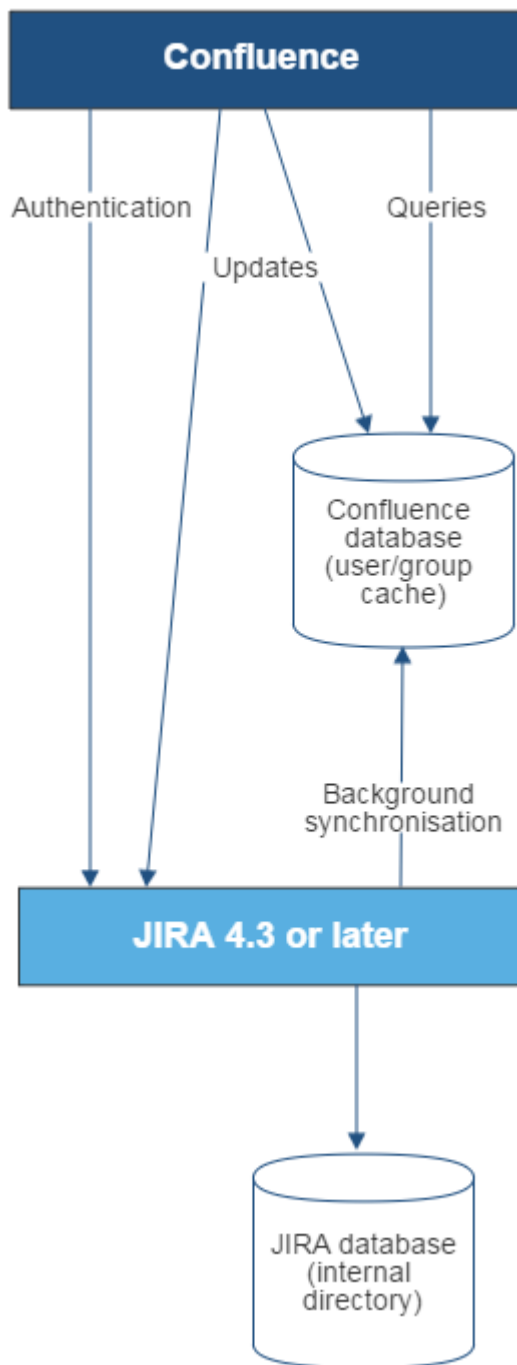


Diagram above: Confluence connecting to JIRA for user management.

Confluence Connecting to JIRA and JIRA Connecting to LDAP

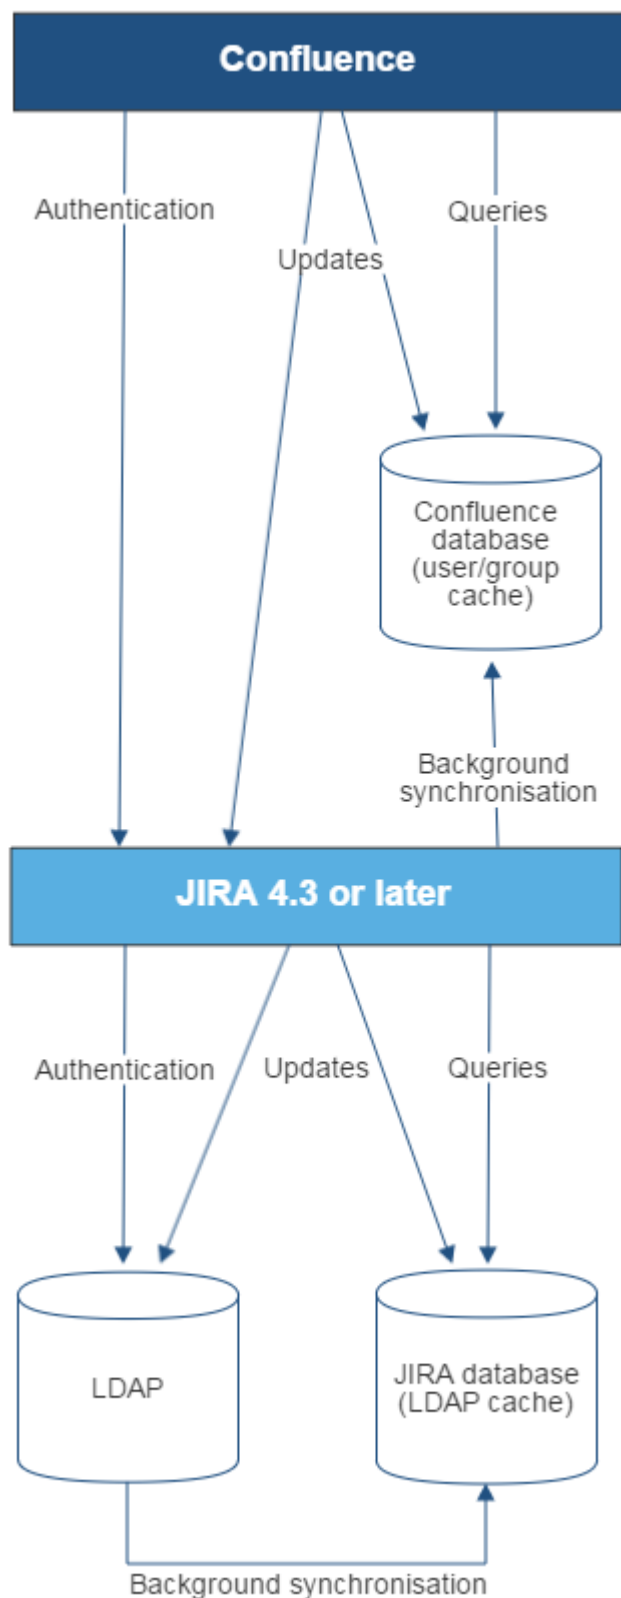


Diagram above: Confluence connecting to JIRA for user management, with JIRA in turn connecting to LDAP.

Confluence and JIRA Connecting to Crowd

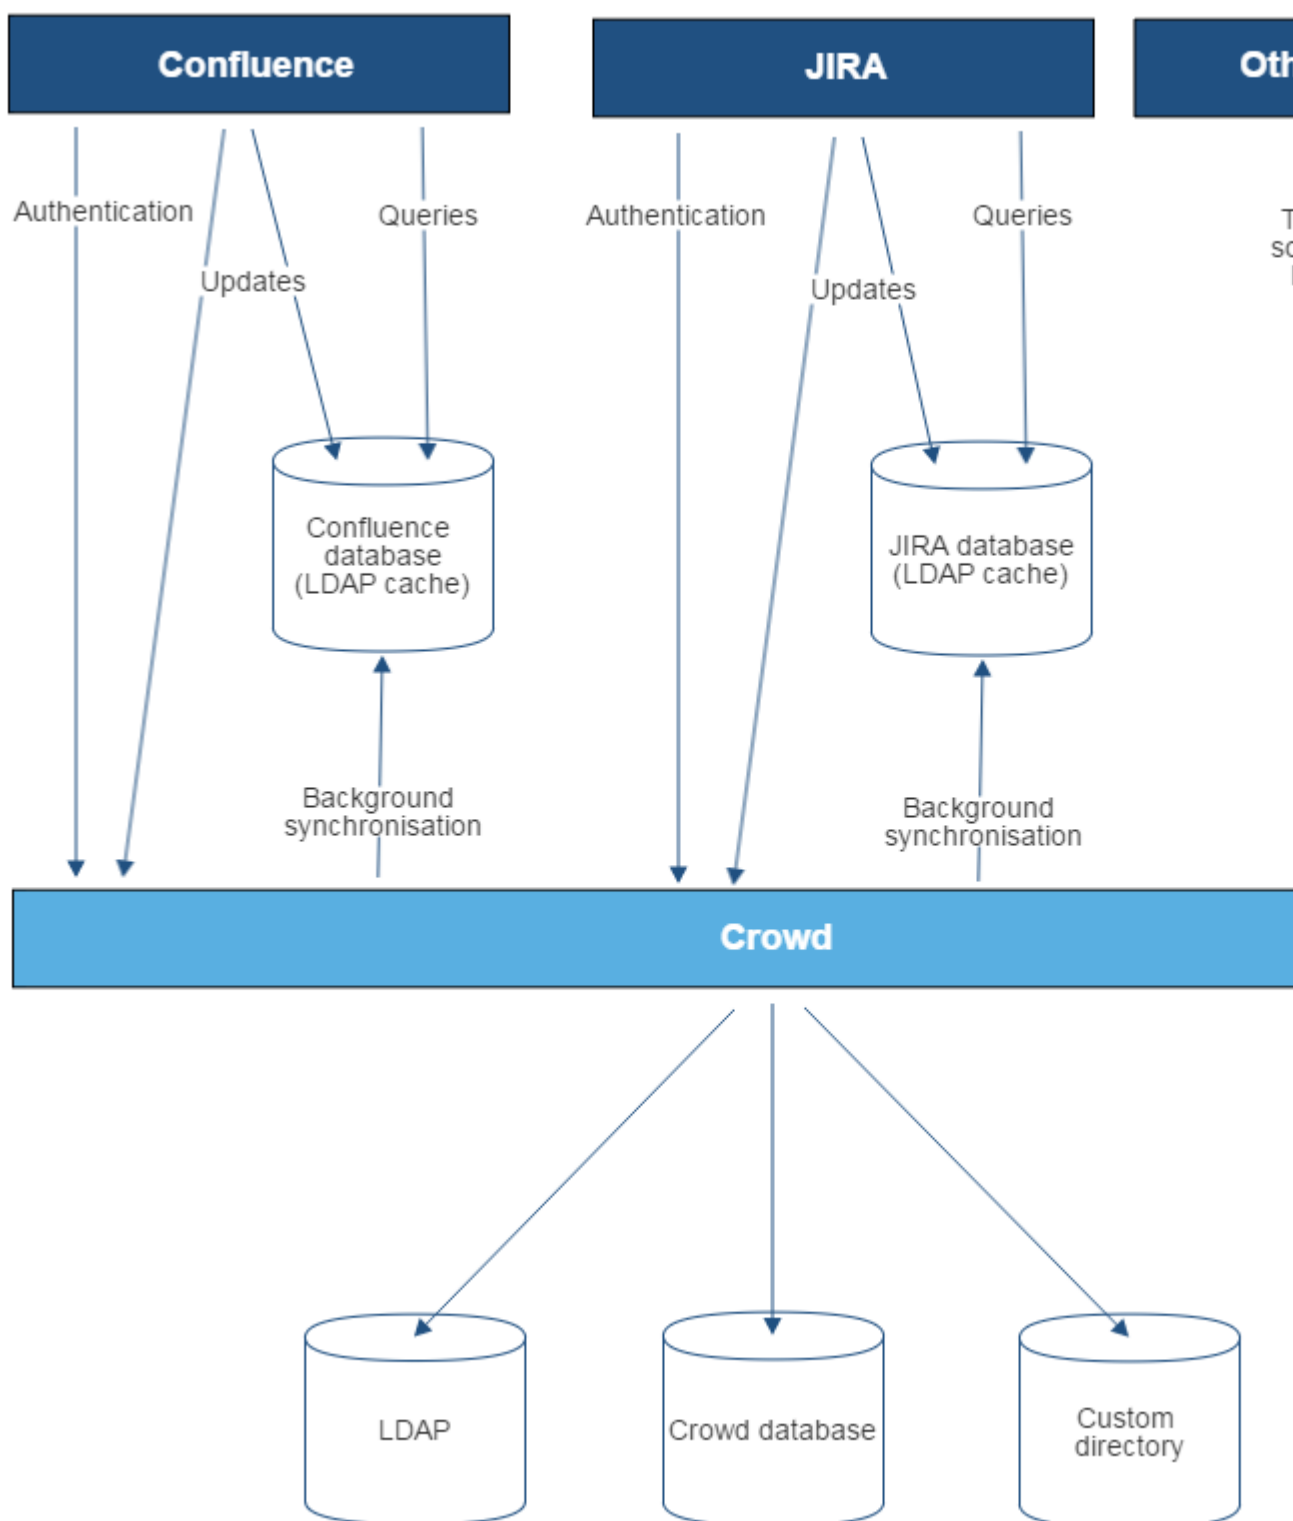


Diagram above: Confluence, JIRA and other applications connecting to Crowd for user management.

User Management Limitations and Recommendations

This page describes the optimal configurations and limitations that apply to user management in Confluence.

General Recommendations

- **Avoid duplicate usernames across directories.** If you are connecting to more than one user directory, we recommend that you ensure the usernames are unique to one directory. For example, we do not recommend that you have a user `jsmith` in both 'Directory1' and 'Directory2'. The reason is the potential for confusion, especially if you swap the order of the directories. Changing the directory order can change the user that a given username refers to.
- **Be careful when deleting users in remote directories.** If you are connecting to an LDAP directory, a Crowd directory or a JIRA directory, please take care when deleting users from the remote directory. If you delete a user that is associated with data in Confluence, this will cause problems in Confluence.

On this page:

- [General Recommendations](#)
- [Recommendations for Connecting to LDAP](#)
 - [Optimal Number of Users and Groups in your LDAP Directory](#)
 - [Redundant LDAP is Not Supported](#)
 - [Specific Notes for Connecting to Active Directory](#)
- [Recommendations for Connecting to JIRA for User Management](#)
 - [Single Sign-On Across Multiple Applications is Not Supported](#)
 - [Custom Application Connectors are Not Supported](#)
 - [Custom Directories are Not Supported](#)
 - [Load on your JIRA instance](#)
 - [JIRA Cloud applications not supported](#)
 - [Recommendations](#)

Related pages:

- [Connecting to an LDAP Directory](#)
- [Connecting to Crowd or JIRA for User Management](#)
- [Configuring User Directories](#)

Recommendations for Connecting to LDAP

Please consider the following limitations and recommendations when connecting to an LDAP user directory.

Optimal Number of Users and Groups in your LDAP Directory

The connection to your LDAP directory provides powerful and flexible support for connecting to, configuring and managing LDAP directory servers. To achieve optimal performance, a background synchronisation task loads the required users and groups from the LDAP server into the application's database, and periodically fetches updates from the LDAP server to keep the data in step. The amount of time needed to copy the users and groups rises with the number of users, groups, and group memberships. For that reason, we recommended a maximum number of users and groups as described below.

This recommendation affects connections to LDAP directories:

- Microsoft Active Directory
- All other LDAP directory servers

The following LDAP configurations are **not** affected:

- Internal directories with LDAP authentication
- LDAP directories configured for 'Authentication Only, Copy User On First Login'

Please choose one of the following solutions, depending on the number of users, groups and memberships in your LDAP directory.

Your environment	Recommendation
------------------	----------------

Up to 10 000 (ten thousand) users, 1000 (one thousand) groups, and 20 (twenty) groups per user	Choose the ' LDAP ' or ' Microsoft Active Directory ' directory type. You can make use of the full synchronisation option. Your application's database will contain all the users and groups that are in your LDAP server.
More than the above	Use LDAP filters to reduce the number of users and groups visible to the synchronisation task.

Our Test Results

We performed internal testing of synchronisation with an AD server on our local network consisting of 10 000 users, 1000 groups and 200 000 memberships.

We found that the initial synchronisation took about 5 minutes. Subsequent synchronisations with 100 modifications on the AD server took a couple of seconds to complete.

Please keep in mind that a number of factors come into play when trying to tune the performance of the synchronisation process, including:

- **Size of userbase.** Use LDAP filters to keep this to the minimum that suits your requirements.
- **Type of LDAP server.** We currently support change detection in AD, so subsequent synchronisations are much faster for AD than for other LDAP servers.
- **Network topology.** The further away your LDAP server is from your application server, the more latent LDAP queries will be.
- **Database performance.** As the synchronisation process caches data in the database, the performance of your database will affect the performance of the synchronisation.
- **JVM heap size.** If your heap size is too small for your userbase, you may experience heavy garbage collection during the synchronisation process which could in turn slow down the synchronisation.

Redundant LDAP is Not Supported

The LDAP connections do not support the configuration of two or more LDAP servers for redundancy (automated failover if one of the servers goes down).

Specific Notes for Connecting to Active Directory

When the application synchronises with Active Directory (AD), the synchronisation task requests only the changes from the LDAP server rather than the entire user base. This optimises the synchronisation process and gives much faster performance on the second and subsequent requests.

On the other hand, this synchronisation method results in a few limitations:

1. **Externally moving objects out of scope or renaming objects causes problems in AD.** If you move objects out of scope in AD, this will result in an inconsistent cache. We recommend that you do not use the external LDAP directory interface to move objects out of the scope of the sub-tree, as defined on the application's directory configuration screen. If you do need to make structural changes to your LDAP directory, manually synchronise the directory cache after you have made the changes to ensure cache consistency.
2. **Synchronising between AD servers is not supported.** Microsoft Active Directory does not replicate the uSNChanged attribute across instances. For that reason, we do not support connecting to different AD servers for synchronisation. (You can of course define multiple different directories, each pointing to its own respective AD server.)
3. **Synchronising with AD servers behind a load balancer is not supported.** As with synchronising between two different AD servers, Microsoft Active Directory does not replicate the uSNChanged attribute across instances. For that reason, we do not support connecting to different AD servers even when they are load balanced. You will need to select one server (preferably one that is local) to synchronise with instead of using the load balancer.
4. **You must restart the application after restoring AD from backup.** On restoring from backup of an AD server, the uSNChanged timestamps are reverted to the backup time. To avoid the resulting confusion, you will need to flush the directory cache after a Active Directory restore operation.
5. **Obtaining AD object deletions requires administrator access.** Active Directory stores deleted objects in a special container called cn=Deleted Objects. By default, to access this container you need

to connect as an administrator and so, for the synchronisation task to be aware of deletions, you must use administrator credentials. Alternatively, it is possible to change the permissions on the cn=Deleted Objects container. If you wish to do so, please see [this Microsoft KB Article](#).

6. **The User DN used to connect to AD must be able to see the uSNChanged attribute.** The synchronisation task relies on the uSNChanged attribute to detect changes, and so must be in the appropriate AD security groups to see this attribute for all LDAP objects in the subtree.

Recommendations for Connecting to JIRA for User Management

Please consider the following limitations and recommendations when connecting to a JIRA server for user management.

Single Sign-On Across Multiple Applications is Not Supported

When you connect to a JIRA application for user management, you will not have single sign-on across the applications connected in this way. JIRA, when acting as a directory manager, does not support SSO.

Custom Application Connectors are Not Supported

JIRA applications, Confluence, FishEye, Crucible and Bamboo can connect to a JIRA server for user management. Custom application connectors will need to use the new REST API.

Custom Directories are Not Supported

Earlier versions of JIRA supported OSUser Providers. It was therefore possible write a special provider to obtain user information from any external user directory. This is no longer the case.

Load on your JIRA instance

If your JIRA instance is already under high load, then using it as a User Server will increase that load.

JIRA Cloud applications not supported

You cannot use JIRA Cloud applications to manage standalone users. Cloud users and users within your self-hosted Atlassian applications need to be managed separately.

Recommendations

Your environment	Recommendation
<p>If all the following are true:</p> <ul style="list-style-type: none"> • Your JIRA application is not under high load. • You want to share user and group management across just a few applications, such as one JIRA Software server and one Confluence server, or two JIRA servers. • You do not need single sign-on (SSO) between your JIRA application and Confluence, or between two JIRA servers. • You do not have custom application connectors. Or, if you do have them, you are happy to convert them to use the new REST API. • You are happy to shut down all your servers when you need to upgrade your JIRA application. 	<p>Your environment meets the optimal requirements for using a JIRA application for user management.</p>

<p>If one or more of the following are true:</p> <ul style="list-style-type: none"> • If your JIRA application is already under high load. • You want to share user and group management across more than 5 applications. • You need single sign-on (SSO) across multiple applications. • You have custom applications integrated via the Crowd SOAP API, and you cannot convert them to use the new REST API. • You are not happy to shut down all your servers when you need to upgrade JIRA. 	<p>We recommend that you install Atlassian Crowd for user management and SSO.</p>
<p>If you are considering creating a custom directory connector to define your own storage for users and groups...</p>	<p>Please see if one of the following solutions will work for you:</p> <ul style="list-style-type: none"> • If you have written a custom provider to support a specific LDAP schema, please check the supported LDAP schemas to see if you can use one of them instead. • If you have written a custom provider to support nested groups, please consider enabling nested groups in the supported directory connectors instead. • If you have written a custom provider to connect to your own database, please consider loading the data into the application's database instead. • If you need to keep the custom directory connection, please consider whether Atlassian Crowd meets your requirements. See the documentation on Creating a Custom Directory Connector.

Requesting Support for External User Management

This page gives guidelines on how to request help from the Atlassian support team if you are having problems with external user management. External user management includes connections to Active Directory, other LDAP servers, Atlassian Crowd or a JIRA application for user management. The information on this page is provided in addition to the more general page on [Troubleshooting Problems and Requesting Technical Support](#).

The cause of such problems may be:

- The LDAP server is not responding.
- The application password is incorrectly configured, causing the LDAP server or other directory to return an authentication error.
- Other LDAP settings are incorrectly configured.

Troubleshooting the Connection to your External User Directory

The configuration screen for external directories in Confluence has a '**Test Settings**' button. This will help you to diagnose problems with user management in Active Directory and other LDAP servers.

To test your directory connection:

1. Choose the **cog icon**

On this page:

- [Troubleshooting the Connection to your External User Directory](#)
- [Problems During Initial Setup](#)
- [Complex Authentication or Performance Problems](#)

Related pages:

- [Troubleshooting Problems and Requesting Technical Support](#)
- [Configuring User Directories](#)



- , then choose **General Configuration** under Confluence Administration
2. Click '**User Directories**' in the left-hand panel.
3. **Edit** the relevant directory.
4. Click '**Test Settings**'.
5. The results of the test will appear at the top of the screen.

Please refer to our knowledge base articles for [troubleshooting user management and login issues](#).

If the above resources do not help, continue below.

Problems During Initial Setup

Raise a [support request](#) and include the following information.

- Download an LDAP browser to make sure you have the right settings in your LDAP directory. Atlassian recommends [LDAP Studio](#). Include screenshots of your user and group DNs.
- If you can start up Confluence and access the Administration Console, review your directory settings. See [Connecting to an LDAP Directory](#). Attach screenshots of all your settings.

Complex Authentication or Performance Problems

Raise a [support request](#) and include the following information.

Confluence Server

Log in to Confluence and access the Administration Console.

- Take a screenshot of the '**System Information**' screen, or save the page as HTML.
- Take a screenshot of the '**Global Permissions**' screen, if people are having problems with logging in.
- Go to '**Space Admin**' for the relevant space and take a screenshot of the '**Permissions**' page, if you are having problems with space or page permissions.

Confluence Configuration Files

- If you have implemented a custom authenticator or in any way modified `seraph-config.xml` or `seraph-paths.xml`, please provide the modified file.

User Management System

- Include the name and version of your LDAP server.
- Does your LDAP server use dynamic or static groups?
- Review your directory settings. See [Connecting to an LDAP Directory](#). Attach screenshots of all your settings.

Diagnostics

- Enable profiling. See [Performance Tuning](#).
- Enable detailed user management logging, by editing `confluence/WEB-INF/classes/log4j.properties`.

Change this section:

```
###
# Atlassian User
###
#log4j.logger.com.atlassian.user=DEBUG
#log4j.logger.com.atlassian.confluence.user=DEBUG
#log4j.logger.bucket.user=DEBUG
#log4j.logger.com.atlassian.seraph=DEBUG
#log4j.logger.com.opensymphony.user=DEBUG
```

Remove the '#' signs at the beginning of the lines, so that it looks like this:

```
###
# Atlassian User
###
log4j.logger.com.atlassian.user=DEBUG
log4j.logger.com.atlassian.confluence.user=DEBUG
log4j.logger.bucket.user=DEBUG
log4j.logger.com.atlassian.seraph=DEBUG
log4j.logger.com.opensymphony.user=DEBUG
```

- After enabling both the above, please attempt a Confluence LDAP account login and attach a copy of the log files that are produced when the problem occurs. To do this, locate your install directory, then zip the full `/logs` directory into a single file for us to examine. The logs directory is located in your Confluence Home directory.

Disabling the Built-In User Management

In some circumstances you may want to disable Confluence's built in user management, and delegate all user management to an external application, such as JIRA Software or JIRA Service Desk. You can disable internal user management by turning on Confluence's **External User Management** setting. You'll need to be a [system administrator](#) to do this.

You might disable Confluence's internal user management:

- When [Crowd's directory permissions](#) are configured so that Confluence cannot update the Crowd directories (as a system error will occur when Confluence attempts to write data into Crowd). See [Connecting to Crowd or JIRA for User Management](#) for more information.
- If you are using a JIRA application for user management. This centralises all user management in that JIRA app. See [Connecting to Crowd or JIRA for User Management](#).

To disable management of users and groups within Confluence:



1. > **General Configuration** > **Security Configuration**.
2. Click **Edit**.
3. Select the **External user management** checkbox then **Save** your change.

Note: If you turn on **External user management**:

- You will not be able to add users or groups in Confluence.
- You will not be able to use public signup in your site.
- The **Forgot Password** link will not appear on the Confluence login page.
- Users will not be able to reset their password in Confluence.

Managing Add-ons or Plugins

An **add-on** is a separately installed component that provides Confluence functionality. The terms 'plugin' and 'add-on' are often used interchangeably.

There are two main types of add-ons:

- System add-ons - these are bundled with Confluence and provide core functionality
- User installed add-ons - these are usually downloaded from [The Marketplace](#) and may have been created by Atlassian or by a third party developer.

For information about developing your own add-ons for Confluence, see the [Confluence Developer documentation](#).

About the Universal Plugin Manager

Add-ons are managed via the Universal Plugin Manager (known as the UPM). The UPM can be found in most Atlassian applications, and provides a consistent experience for administering add-ons. To visit the UPM, go to



> **Add-ons** in the Confluence header.

The UPM allows you to:

- Discover and install new add-ons from the [Atlassian Marketplace](#).
- Install or remove add-ons.
- Configure add-on settings.
- Enable or disable add-ons and their component modules.
- Confirm add-on compatibility before upgrading Confluence.

You'll need Confluence Administrator permissions to access the UPM. See [Request Add-ons](#) for information on how users can find and request add-ons.

See the [Universal Plugin Manager documentation](#) for more information on using the UPM.

Writing User Macros

User macros are useful if you want to create your own custom macros. These can be to perform specific actions, apply custom formatting and much more.

User macros are created and managed within Confluence itself, you do not need to develop an add-on. You will need some coding skills though.

You'll need System Administrator permissions to create and manage user macros.

On this page:


- [Create a User Macro](#)
- [Edit a user macro](#)
- [Delete a user macro](#)
- [Best practices](#)
- [Example user macros](#)

Related Pages:

- [User Macro Module](#) (Developer documentation)

Create a User Macro

To add a new user macro:

1. Go to  > **General Configuration** > **User Macros**
2. Choose **Create a User Macro**
3. Enter the macro details (see table below)
4. Click **Add**


Macro details field	Description
Macro name	This is the name of the macro, as it appears in the code.
Visibility	<p>This controls who can see this macro in the macro browser or auto-complete. Options are:</p> <ul style="list-style-type: none"> • Visible to all users • Visible only to system administrators <p>Note that if you select Visible only to system administrators, users will still see the output of the macro on a page, and the macro placeholder will still be visible when a user edits a page. It is only hidden in the macro browser and autocomplete.</p> <p>All macro information is discoverable, including the macro title, description, parameter names and other metadata. Do not include confidential data anywhere in the definition of a user macro, even if it is marked as visible only to system administrators.</p>
Macro Title	This is the title that will appear in the macro browser and auto-complete.
Description	This is the description that will appear in the macro browser. The macro browser's search will pick up matches in both the title and description.

Categories	Select one or more macro browser categories for your macro to appear in.
Icon URL	Enter an absolute URL (for example <code>http://mysite.com/mypath/status.png</code>) or path relative to the Confluence base URL (for example <code>/images/icons/macrobrowser/status.png</code>) if you want the macro browser to display an icon for your macro.
Documentation URL	If you have documentation for your macro, enter the URL here.
Macro Body Processing	<p>Specify how Confluence should process the body before passing it to your macro.</p> <p>The macro body is the content that is displayed on a Confluence page. If your macro has a body, any body content that the user enters will be available to the macro in the <code>\$body</code> variable.</p> <p>Options for processing the macro body include:</p> <ul style="list-style-type: none"> • No macro body Select this option if your macro does not have a body. • Escaped Confluence will add escape characters to the HTML markup in the macro body. Use this if you want to show actual HTML markup in the rendered page. For example, if the body is <code>Hello World</code> it will render as <code>Hello World</code>. • Unrendered HTML in the body will be processed within the template before being output. Ensure that HTML is ultimately output by the template. • Rendered Confluence will recognise HTML in the macro body, and render it appropriately. For example, if the body is <code>Hello World</code> it will render as Hello World.
Template	<p>This is where you write the code that determines what the macro should do.</p> <ul style="list-style-type: none"> • Use HTML and Confluence-specific XML elements in the macro template. Details of Confluence's storage format are in Confluence Storage Format. • You can use the Velocity templating language. Here is more information on the Velocity project. • If your macro has a body, your template can refer to the macro body text by specifying <code>'\$body'</code>. • Each parameter variable you use must have a matching metadata definition. Use <code>@param</code> to define metadata for your macro parameters. • When using the information passed using parameters, refer to your parameters as <code>\$paramXXX</code> where 'XXX' is the parameter name that you specified in the <code>@param</code> metadata definition. • Use <code>@noparams</code> if your macro does not accept parameters. <p>See User Macro Template Syntax for more information and examples.</p>

Do you need a plugin instead?

If you want to distribute your user macro as a plugin, please refer to the developer's guide to the [User Macro plugin module](#). If you want to create more complex, programmatic macros in Confluence, you may need to write a [Macro plugin](#).


Edit a user macro**To edit a user macro:**

1. Go to  **> General Configuration > User Macros**
2. Click **Edit** next to the relevant macro
3. Update the macro details

4. Click **Save**

Delete a user macro

To delete a user macro:

1. Go to  **> General Configuration > User Macros**
2. The currently configured user macros will appear
3. Click **Delete** next to the relevant macro

Before deleting a user macro, you should [search](#) for all occurrences of the macro in pages and blog posts. Users will see an 'unknown macro' error if you delete a user macro that is still in use on a page.

Best practices

This section contains tips and suggestions for best practices when creating your own user macros.

Add a descriptive header to your macro template

We recommend that you include a short description as a comment at the top of the **Template** field as shown below.

```
## Macro title: My macro name
## Macro has a body: Y or N
## Body processing: Selected body processing option
## Output: Selected output option
##
## Developed by: My Name
## Date created: dd/mm/yyyy
## Installed by: My Name

## Short description of what the macro does
```

Expose your parameters in the macro browser

The macro browser is the easiest way for users to configure your macro. You can specify the macro category, link to an icon, define the parameters that the macro browser will use to prompt the user for information, and more.

Supply default values for macro parameters

As you can't guarantee that a user has supplied parameters, one of the first things to do in the macro is check that you have received some value if you expect to rely on it later on in the macro code.

In the example below, the macro expects three parameters, and substitutes sensible defaults if they are not supplied.

```
#set($spacekey= $paramspacekey)
#set($numthreads= $paramnumthreads)
#set($numchars= $paramnumchars)

## Check for valid space key, otherwise use current
#if (!$spacekey)
  #set ($spacekey=$space.key)
#end

## Check for valid number of threads, otherwise use default of 5
#if (!$numthreads)
  #set ($numthreads=5)
#end

## Check for valid excerpt size, otherwise use default of 35
#if (!$numchars)
  #set ($numchars=35)
#end
```

Example user macros

▼ [Hello World](#)

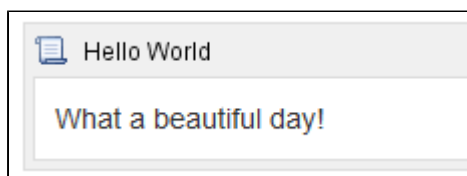
This example demonstrates how to create a user macro that displays the text 'Hello World!' and any text that the user places in the body of the macro.

Field	Value
Macro name	helloworld
Visibility	Visible to all users in the Macro Browser
Macro Title	Hello World
Description	Displays "Hello World" and the macro body.
Categories	Confluence Content
Icon URL	You can leave this field blank
Documentation URL	You can leave this field blank
Macro body processing	Rendered

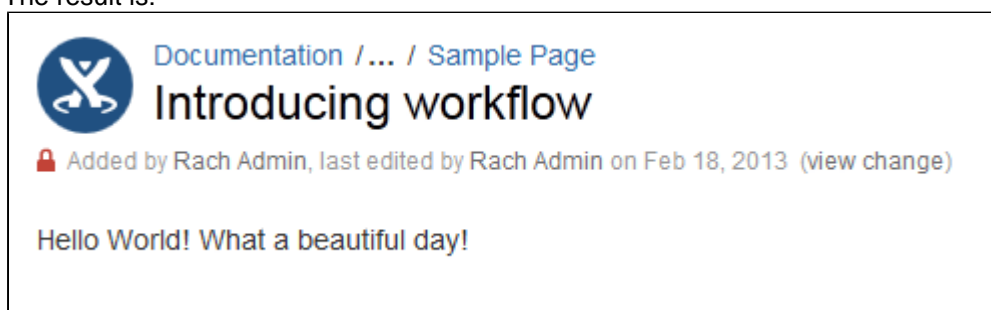
Template	<p>Enter the code below in the template field - this example will print the text straight onto the page.</p> <pre>## @noparams Hello World! \$body</pre> <p>If you wanted the text to appear in a panel you could include the relevant AUI message class as shown here.</p> <pre>## @noparams <div class="aui-message closeable"> Hello World! \$body </div></pre>
-----------------	--

Using the 'Hello World' macro on a page

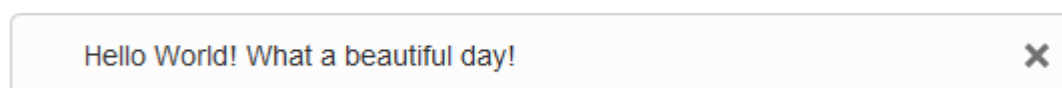
Now you can add the macro to your Confluence page using the Macro Browser, or by typing {hello in the editor and selecting the macro from the list of suggestions.



The result is:



If you chose to include a panel, the result would be:



▼ NoPrint

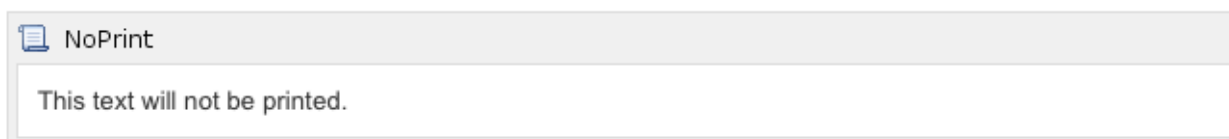
This example demonstrates how to create a user macro that can contain text that is visible when viewing a page, but does not print.

Field	Value
Macro name	noprint
Visibility	Visible to all users in the Macro Browser
Macro Title	No Print

Description	Hides text from printed output.
Categories	Confluence Content
Icon URL	You can leave this field blank
Documentation URL	You can leave this field blank
Macro body processing	Rendered
Template	Enter the code below in the template field. <pre>## @noparams <div class="NoPrint">\${body}</div></pre>

Using the 'NoPrint' Macro on a page

Now you can add the macro to your Confluence page using the Macro Browser. Text entered into the body of the macro placeholder will not be printed, but will appear when the page is viewed online.



Making the PDF export recognise the NoPrint macro

See [Advanced PDF Export Customisations](#).

▼ Colour and Size

This example demonstrates how you can pass parameters to your macro. We'll create a font style macro which has two parameters to allow the user to specify the colour and size of the text contained in the macro body.

Field	Value
Macro name	stylish
Visibility	Visible to all users in the Macro Browser
Macro Title	Stylish
Description	Applies colour and size to text.
Categories	Confluence Content
Icon URL	You can leave this field blank
Documentation URL	You can leave this field blank
Macro body processing	Rendered

Template	<p>Enter the code below in the template field. If your macro requires more than one parameter, you can use variables \$param0 to \$param9 to represent them.</p> <pre>## @param 0:title=colour type=string ## @param 1:title=size type=string \$body</pre> <p>Alternatively, you can also use explicitly-named parameters in your macro. These macro parameters will appear as variables with the name \$param<x> where <x> is the name of your parameter.</p> <pre>## @param Colour:title=colour type=string ## @param Size:title=size type=string \$body</pre>
-----------------	--

▼ Formatted Panel

This example demonstrates how to write a user macro that creates a panel that is preformatted with specific colours. It will create a panel that looks like this:

(Title)

Note: The panel's title will be empty if the user does not give a value for the title parameter.

Field	Value
Macro name	formpanel
Visibility	Visible to all users in the Macro Browser
Macro Title	Formatted Panel
Description	Creates a panel preformatted with specific colours
Categories	Formatting
Icon URL	You can leave this field blank
Documentation URL	You can leave this field blank
Macro body processing	Escaped

Template	<p>Enter the code below in the template field. See below for a more detailed explanation of the code below.</p> <pre>## @param Title:title=Title type=string desc=Title <ac:structured-macro ac:name="panel"> <ac:parameter ac:name="titleBGColor">#ccc</ac:parameter> <ac:parameter ac:name="borderStyle">solid</ac:parameter> <ac:parameter ac:name="borderColor">#6699CC</ac:parameter> <ac:parameter ac:name="borderWidth">2</ac:parameter> <ac:parameter ac:name="titleColor">#000000</ac:parameter> <ac:parameter ac:name="title">\${paramTitle}</ac:parameter> <ac:rich-text-body>\$body</ac:rich-text-body> </ac:structured-macro></pre>
-----------------	--

Explanation of the code in the macro template

Below is a breakdown of the user macro template code.

Item	Description
## @param Title:title=Title type=string desc=Title	<p>@param defines the metadata for your macro parameters.</p> <p>@param Title</p> <p>This parameter is called "Title".</p> <p>title=Title</p> <p>defines the parameter title that will appear in the macro browser as "Title".</p> <p>type=string</p> <p>defines the field type for the parameter as a text field.</p> <p>desc=Title</p> <p>defines the description of the parameter in the macro browser.</p>

<pre><ac:structured-macro ac:name="panel"></pre>	<p>This calls the Confluence Panel macro.</p> <p>The easiest way to find out the code name of a Confluence macro by viewing the Storage Format of a page containing the macro. You'll need Confluence Administrator permissions to view the storage format.</p>
<pre><ac:parameter ac:name="titleBGColor">#ccc</ac:parameter> <ac:parameter ac:name="borderStyle">solid</ac:parameter> <ac:parameter ac:name="borderColor">#6699CC</ac:parameter> <ac:parameter ac:name="borderWidth">2</ac:parameter> <ac:parameter ac:name="titleColor">#000000</ac:parameter></pre>	<p>Sets the parameters for the macro: the background colour, border style, border colour, border width and title colour.</p> <p>To discover the names of the parameters for a Confluence macro, view the storage format as described above.</p>
<pre><ac:parameter ac:name="title">\${!paramTitle}</ac:parameter></pre>	<p>Enters the value stored in the 'Title' parameter into the title section of the macro.</p> <p>The ! tells the macro to leave the title blank, when there is no data in the "Title" parameter.</p>
<pre><ac:rich-text-body>\$body</ac:rich-text-body></pre>	<p>Users can enter data that is stored in the body of the macro. This line enables the macro to access and store the body content passed to your macro.</p>
<pre></ac:structured-macro></pre>	<p>This command marks the end of the macro.</p>

User Macro Template Syntax

See [Writing User Macros](#) for an introduction to writing a user macro.

This page provides information about the code you can enter in a user macro template.

Accessing your macro's body

Use the `$body` object within your user macro template to access the content passed to your macro in the macro body.

The `$body` object is available if you have specified that your macro has a body (in other words, if you have *not* selected **No macro body**).

Example: Let's assume your macro is called `hello world`.

Enter the following code in your template:

```
Hello World: $body
```

A user, when editing a Confluence page, chooses your macro in the macro browser and then enters the following in the macro placeholder that is displayed in the edit view:

```
From Matthew
```

The wiki page will display the following:

```
| Hello World: From Matthew
```

Using parameters in your user macro

You can specify parameters for your macro, so that users can pass it information to determine its behaviour on a Confluence page.

How your macro parameters are used on a Confluence page

When adding a macro to a Confluence page, the macro browser will display an input field for each macro parameter. The field type is determined by the parameter type you specify.

Defining the parameters

A parameter definition in the template contains:

- `@param`
- The parameter name
- A number of attributes (optional).

Format:

```
## @param MYNAME:title=MY TITLE|type=MY TYPE|desc=MY  
DESCRIPTION|required=true|multiple=true|default=MY DEFAULT VALUE
```

Additional notes:

- The order of the parameters in the template determines the order in which the macro browser displays the parameters.
- We recommend that you define the parameters at the top of the template.
- There may be additional attributes, depending on the parameter type you specify.

On this page:

- [Accessing your macro's body](#)
- [Using parameters in your user macro](#)
- [Objects available to your macro](#)
- [Controlling parameter appearance in the editor placeholder](#)

Related pages:

- [Writing User Macros](#)

The sections below describe each of the attributes in detail.

Attribute name	Description	Required / Recommended / Optional
(an unnamed, first attribute)	A unique name for the parameter. The parameter name is the first attribute in the list. The name attribute itself does not have a name. See the section on name below.	Required
title	The parameter title will appear in the macro browser. If you do not specify a title, Confluence will use the parameter name.	Recommended
type	The field type for the parameter. See the section on type below.	Recommended
desc	The parameter description will appear in the macro browser.	Optional
required	Specifies whether the user must enter information for this parameter. Defaults to false.	Optional
multiple	Specifies whether the parameter accepts multiple values. Defaults to false.	Optional
default	The default value for the parameter.	Optional

Parameter name

The parameter name is the first attribute in the list. The name attribute itself does not have a name.

Example: The following code defines 2 parameters, named 'foo' and 'bar':

```
## @param foo
## @param bar
```

Parameter type

The field type for the parameter. If you do not specify a type, the default is `string`.

Parameter type	Description
boolean	Displays a checkbox to the user and passes the value 'true' or 'false' to the macro as a string.
enum	Offers a list of values for selection. You can specify the values to appear in a dropdown in the macro browser. Example of specifying the enum values: <div data-bbox="494 1655 1482 1722" data-label="Text"> <pre>## @param colour:title=Colour type=enum enumValues=Grey,Red,Yellow,Green</pre> </div> <p><i>Note about i18n:</i> Confluence does not support internationalisation of the enum values. The user sees is the one passed to the macro as the parameter value, with the capitalisation case 'Grey', 'Red', etc.</p>

string	<p>A text field. This is the default type. Example with a required field:</p> <pre>## @param status:title=Status type=string required=true desc=Status to display</pre>
confluence-content	<p>Offers a control allowing the user to search for a page or blog post. Example:</p> <pre>## @param page:title=Page type=confluence-content required=true desc=S a page do use</pre>
username	<p>Search for user.</p> <pre>## @param user:title=Username type=username desc=Select user to display</pre>
spacekey	<p>Offers a list of spaces for selection. Passes the space key to the macro. Example:</p> <pre>## @param space:title=Space type=spacekey</pre>
date	<p>Confluence accepts this type, but currently treats it in the same way as 'string'. Example:</p> <pre>## @param fromDate:title=From Date type=date desc=Date to st from. Format: dd/mm/YYYY</pre> <p><i>Note about dates:</i> A user can enter a date in any format, you should validate the date for user macro.</p>
int	<p>Confluence accepts this type, but treats it in the same way as 'string'. Example with a default value:</p> <pre>## @param numPosts:title=Number of Posts type=int default=15 desc=Number of posts to display</pre>
percentage	<p>Confluence accepts this type, but treats it in the same way as 'string'. Example:</p> <pre>## @param pcent:title=Percentage type=percentage desc=Number posts to display</pre>

Using the parameters in your macro code

The parameters are available in your template as `$paramfoo`, `$parambar` for parameters named "foo" and "bar".

Normally, a parameter like `$paramfoo` that is missing will appear as '`$paramfoo`' in the output. To display nothing when a parameter is not set, use an exclamation mark after the dollar sign like this: `$!paramfoo`

Using no parameters

If your macro does not accept parameters, you should use `@noparams` in your template.

If the user macro contains no parameters and does not specify `@noparams`, then the macro browser will display a free-format text box allowing users to enter undefined parameters. This can be confusing if the macro does not accept parameters.

Example: Add the following line at the top of your template:

```
## @noparams
```

Objects available to your macro

Including the macro body and parameters, the following Confluence objects are available to the macro:

Variable	Description	Class Reference
<code>\$body</code>	The body of the macro (if the macro has a body)	String
<code>\$paramfoo</code> , <code>\$parambar</code> , ... <code>\$param<name></code>	Named parameters ("foo", "bar") passed to your macro.	String
<code>\$config</code>	The <code>BootstrapManager</code> object, useful for retrieving Confluence properties.	BootstrapManager
<code>\$renderContext</code>	The <code>PageContext</code> object, useful for (among other things) checking <code>\$renderContext.outputType</code>	PageContext
<code>\$space</code>	The <code>Space</code> object that this content object (page, blog post, etc) is located in (if relevant).	Space
<code>\$content</code>	The current <code>ContentEntity</code> object that this macro is a included in (if available).	ContentEntityObject

Macros can also access objects available in the default Velocity context, [as described in the developer documentation](#).

Controlling parameter appearance in the editor placeholder

You can determine which macro parameters should appear in the placeholder in the Confluence editor.

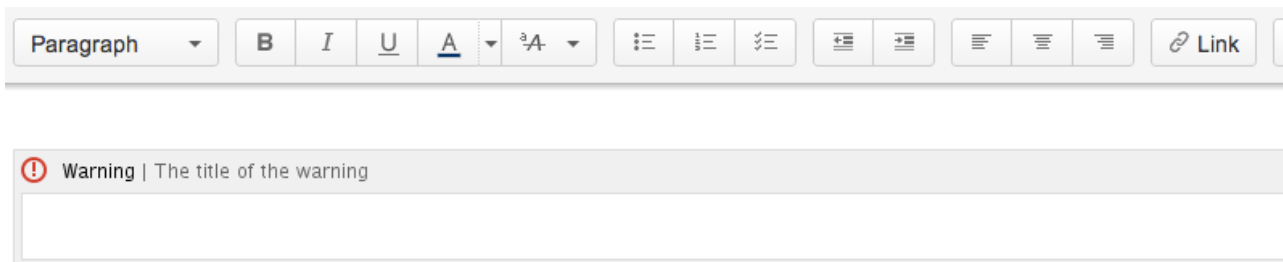
By default as many parameters as can fit will be displayed in the placeholder, as shown here:



You can control which parameters you want to display here, to ensure the most relevant information is visible to the author.

For example, the Confluence Warning macro has two parameters, *title* and *icon*. We consider *title* to be the most interesting parameter, so we have configured the Warning macro to show only the value of the *title* parameter.

Let's assume an author adds the Warning macro to a page, and gives it a title of 'The title of the warning'. The macro configuration leads to a placeholder as shown here:



To configure the macro placeholder for a user macro, you will add attributes to the @param entry in the template.

For example, if our Warning macro is a user macro, the configuration for the *title* parameter is as follows:

```
## @param
title:type=string|option-showNameInPlaceholder=false|option-showValue
InPlaceholder=true
```

The attribute `showNameInPlaceholder` specifies that the *title* parameter's *name* should not be shown.

The attribute `showValueInPlaceholder` specifies that the *title* parameter's *value* should be shown.

If none of the parameters in a macro include any of the above attributes, then the default behaviour is to show all the parameters that fit in the placeholder: full title and value.

If one or more parameters has either attribute set, then all parameters that do not include the attributes will default to false (that is, they will not be shown).

Customising your Confluence Site

This page is an introduction to customising Confluence at site level. This is of interest to Confluence administrators – people with System Administrator or Confluence Administrator permissions.

For guidelines on customisations at a personal and space level, see [Your User Profile](#) or [Change the Look and Feel of a Space](#).

We've documented the customisations under two broad headings:

- You can change the **appearance** of Confluence by customising the dashboard, adjusting the colours, adding a site logo, and more. See [Changing the Look and Feel of Confluence](#).
- You can determine the **default behaviour** by setting various options, or define the **default content** that appears in new spaces, on the dashboard, and in other Confluence locations. See [Changing the Default Behaviour and Content in Confluence](#).

Related pages:

- [Integrating Confluence with Other Applications](#)
- [Tracking Customisations Made to your Confluence Installation](#)
- [Confluence Administrator's Guide](#)

Changing the Look and Feel of Confluence

You can customise the 'look and feel' of Confluence at both the site (global) and space levels.

Any changes you make to the look and feel at site level will be applied as the default look and feel for all the spaces in the site. This means that any

customisations will only be reflected in the "Default" theme. No other theme will have an impact from this change. An individual space can be configured to have its own look and feel through the [space administration screens](#).

Related pages:

- [Administering Site Templates](#)
- [Confluence Administrator's Guide](#)
- [Working With Decorator Macros](#)
- [Customising a Specific Page](#)
- [Upgrading Customised Site and Space Layouts](#)

Ways to customise the look and feel of your site:

- Add your own site logo. See [Changing the Site Logo](#).
- Change the **colour scheme** of the user interface. See [Customising Colour Schemes](#).
- Use **themes** for advanced layout customisation. See [Working with Themes](#).
- Change the appearance of the **dashboard**. See [Customising the Confluence Dashboard](#).
- Change the **site or space layouts**, which determine how the controls are laid out in the site. This does not change the actual page layouts, but it does change the way the surrounding controls appear in the page. See [Customising Site and Space Layouts](#).

Customising the Confluence Dashboard

The dashboard is the default landing page for your Confluence site. It gives people all the tools they need to discover pages, resume their work and quickly jump to their favourite spaces and pages.

Editing the site welcome message

The site welcome message appears on the right hand side of the dashboard and is the perfect place to inject some of your organisation's personality.

See [Editing the Site Welcome Message](#) to find out how to add announcements, useful links, images, macros and more.

You'll need Confluence administrator permissions to edit the site welcome message.

Using a page as the site landing page

If you want more control, you can choose to use an ordinary Confluence page as your site landing page, instead of sending people to the dashboard. See [Configuring the Site Home Page](#) to find out more.

Using a page instead of the dashboard can be useful if most people will be reading, rather than creating, pages in your site. However, for sites where you want to encourage teams to collaborate, the dashboard provides the best tools for resuming work in progress and keeping up with what is happening in the site.

Advanced customisations

You can further customise the dashboard by editing the global layout file. See [Customising Site and Space Layouts](#) for more information on how to do this. You'll need some knowledge of [Velocity](#) to modify the layout files.

There are two locations that you can add content to:

- Web panels added to `atl.dashboard.secondary` will appear below the site welcome message.
- Web items added to `system.dashboard.button` will appear next to the Create space and Invite users button at the top right of the dashboard.

If you modify layouts in Confluence you will need to reapply your modifications each time you upgrade Confluence. The more dramatic your customisations are, the harder it may be to reapply the changes when upgrading. See [Upgrading Customised Site and Space Layouts](#) to find out what will be involved before

On this page:

- [Editing the site welcome message](#)
- [Using a page as the site landing page](#)
- [Advanced customisations](#)

Related pages:

- [Favourite Pages](#)
- [Changing the Look and Feel of Confluence](#)

modifying the layouts.

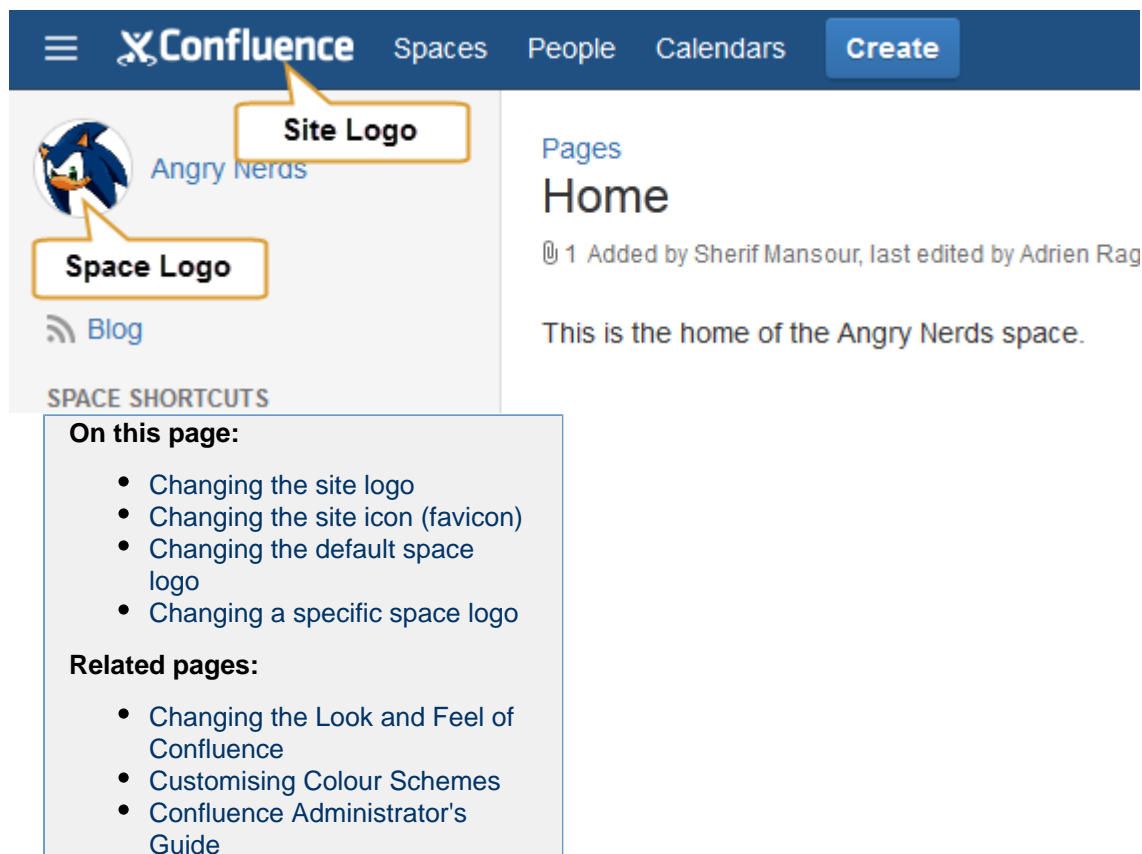
Changing the Site Logo

You can customise the look and feel of your Confluence site by changing the logos.

You can change:

- the **site logo**
- the **default space logo** for all spaces
- the **space logo** for individual spaces.

Screenshot: Location of the Site Logo and Space Logo in Confluence.



Changing the site logo

The Site Logo appears in the header and is visible throughout Confluence. You need Confluence Administrator permissions to change the site logo.

To change the site logo:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Site Logo and Favicon**.
3. Choose **Browse** to upload a new logo.
4. Choose **Show Logo Only** or **Show Logo and Title** depending on whether you wish the Site Title to display in the header.
5. Choose **Save**.

Confluence's Auto Look and Feel will detect the colours in your new logo, and change the site colour scheme to match.

If you would prefer to use the default [colour scheme](#) with your custom logo go to



> **General Configuration** > **Colour Scheme** > **Edit** and then choose **Reset** to revert back to the default

scheme.

Screenshot: Header showing Site Logo, Site Title and auto look and feel changes to the colour of the header



Changing the site icon (favicon)

You can also change the site favicon (the icon that appears in your browser tab). You need Confluence Administrator permissions to do this.

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Site Logo and Favicon**
3. Locate your image file and choose **Upload**.

You can upload PNG, GIF, JPEG, or ICO files. For best results images should be square, and at least 48x48 pixels.

Changing the default space logo

The Space Logo appears in the sidebar and as an icon in the Sites Directory. If you are using the Documentation theme the Space Logo displays beside the Space Title.

The default space logo applies to all spaces that do not have a custom space logo applied - see [Changing a Space's Logo](#).

You need to be a Confluence Administrator to change the default space logo.

To change the default space logo:

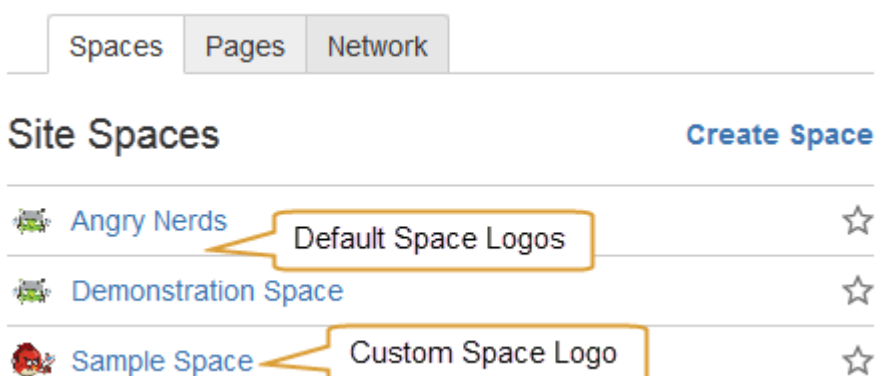
1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Default Space Logo** in the left-hand panel.
3. Choose **Logo:ON**
4. Choose **Browse** to upload a new logo
5. Choose **Upload Logo**
6. Choose **Save**.

Screenshot: Confluence spaces showing the default logo, and a space with a customised logo



Changing a specific space logo

Space Administrators can change the logo for their space. This overrides the default space logo and any changes to the default space logo will not appear in these spaces. See example above - 'Sample Space' has a custom logo.

See [Changing a Space's Logo](#) to find out how to change the logo in a specific space.

Customising Colour Schemes

Confluence administrators can configure a new colour scheme for the site. The default colour scheme for the site will also become the default for all spaces within it. Space administrators can configure a different colour scheme for spaces. The space colour scheme will override the site-wide colour scheme.

To change the site's colour scheme:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Colour Scheme** in the left-hand panel.
3. Choose **Edit**.
4. Enter standard HTML/CSS2 colour codes, or use the colour-picker



to choose a new colour from the palette provided.

5. Choose **Save**. Any changes you make will immediately be reflected across the Confluence site.

On this page

- [Reset your colour scheme after uploading a site logo](#)
- [Notes](#)

Related pages:

- [Page Templates](#)
- [Working with Themes](#)
- [Changing the Look and Feel of Confluence](#)
- [Confluence Administrator's Guide](#)

Some UI elements below are for specific themes, and colour changes may not take effect for other themes.

- **Top Bar** - the top navigation bar background
- **Top Bar Text** - the text on the top navigation bar
- **Header Button Background** - buttons on the top navigation bar (e.g. Create button)
- **Header Button Text** - the text on buttons on the top navigation bar
- **Top Bar Menu Selected Background** - background colour of top navigation bar menu items when selected (e.g. spaces)
- **Top Bar Menu Selected Text** - text colour of top navigation bar menu items when selected
- **Top Bar Menu Item Text** - text on top navigation bar drop down menus (e.g. help or cog)
- **Menu Item Selected Background** - highlight colour on top navigation bar drop down menu items
- **Menu Item Selected Text** - text colour on highlighted top navigation bar drop down menu items
- **Search Field Background** - the background colour of the search field on the header
- **Search Field Text** - the colour of the text in the search field on the header
- **Page Menu Selected Background** - the background colour of the drop down page menu when selected
- **Page Menu Item Text** - the text of the menu items in the drop down page menu
- **Heading Text** - all heading tags throughout the space
- **Space Name Text** - the text of the current space name located above the page title
- **Links** - all links throughout the space
- **Borders and Dividers** - table borders and dividing lines
- **Tab Navigation Background** - the background colour of the tab navigation
- **Tab Navigation Text** - the text of the tab navigation when highlighted
- **Tab Navigation Background Highlight** - the background colour of the tab navigation when highlighted
- **Tab Navigation Text Highlight** - the text of the tab navigation elements when highlighted

Screenshot: Editing the colour scheme

Custom Colour Scheme

A custom colour scheme which can be edited.



Edit

The following colours can be customised for this colour scheme.

Top Bar		<input type="text" value="#205081"/>	
Top Bar Text		<input type="text" value="#FFFFFF"/>	
Header Button Background		<input type="text" value="#3572B0"/>	
Header Button Text		<input type="text" value="#FFFFFF"/>	
Top Bar Menu Selected Background		<input type="text" value="#3B73AF"/>	
Top Bar Menu Selected Text		<input type="text" value="#FFFFFF"/>	
Top Bar Menu Item Text		<input type="text" value="#333333"/>	
Menu Item Selected Background		<input type="text" value="#3B73AF"/>	
Menu Item Selected Text		<input type="text" value="#FFFFFF"/>	
Search Field Background		<input type="text" value="RGBA(0, 0, 0, 0.2)"/>	
Search Field Text		<input type="text" value="#FFFFFF"/>	
Page Menu Selected Background		<input type="text" value="#3B73AF"/>	
Page Menu Item Text		<input type="text" value="#333333"/>	
Heading Text		<input type="text" value="#333333"/>	
Space Name Text		<input type="text" value="#999999"/>	
Links		<input type="text" value="#3572B0"/>	
Borders and Dividers		<input type="text" value="#CCCCCC"/>	
Tab Navigation Background		<input type="text" value="#3B73AF"/>	
Tab Navigation Text		<input type="text" value="#FFFFFF"/>	
Tab Navigation Background Highlight		<input type="text" value="#205081"/>	
Tab Navigation Text Highlight		<input type="text" value="#FFFFFF"/>	

Reset

Save

Cancel

Reset your colour scheme after uploading a site logo

When you upload a site logo, Confluence automatically detects the colours in your logo and customises the colour scheme for you.

You can change the colour scheme as above, or reset your colour scheme back to the default (and still keep your new site logo).

To reset the colour scheme:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Colour Scheme** in the left-hand panel.
3. Choose **Edit**.
4. Choose **Reset**.

Notes

- If you make a mistake, just choose **Reset** and then try again.
- Some UI elements are specific to the default theme and may not take effect for other themes.

Styling Confluence with CSS

This page explains the facility for changing the look and feel of Confluence with CSS.

Introduction

Cascading Style Sheets (CSS) are an industry-standard way of styling a web page. The content of a page is rendered with HTML, and its look and feel is determined by CSS files. You can upload a CSS text file, or simply type in a stylesheet, and apply it to a space or even a whole Confluence site.

Note: By default, only system administrators can edit the CSS for a space or for the site. To allow any user with Space Admin permissions to edit the CSS for a space, go to



> **General Configuration** > **Security Configuration** and select **Custom Stylesheets for Spaces**.

Creating CSS styles that work seamlessly across different browsers is a delicate task for basic web sites, and reasonably challenging when customising web applications like Confluence. It is important to test each change that you make and ensure it works as expected in all areas of Confluence – for example, on the Confluence dashboard as well as on regular pages.

In order to get you started, we have compiled this introduction, a [basic styling tutorial](#).

Considerations for Using Custom CSS

CSS Knowledge is Required

If you are not familiar with CSS, see the links in the [CSS Resources](#) section below. You should spend some time to become confident with Cascading Style Sheets before you start editing your Confluence style sheets.

Security

Custom CSS can be used to inject scripts into a page, opening the risk of cross-site scripting (XSS) attacks. With this feature enabled, space administrators could upload styles that steal other users' login credentials, trick their browsers into performing actions on the wiki without their knowledge, or even obtain global

On this page:

- [Introduction](#)
- [Considerations for Using Custom CSS](#)
- [Getting Started](#)
- [CSS Resources](#)

Related pages:

- [Basic Styling Tutorial](#)
- [Styling Fonts in Confluence](#)

administration privileges. As such, this feature is disabled by default. Confluence administrators should only enable custom CSS if they are comfortable with the risks listed in this paragraph.

Scaling

Each page needs to scale. Depending on the resolution of the user's screen, the content should render intelligently. Your designs needs to degrade gracefully. Try resizing each page that exists in Confluence. There are quite a few pages in the browse-space-section, like drafts, labels, page hierarchy, and so on. Your style has to work everywhere, not just in the first page you happen to be looking at.

Features Cannot Be Disabled

It is easy to turn off certain links, headers, or even menu items by simply setting their style to 'hidden'. This can help you to roll out Confluence to users that may not be very Wiki-savvy yet. The simpler the UI, the easier it may be for them to use. However, please remember that removing the link to a part of the application does not mean that the functionality is not available. Every user can still change their style from within their browsers, or access the URL directly. Don't rely on CSS to disable parts of Confluence.

Features Should Not Be Disabled

Users familiar with Confluence will expect to find the same controls that they are accustomed to. Removing buttons or controls from the interface is not advised as it may frustrate your users and cause them to circumvent your design by using direct URL access, as mentioned above.

Custom CSS does not apply to Admin screens

Any CSS styling applied to your site will not be applied to the Administration console. This is to ensure changes to CSS do not prevent administrators from accessing Admin functions in future.

Confluence Version Compatibility

Be aware of any plans to upgrade your Confluence instance. Future versions of Confluence may not be compatible with your custom CSS — this may cause your CSS to break, requiring maintenance when Confluence is upgraded. Ask your Confluence administrator for more information.

Test on Different Web Browsers

As a rule you should test your modifications on different web browsers. Internet Explorer, Firefox, Opera and Safari (on Mac OS X) are some of the more popular browsers.

Note about supported web browsers: Please ensure that you are using one of the web browsers supported by Confluence. If you are using an unsupported browser or browser version, some features may not work correctly. Check the [Supported Platforms](#) page to find the list of supported web browsers and browser versions on this page.

CSS Customisation is Not Supported

As creating custom CSS has potentially limitless possibilities, Atlassian will not support issues that are caused by or related to CSS customisation.

Getting Started

Editing the CSS

To edit a space's CSS style sheets:

1. Go to the space and choose **Space tools > Look and Feel** from the bottom of the sidebar
2. Choose **Stylesheet** then **Edit**.
3. Paste your custom CSS into the text field.
4. Save your changes. The new CSS will be visible on all content pages in the space.


If your space uses the Documentation theme:

1. Choose **Browse > Space Admin** from the header

Note: The **Space Admin** option appears only if you have [space admin permissions](#), or if you're part of the 'confluence-administrators' group.

2. Click **Stylesheet** in the left-hand panel under the heading 'Look and Feel'.
3. Choose **Edit**.
4. Paste your custom CSS into the text field.
5. Save your changes. The new CSS will be visible on all content pages in the space.

To edit your global CSS stylesheet:

1. Choose the **cog icon**  , then choose **General Configuration** under Confluence Administration
2. Choose **Stylesheet**.
3. Choose **Edit**.
4. Paste your custom CSS into the text field.
5. Choose **Save**.

Note:

- The new CSS will be visible across all spaces, provided they do not define their own custom stylesheet and are not using a theme. This CSS will also overwrite all styles defined in custom global themes.
- You may be able to add CSS to your site by choosing **Custom HTML** in the administration section, and adding your CSS definitions to the HEAD or BODY of the page. You should only use this option if you cannot achieve the desired results via the global stylesheet.

Follow the Tutorial

Follow the examples in the [Basic Styling Tutorial](#) to get started.

CSS Resources

- [W3C CSS Standards](#)
- [Mozilla Developer Network](#)
- [W3resource.com](#)

Basic Styling Tutorial

This page contains instructions on how to get started with custom CSS styling in Confluence.

CSS Editing Quick-Start

To edit a space's CSS style sheets:

1. Go to the space and choose **Space tools > Look and Feel** from the bottom of the sidebar
2. Choose **Stylesheet** then **Edit**.
3. Paste your custom CSS into the text field.
4. Save your changes. The new CSS will be visible on all content pages in the space.

If your space uses the Documentation theme:

1. Choose **Browse > Space Admin** from the header

Note: The **Space Admin** option appears only if you have [space admin permissions](#), or if you're part of the 'confluence-administrators' group.

2. Click **Stylesheet** in the left-hand panel under the heading 'Look and Feel'.
3. Choose **Edit**.
4. Paste your custom CSS into the text field.
5. Save your changes. The new CSS will be visible on all content pages in the space.

Tutorial: Changing the Header Background

On this page:

- [CSS Editing Quick-Start](#)
- [Tutorial: Changing the Header Background](#)
- [CSS Editing Tips](#)
- [Notes](#)

Related pages:

- [Styling Confluence with CSS](#)

The header is the menu area at the top of a default Confluence page where the **Breadcrumb Links**, **Browse** menu, **User** menu and the **Quick Search** box reside. In this example, we are going to change the background of the header to include a custom graphic.

1. Create a custom graphic. For this example, we created a custom header graphic of 1046 x 61 pixels.
2. Upload the custom graphic to a page in the space that you are customising.
3. Note the page ID of the page where you uploaded the new graphic. (in this example, the page ID was '**658833839**').
4. Compose your custom CSS for the header. The example below loads the new graphic (called '**header.png**') from a specific page (denoted by page ID '**658833839**') in the same space.

```
#header .mui-header {  
  
background-image:url('../download/attachments/658833839/header.png');  
background-repeat: no-repeat;  
}
```

5. Log in as the Space Administrator.
6. Open the Space Admin page.
7. Click **Stylesheet**.
8. Click **Edit** to change the code in the text field.
9. Paste your custom CSS into the text field.
10. Click **Save** and then reload the page (you may have to shift-reload). The background of the header will change.
11. The custom header will be visible on all content pages in the space. To revert your change, simply delete the custom code from the 'Stylesheet' page and click **Save**.

CSS Editing Tips

Begin With a Space Stylesheet

A space stylesheet is a good starting point for CSS customisation, as it already includes all of the elements that can be changed. When you work on the space stylesheet it styles all content pages in the space. Build and test it at space-level, before considering applying the new stylesheet to your entire site. Once you are satisfied with your space design, test it thoroughly until you are confident that it has no problems. Then, you can look into advanced customisation of the Confluence CSS such as adjusting the Search page, the Dashboard and other integral pages.

Use the Right Tools

As the Confluence CSS is reasonably sophisticated, web development applications will help you to understand how the page styles have been created. In particular, you will need to view the existing source for the pages you're starting to work on. If you don't already have some, tools such as the following free applications will allow you to do this.

1. [Firebug](#)

Firebug, a plugin for the Firefox web browser, allows you to take a look at the style of each element on your page. This is very useful to see what styles are currently applied, for example styles applied to the header only.

2. [Web Developer](#)

The Web Developer plugin for Firefox allows you to edit CSS inline and create new page designs.

3. [CSS Edit](#)

CSS Edit is a stand-alone CSS editor for Macintosh that extracts all existing styles from a given page and allows you to overwrite these.

Edit Simple Elements First

Begin by editing simple elements and checking that they work. By making changes, then checking that each one worked, you can easily isolate any CSS code that is causing problems. Be aware that some page elements are more suited to customisation than others. For example, adding a gradient to the toolbar is less

likely to 'break' the page than changing the page width. Editing reasonably static elements such as background graphics will render more predictably than designs which attempt to completely change the user interface or the Javascript-powered drop-down menus (which we don't recommend editing).

Notes

Note: By default, only system administrators can edit the CSS for a space or for the site. To allow any user with Space Admin permissions to edit the CSS for a space, go to



> **General Configuration** > **Security Configuration** and select **Custom Stylesheets for Spaces**.

Styling Fonts in Confluence

Confluence provides the ability to adjust its visual style via Cascading Style Sheets (CSS). This tutorial shows you to change the fonts and font sizes of a Confluence page, using a few lines of CSS.

Related pages:

- [Basic Styling Tutorial](#)
- [Styling Confluence with CSS](#)

Below is the code for the custom font. Copy and paste it into the Space Stylesheet form within the Space Administration section.

Changing the fonts

In order to customise the fonts in Confluence, you first need to set the body font to the font you want. Secondly, you may want to adjust the font size because different fonts have different relative sizes.

The relevant CSS is shown below. It changes Confluence's font from the default of Helvetica/Arial – *sans serif* to Times/Times New Roman – *serif*. To adjust for the fact that Times is a bit smaller than Helvetica, we increase the font size to 14 pixels. The many styles that 'wiki-content' in their definition are necessary to change the font size for all the tags in the wiki content.

```
body {
    font-family: Times, "Times New Roman", serif;
    font-size: 14px;
}
.wiki-content,
.wiki-content p,
.wiki-content table,
.wiki-content tr,
.wiki-content td,
.wiki-content th,
.wiki-content ol,
.wiki-content ul,
.wiki-content li {
    font-size: 14px;
}
```

Notes

Note: By default, only system administrators can edit the CSS for a space or for the site. To allow any user with Space Admin permissions to edit the CSS for a space, go to



> **General Configuration** > **Security Configuration** and select **Custom Stylesheets for Spaces**.

Working with Themes

Themes are pre-defined style sets that you can apply to Confluence, to alter the appearance of your site. This is a way of personalising the 'look and feel' of Confluence. You can apply a theme to your entire Confluence site and to individual spaces. Choose a specific theme if you want to add new functionality or significantly alter the appearance of Confluence.

Confluence comes with a selection of themes. After a theme is installed on

your Confluence site, any space administrator can apply a theme to a space.

A site administrator can install new themes as add-ons via the Confluence Administration Console.

By default when you create a new space, the space will have the Confluence default theme.

To look at the themes installed on your Confluence site:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Themes** in the left-hand panel.
3. You will see a list of all installed themes.

Related pages:

- [Apply a Theme to a Space](#)
- [Applying a Theme to a Site](#)
- [Configure the Documentation Theme](#)
- [Confluence Administrator's Guide](#)
- [Creating a Theme](#)

Useful add-ons

Before installing an add-on (also called a plugin) into your Confluence site, please check the add-on's information page to see whether it is supported by Atlassian, by another vendor, or not at all. See our guidelines on [add-on support](#). Visit the Atlassian Marketplace to [search for additional themes](#) you can add to your site.

Applying a Theme to a Site

You can use a theme to personalise the 'look and feel' of Confluence. Some themes simply change the basic styling, others add new functionality or significantly alter the appearance of Confluence. You can apply a theme to your entire Confluence site and to individual spaces.

Confluence comes with a selection of themes. In addition, a site administrator can install new themes as plugins via the Confluence Administration Console. Provided that the theme is available in your Confluence site, any space administrator can apply a theme to a space. By default when you create a new space, the space will have the Confluence default theme.

To apply a theme across the site:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Themes** in the left-hand panel.
3. The screen will display all available themes. Select a radio button to choose a theme.
4. Choose **Confirm**.

Related pages:

- [Apply a Theme to a Space](#)
- [Configure the Documentation Theme](#)
- [Confluence Administrator's Guide](#)
- [Creating a Theme](#)

Screenshot: Applying a theme

Site Theme

Current Theme

The current theme controls the layout and colours of this space.



Default Theme

This is the original Confluence look and feel. Page content spans the full width of the screen.

Choose New Theme

To change the theme of this space, select one below.



☐ Documentation Theme

This theme is well suited for structured content, such as documentation. It features a table of contents (page tree) on the left, making it easier to see the structure of a space and move from page to page. You can customise the left-hand panel, page header and page footer.



[Find more themes...](#)

Confirm

Creating a Theme

If you want to create your own theme, you will need to write a Confluence plugin. Please refer to the following pages in our developer documentation:

- Get started with [plugin development](#).
- Follow the developer's tutorial for [writing a Confluence theme](#).
- Create a theme using the [theme plugin module](#).

Related pages:

- [Applying a Theme to a Site](#)
- [Apply a Theme to a Space](#)
- [Configure the Documentation Theme](#)
- [Confluence Administrator's Guide](#)

Customising Site and Space Layouts

You can modify Confluence's look and feel by editing layout files (also known as decorators). Editing these files allows you to change the look and feel of the whole Confluence site, or just an individual space.

When you edit a site layout, you'll be modifying the default decorators in every space in your site, except

for those that have already been edited in a space. See [Customise Space Layouts](#) for more information on how to edit the decorators for a single space.

You'll need [System Administrator](#) permissions to edit site layouts.

On this page:

- [Editing a site decorator file](#)
- [Using Velocity macros](#)
- [Advanced customisations](#)

Related pages:

- [Velocity Template Overview](#)
- [Basic Introduction to Velocity](#)
- [Customising your Confluence Site](#)

If you modify layouts in Confluence you will need to reapply your modifications each time you upgrade Confluence. The more dramatic your customisations are, the harder it may be to reapply the changes when upgrading. See [Upgrading Customised Site and Space Layouts](#) to find out what will be involved before modifying the layouts.

Confluence is built on top of the open source [SiteMesh](#) library, a web-page layout system.

To edit the layout of Confluence, you will need to modify these decorator files. A decorator file is a `.vmd` file and is written in Velocity. You can learn more from the [Velocity User Guide](#).


Once you are familiar with Velocity, you can edit the decorator files to personalise the appearance of Confluence.

The decorator files in Confluence are grouped into the following categories:

- **Site layouts:** These are used to define the controls that surround each page in the site. For example, the header, footer and dashboard.
- **Content layouts:** These control the appearance of content such as pages and blog posts. They do not change the way the pages themselves are displayed, but allow you to alter the way the surrounding comments or attachments are displayed.
- **Export layouts:** These control the appearance of spaces and pages when they are exported to HTML.

Editing a site decorator file

To edit a site decorator:

1. Go to  **> General Configuration > Layouts** (under **Look and Feel**)
2. Click **Create Custom** next to the decorator `.vmd` file you want to modify.
3. Make your changes and click **Update**.

If something goes wrong: Hit **Reset Default** to revert to the original layouts.

Using Velocity macros

When editing Custom Decorator Templates, there are a number of macros available to define complex or variable parts of the page such as menus and breadcrumbs. You may insert these macros anywhere in your templates. More information on [Working With Decorator Macros](#).

Advanced customisations

Overriding Velocity templates

The `velocity` directory is at the front of Confluence's Velocity template search path. As such, you can override *any* of Confluence's Velocity templates by placing an identically named file in the right place. While we don't recommend you do this unless you know exactly what you're doing, it does give you complete

control over the look of every aspect of Confluence. It also means that you can edit your templates in a text-editor if you wish, rather than through the web interface.

Caching

Velocity is configured to cache templates in memory. When you edit a page from within Confluence, it knows to reload that page from disk. If you are editing the pages on disk, you will either have to turn off velocity's caching temporarily in `WEB-INF/classes/velocity.properties`, or restart the server to make your changes visible.

Location of Velocity files

You will find the Velocity files in your Confluence installation directory. The primary Velocity files are located in the `<CONFLUENCE-INSTALLATION>\confluence\decorators` directory. For example, you will find the following files in that directory: `main.vmd`, `space.vmd`, `form-aui.vmd`, `global.vmd`, and more.

Finding the layout via the URL

If the layout has changed so extensively as to not be visible, you can browse to the URL directly:

```
http://<confluence base url>/admin/resetdecorator.action?decoratorName=decorators/main.vmd
```

Substitute the base URL and the appropriate `.vmd` file.

Upgrading Customised Site and Space Layouts

As Confluence evolves, so do the default site and space layouts that drive the rendering of every page. As new functionality is added or current functionality is changed, the default layouts are modified to support these changes.

Related pages:

- [Customising Site and Space Layouts](#)

If you are using [custom layouts](#) based on defaults from a previous Confluence version, you run the risk of **breaking functionality**, or worse, **missing out on great new features!**

Take care on each new release of Confluence to reapply your changes to the new default templates.

To reapply your custom layouts, you need to:

1. Obtain the source of your custom layouts from your current version of Confluence.
2. Reapply your customisations to the new default layouts.

Step 1. Obtain your Custom Layouts

Ideally, you should keep a record of each customisation you have applied to each of your Confluence site or space layouts.

If not, you should be able to find your customisations using the following method. This method extracts all site- and space-level layouts from your Confluence site as a single output. From this output, you should be able to identify your customisations.

This method is handy to use if you have:

- Many spaces with space layout customisations, or
- Do not have an independent record of your site or space layout customisations.

Custom layouts are stored in the `DECORATOR` table within your Confluence database. You can `SELECT` for the source of the layout using SQL like this:

```
mysql> select SPACEKEY,DECORATORNAME,BODY from DECORATOR;
+-----+-----+-----+
| SPACEKEY | DECORATORNAME | BODY |
+-----+-----+-----+
| NULL     | decorators/main.vmd | ... |
+-----+-----+-----+
1 row in set (0.03 sec)
```

This example was tested on **MySQL**, but should be applicable to all SQL databases.

Step 2. Reapply your Customisations

When you upgrade Confluence to another major release of Confluence, you will need to manually reapply any customisations you made to any site-wide or space-specific layouts. Unless otherwise stated, you should not need to reapply customisations after conducting a minor release upgrade of Confluence.

What are 'major' and 'minor' releases? Major release upgrades are ones where the 1st digit of Confluence's version number or the 1st digit after the 1st decimal place differ after the upgrade, for example, when upgrading from Confluence 3.0 to 3.1, or 2.8 to 3.0. Minor release upgrades are ones where the 1st digit of Confluence's version number and the 1st digit after the 1st decimal place remain the same after the upgrade, for example, when upgrading Confluence 3.0 to 3.0.1.

If you have made Confluence site-wide layout customisations:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Select **Layouts** in the left-hand navigation panel. The decorators are grouped under **Site**, **Content** and **Export** layouts.
3. Ensure you have all your customisations available (preferably in a form which can be copied and pasted).
4. Click **Reset Default** next to the layout whose customisations need to be reapplied.
5. Click **Create Custom** next to the same layout and reapply your customisations (by copying and pasting them) into the appropriate locations within the new default layout.
6. Click the **Save** button.
7. Repeat this procedure from step 4 for each layout whose customisations need to be reapplied.

If you have made space-specific layout customisations:

1. Go to the space and choose **Space tools > Look and Feel** from the bottom of the sidebar
2. Choose **Layout**. The decorators are grouped under **Site**, **Content** and **Export** layouts.
3. Ensure you have all your customisations available (preferably in a form which can be copied and pasted).
4. Click **Reset Default** next to the layout whose customisations need to be reapplied.
5. Click **Create Custom** next to the same layout and reapply your customisations (by copying and pasting them) into the appropriate locations within the new default layout.
6. Click the **Save** button.
7. Repeat this procedure from step 5 for each layout whose customisations need to be reapplied.

If your space is using the Documentation theme:

1. Go to a page in the space.
2. Choose **Browse > Space Admin** from the header

Note: The **Space Admin** option appears only if you have [space admin permissions](#), or if you're part of the 'confluence-administrators' group.

3. Choose **Layout** from the left hand panel.
4. Follow the steps above.

Step 3. Test your Modifications Carefully

Changes may interact unpredictably with future versions of Confluence. When upgrading, you should always

test your custom modifications thoroughly before deploying them on a live site. It's beyond the scope of Atlassian Support to test and deploy these changes.

Turning Off Caching

Velocity is configured to cache templates in memory. When you edit a page from within Confluence, it knows to reload that page from disk. If you are editing the pages on disk, you will either have to turn off Velocity's caching temporarily in `WEB-INF/classes/velocity.properties`, or restart the server to make your changes visible.

The `velocity.properties` file is available in the `confluence-x.x.x.jar` file, where `x.x.x` is the Confluence version number. The JAR file is located in the `WEB-INF/lib` directory. If you wish to make modification to the files in the JAR, we recommend the following steps:

1. Stop Confluence.
2. Make a backup copy of the JAR file.
3. Un-jar the file
4. Locate and edit the appropriate file that you wish to modify.
5. Re-jar the `confluence-x.x.x.jar` file.
6. Relocate the JAR file to the appropriate directory.
7. Restart Confluence.

Working With Decorator Macros

Decorator Macros are [Velocity](#) macros which are used to draw complex or variable parts of the page such as menus and breadcrumbs when editing [Custom decorators](#). Decorator macros can be inserted anywhere in your templates.

The macro is called by inserting a string of the form: `#macroName("argument1" "argument2" "argument3")`. There are no commas between the arguments. Unless otherwise noted, these macros take no arguments.

NOTE: These macros will only work reliably when customising `main.vmd`. They may not work in other Velocity decorators. Decorator macros will not work inside normal confluence pages.

Macro	Usage
<code>#breadcrumbs()</code>	Draws the "You are here" breadcrumbs list, like the one found above the page name in the default template.
<code>#includePage(pageTitle)</code>	Includes a confluence page with the specified title. If you have 2 or more pages with the same title across multiple spaces, this macro will include the page belonging to the space you are currently viewing.
<code>#searchbox()</code>	Inserts a search box into the page, like the one to the far right of the breadcrumbs in the default template.
<code>#globalnavbar(type)</code>	Draws the global navigation bar, as found in the top right-hand corner of the default template. The navigation bar can be displayed in two modes:
<code>#globalnavbar("table")</code>	Displays the navigation bar in its default mode: drawn as a table of links with coloured backgrounds and mouse-over effects.
<code>#globalnavbar("text")</code>	Displays the navigation bar as series of text links separated by characters.
<code>#usernavbar()</code>	Draws the user-specific navigation-bar. This bar contains the links to the user's profile and history, or to the login and signup pages if the user is not logged in.
<code>#helpicon()</code>	Draws the help icon, and link to the Confluence help page.

<code>#printableicon()</code>	On pages where a printable version is available, draws the printable page icon, linking to the printable version of the page. Otherwise, draws nothing
<code>#pagetitle(class)</code>	When you are viewing a page in a Confluence space, draws the name of the space that page is in. Otherwise, writes the word "CONFLUENCE". The "class" argument is the CSS class that the title should be drawn in. Unless you have customised your Confluence installation's CSS file, you should call this with "spacename" as the class: <code>#pagetitle("spacename")</code>
<code>#poweredby()</code>	Writes out the "Powered by Confluence" and Confluence version-number boilerplate found at the bottom of the default template.
<code>#bottomshadow()</code>	Draws the fading shadow-effect found at the bottom of the content area in the default template.
<code>#dashboardlink()</code>	Inserts a link to the dashboard page.

Custom Decorator Templates

About Decorators

Confluence is built on top of the Open Source SiteMesh library, a web-page layout system that provides a consistent look and feel across a site. SiteMesh works through "decorators" that define a page's layout and structure, and into which the specific content of the page is placed. If you are interested, you can read more in the [SiteMesh](#) documentation.

What this means for Confluence is that you can customise the look and feel of parts of your Confluence site through editing decorators, for example:

- The "Main" decorator defines the generic header and footer
- The "Page" decorator defines how a page is displayed
- The "Printable" decorator defines the look and feel of the printable versions of pages.

You can view and edit these decorators from within Confluence. Changes to the decorators will affect all spaces in that Confluence installation.

The decorator that is used to draw Confluence's administrative pages cannot be edited from within Confluence. This means that if you make a mistake that renders the rest of the site unuseable, the administrative pages should still be available for you to fix the template.

Browsing the Default Decorators

At any time, you can browse the default decorators that come packaged with Confluence by following the "View Default" links on the "Site Layouts" page. The template browser also allows you to view the "#parsed" templates that are included within the template when it is compiled. While you can't edit these included templates, you will probably have to copy some or all of them into your custom template as you do your customisation.

Editing Custom Decorators

To edit Confluence decorators you will need a good knowledge of [HTML](#), and some understanding of the [Velocity templating language](#).

To edit a decorator:

1. Go to **Confluence Admin > Layouts**.
2. Choose **Create Custom** beside the decorator you wish to edit.
3. Save your changes.

If you make a mistake or want to undo your changes, choose **Reset Default** beside the edited decorator.

Alternatively, the custom templates are stored in the DECORATOR table in the database. If you have somehow managed to render Confluence completely unuseable through editing your templates, delete the relevant entries from the DECORATOR table.

Macros

Some parts of the page are drawn using Velocity macros, including the navigation bar. The macros you should know about when editing decorators are described in [Working With Decorator Macros](#).

For Advanced Users

The `velocity` directory is at the front of Confluence's velocity template search path. As such, you can override *any* of Confluence's velocity templates by placing an identically named file in the right place.

While we don't recommend you do this, it does give you complete control over the look of every aspect of Confluence. It also means that you can edit your templates in a text-editor if you wish, rather than through your browser.

There are, however, two important caveats:

1. Velocity is configured to cache templates in memory. When you edit a page from within Confluence, it knows to reload that page from disk. If you are editing the pages on disk, you will either have to turn off velocity's caching temporarily in `WEB-INF/classes/velocity.properties`, or restart the server to make your changes visible.
2. Changes may interact unpredictably with future versions of Confluence. When upgrading, you should always test your custom modifications thoroughly before deploying them on a live site.

Customising a Specific Page

If you'd like to change the appearance of a specific page, you can modify the corresponding Velocity template. Here's how to find out which one:

1. Access the page. Note the name of the action. For example, the "Contact Administrators" page is `<baseUrl>/administrators.action`.
2. Browse to `<confluence-install>/confluence/WEB-INF/lib/confluence-x.y.jar`. Copy the file.
3. Unzip or unjar the file using a standard unzipper or the [java jar utility](#).
4. Open `xwork.xml`. Search the file for the name of the action corresponding to the page you'd like to modify. You'll see an entry like:

```
<action name="administrators"
class="com.atlassian.confluence.user.actions.AdministratorsAction">
    <interceptor-ref name="defaultStack"/>
    <result name="success" type="velocity">/administrators.vm</result>
</action>
```

5. The file to look for is the `vm` or `vmd` file. In the above example, it's `administrators.vmd`. Because there is no context path (just a `/` before the name of the file), it's in the root of the Confluence webapp. For the stand-alone, that's `<confluence-install>/confluence` folder.
6. Modify the file.

For details on how to configure the file, check the [Velocity Template Overview](#).

Customising the Login Page

This page gets you started on customising the Confluence login page, to add your own logo or custom text. This will not customise the login *process*, just what users sees when they log in.

Notes:

- Customisations to the Confluence login page will need to be reapplied when you upgrade Confluence. Consider this before making drastic changes to the layout, and be sure to keep a list of what you have changed for your upgrade process later.
- Please test your changes on a **test** Confluence site first.

Only administrators with access to the server where Confluence is running can modify the Confluence login page.

Related pages:

- [Changing the Site Logo](#)
- [Velocity Template Overview](#)
- [Customising Site and Space Layouts](#)
- [Changing the Look and Feel of Confluence](#)
- [Modify Confluence Interface Text](#)

To change the login page:

1. Shut down your Confluence server.
2. In the Confluence installation directory, find the file `confluence/login.vm`.
3. Make a copy of this file as a backup.
4. Edit the file with a text editor to make the required changes. The content contains a mixture of HTML and Velocity. See [Velocity Template Overview](#) (in our developer documentation).
5. Start Confluence and test your changes.

The same process can be applied to modify most of the templates in the Confluence web application. Be careful to test your changes before applying them to a live site. The templates contain code that is vital for Confluence to function, and it is easy to accidentally make a change that prevents use of your site.

Modify Confluence Interface Text

All Confluence UI text is contained in a single Java properties file. This file can be modified to change the default text, and also to translate Confluence into languages other than English.

The UI text file is `ConfluenceActionSupport.properties`. From your Confluence install directory:

```
\confluence\WEB-INF\lib\confluence-x.x.x.jar
```

Replace "x.x.x" with your Confluence version, for example for 4.3.2, it will be named "confluence-4.3.2.jar".

Within this File, the relevant file to edit is

```
:\com\atlassian\confluence\core\ConfluenceActionSupport.properties.
```

Refer to [Editing jar files](#) for reference.

The file contains parameters with `name=value` pairs, in the format:

```
parameter.name=Parameter value
```

Parameter names are any text before the '=' character and should never be modified. Any text after the '=' character is the parameter value, which can be modified freely and can also contain variables. An example involving variables is:

```
popular.labels=The three most popular labels are {0}, {1} and {2}.
```

For more information on replacing values, check out [Translating ConfluenceActionSupport Content](#). Note that plugins store their text internally, so you must modify plugin text individually.

Steps For Modification

1. Stop Confluence
2. Under your install directory, open `\confluence\WEB-INF\lib\confluence-x.x.x.jar\com\atlassian\confluence\core\ConfluenceActionSupport.properties`
3. Search for the text you wish to modify, replace it and save the file in `<Confluence-Install>\confluence\WEB-INF\classes\com\atlassian\confluence\core`. Please create this folder structure, if it does not exist already.

If you re-bundle the JAR file, rather than re-deploy the class in the `WEB-INF\classes` directory, make sure to move the backup JAR file out of the `/lib` directory, or the backup may be deployed by mistake.

4. Restart Confluence

Modify Keyboard Shortcuts

Confluence provides a set of [keyboard shortcuts](#). You could customise the shortcuts by making modifications inside the `ConfluenceActionSupport.properties` file.

- To disable a particular shortcut, you can simply just comment out a respective line of code. One may like to disable the shortcut to one of the navigation links: *View*, *Edit*, *Attachments*, *Info*. For instance, to disable shortcut to *Attachments* one would comment out the following line:

```
#navlink.attachments.accesskey=a
```

- To modify an access key, one could simply just change the letter, bearing in mind the fact that the letter must be unique.

Customising the eMail Templates

Customising the Confluence email templates is **not supported**. If you do decide to edit the templates we strongly recommend you use a test instance of Confluence.

Any customisations you make to the Confluence email notification templates will need to be reapplied after upgrading Confluence.

Email notification templates are contained within the `confluence-email-notifications` plugin, which is a bundled plugin (add-on) that is installed automatically when you install Confluence.

Only administrators with access to the Confluence installation directory can modify the Confluence email templates.

Confluence uses Soy templates (also known as Closure templates) for email notifications. You can find out more in the [Google Developer docs](#) or see our [developer tutorial](#) which contains a short introduction to using Soy templates.

To change the email notification templates:

1. In the Confluence web application folder, find the file `/confluence/WEB-INF/atlassian-bundled-plugins/confluence-email-notifications-plugin-x.x.jar`
Note: This plugin is independently versioned, the version number will not necessarily match Confluence's version number.
2. Copy this file to a working location and extract the jar file. Find out more about how to [edit files within .jar archives](#).
3. Within the jar file, templates are stored in the `/templates/` folder. Edit the Soy templates to make your changes.
4. Zip all the files and change the file extension to `.jar` (or refer to the [guide on editing files within .jar archives](#) for other methods).
5. Drop the new jar file into the `/confluence/WEB-INF/atlassian-bundled-plugins` folder (replacing the original file - you might want to make a copy of the original file for easy roll back) and then restart your instance.
6. Test your changes carefully before installing the updated plugin in production.

We strongly recommend you use a test instance for editing the templates contained within the plugin. If you are unable to enable the plugin, check the Confluence logs for information, it may be that there are problems with your edits to the Soy templates.

RELATED TOPICS

- [Customising Site and Space Layouts](#)
- [Changing the Look and Feel of Confluence](#)
- [Modify Confluence Interface Text](#)

Changing the Default Behaviour and Content in Confluence

Confluence comes with some handy default settings that determine what people see when they first enter the Confluence site, and the default content that is put into new spaces and other areas of Confluence.

Confluence administrators can change the settings to customise the behaviour and the default content of their Confluence site:

- [Administering Site Templates](#)
- [Importing Templates](#)
- [Changing the Site Title](#)
- [Choosing a Default Language](#)
- [Configuring the Administrator Contact Page](#)
- [Configuring the Site Home Page](#)
- [Customising Default Space Content](#)
- [Editing the Site Welcome Message](#)

Related pages:

- [Changing the Look and Feel of Confluence](#)
- [Customising your Confluence Site](#)
- [Confluence Administrator's Guide](#)

Administering Site Templates


A template is a predefined page that can be used as a prototype when creating new pages. Templates can be created by users, or provided by a blueprints. See [Page Templates](#) and [Blueprints](#).

Administrators can import templates, to make them available to other people using Confluence. See [Importing Templates](#).

Confluence also provides 'system templates' which contain default content for the site welcome message (see [Editing the Site Welcome Message](#)) and default space content (see [Customising Default Space Content](#)).

Administrators can also disable templates and blueprints, to stop them appearing in the Create and Create Space dialogs anywhere in their Confluence site.

To **disable** a template or blueprint across the entire Confluence site:

- Choose the **cog icon**  , then choose **General Configuration** under Confluence Administration
- Choose **Global Templates and Blueprints**.
- Choose **Disable** next to the template, page blueprint or space blueprint you wish to disable.

Administrators can re-enable these templates and blueprints at any time.

Importing Templates

A template is a predefined page that can be used as a prototype when creating new pages. Templates are useful for giving pages a common style or format.

You can create your own templates within Confluence. See [Create a Template](#).

In addition, you can download pre-defined templates from the [Atlassian Marketplace](#) in the form of a template bundle. Each template bundle contains one or more templates, created by Atlassian or third parties. Here is a summary of the steps required:

- Download the template bundle from the Atlassian Marketplace.
- Install the template bundle into your

On this page:

- [Step 1. Check the template bundles installed on your Confluence site](#)
- [Step 2. \(Optional\) Download and install additional template bundles from the Atlassian Marketplace](#)
- [Step 3. Import the templates to make them available to users](#)
- [Notes](#)

Related pages:

- [Pages and Blogs](#)
- [Page Templates](#)

- Confluence site.
- Make the templates available by importing them into the site or into an individual space.

You need 'System Administrator' permission to install template bundles into your Confluence site. You need 'Confluence Administrator' permission to manage the existing template bundles on your Confluence site. See [Global Permissions Overview](#).

Step 1. Check the template bundles installed on your Confluence site

To see the template bundles that are currently available for import on your Confluence site:

1. Log in to Confluence as a System Administrator or Confluence Administrator.
2. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

3. Choose **Import Templates** in the left-hand panel. You will see a list of the template bundles installed on your Confluence site, and the templates included in each bundle.

Step 2. (Optional) Download and install additional template bundles from the Atlassian Marketplace

Follow the steps below if you want to add more template bundles to your site.

Before installing an add-on (also called a plugin) into your Confluence site, please check the add-on's information page to see whether it is supported by Atlassian, by another vendor, or not at all. See our guidelines on [add-on support](#).

To upload more templates:

1. Go to the [Atlassian Marketplace](#) and download the template bundle that you need. It will be in the form of a JAR file. Save the JAR file somewhere in your file system.
2. Log in to Confluence as a System Administrator.
3. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

4. Choose **Manage Add-ons** in the left-hand panel.
5. Choose **Upload Add-on**.
6. Browse to find the template bundle that you downloaded, and upload it to Confluence. The template bundle will appear in the list under 'User-installed Add-ons'.

Step 3. Import the templates to make them available to users

You now have one or more template bundles on your site. The templates are not available until you have 'imported' them.

To import a template:

1. Log in to Confluence as a System Administrator or Confluence Administrator.
2. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

3. Choose **Import Templates** in the left-hand panel. You will see the template bundles installed on your Confluence site and the templates included in each bundle.

Note: You can see a preview of the template by choosing the template name.

4. Select the templates to be imported by ticking the check boxes next to the relevant template names.
5. Choose the import destination for the templates in the **Import To** dropdown menu. If you want the templates to be available to only a specific space, choose the name of the space, otherwise choose **Global Templates** to make the templates available to all spaces.
6. Choose **Import**.

Screenshot: Importing a template

Import Templates

The following template package plugins were found. To import templates, tick the checkboxes for the desired templates, select where to import the templates to, and click the Import button.

Human Resources Templates (8)

[Check All](#) [Uncheck All](#)

- ☐ HOW-TO Guide
- ☐ Induction Tasks
- ☐ Job Description
- ☐ Meeting Minutes
- ☐ Recruitment Dashboard
- ☐ Space Home Page
- ☐ Time Sheet
- ☒ Wiki Induction

Import To: Global Templates ▼

Global Templates
Spaces
 Demonstration Space
 Documentation

Notes

- **Building your own template bundles.** You can build a template bundle as an add-on (also called a 'plugin') and then upload it to your Confluence site. You can then import the templates from your custom template bundle, as described above. You will need some programming knowledge to develop a template bundle. See [Creating A Template Bundle](#).
- **Duplicate template names.** If a template with the same name already exists on import, a duplicate template of the same name will be created. You will need to check the templates and rename them manually.
- **Removing the template.** Removing the add-on that contains a template will not remove the template from your Confluence site if you have already imported it. You will need to remove the template manually via the administration console or space administration screen.

Changing the Site Title

The site title appears in your browser's title bar. By default, it is set to 'Confluence'.

To change the title of your Confluence site:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose '**General Configuration**' in the left-hand panel.
3. Choose '**Edit**' at the top of the '**Site Configuration**' screen.
4. Enter a new title for your site in the input field next to '**Site Title**'.
5. Choose '**Save**'.

Related pages:

- [Changing the Site Logo](#)
- [Editing the Site Welcome Message](#)
- [Customising your Confluence Site](#)
- [Confluence Administrator's Guide](#)

Choosing a Default Language

Administrators can define a default language to be applied to all spaces in your Confluence site. Note that individual users can select a language preference for their session.

Related pages:

- [Edit Your User Settings](#)
- [Recognised System Properties](#)
- [Configuring Indexing Language](#)
- [Installing a Language Pack](#)

Setting the default language

To change the default language for the Confluence site:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Select '**Languages**' in the 'Configuration' section of the left-hand panel
3. Choose **Edit** and select the language you want to use as the default language for your Confluence site

Screenshot: Default languages in Confluence



Other settings that affect the language

Individual users can choose the language that Confluence will use to display screen text and messages. Note that the list of supported languages depends on the language packs installed on your Confluence site.

The language used for your session will depend on the settings below, in the following order of priority from highest to lowest:

- The language preference defined in your user profile. Note that you need to be logged in for this setting to take effect.
- The language that you choose by clicking an option at the bottom of the Confluence login screen. Confluence stores this value in a cookie. When the cookie expires, the setting will expire too.
- The language set in your browser. The browser sends a header with a prioritised list of languages. Confluence will use the first supported language in that list. Your Confluence administrator can disable this option by setting a system property.
- The default language for your site, as defined by your Confluence site administrator.

Showing User Interface Key Names for Translation

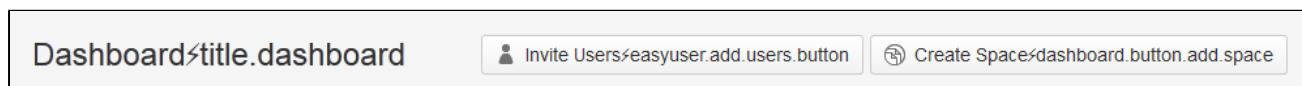
This feature is useful if you are working on creating translations of the Confluence user interface. After opening the Confluence dashboard, you can add this text to the end of your Confluence URL:

```
?i18ntranslate=on
```

Then press Enter.

This will cause each element of the user interface to display its special **key name**. This makes it easier to find the context for each key within the user interface. You can then search for the key on <http://translations.atlassian.com> where you can enter an appropriate translation for your custom language pack.

The key names are displayed with a 'lightning bolt' graphic. For example:



To turn off the translation view, add this code to the end of the Confluence URL:

```
?i18ntranslate=off
```

Configuring the Administrator Contact Page

The administrator contact page is a form that allows a user of Confluence to send a message to the administrators of their Confluence site. (In this context, administrators are the members of the default administrators group.)

See the explanation of [Confluence Groups for Administrators](#).

The title of the administrator contact page is 'Contact Site Administrators'. Typically, Confluence users may get to this page by clicking a link on an error screen such as the '500 error' page.

On this page:

- [Customising the Administrator Contact Message](#)
- [Disabling the Administrator Contact Form](#)
- [Configuring Spam Prevention](#)

Related pages:

- [Configuring Captcha for Spam Prevention](#)

Customising the Administrator Contact Message

You can customise the message that is presented to the user on the '**Contact Site Administrators**' page.

To edit the administrator contact message:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit** at the top of the 'Site Configuration' section.
4. Enter your text in the **Custom Contact Administrators Message** box. You can enter any text or [Confluence wiki markup](#).
5. Choose **Save**.


The Default Administrator Contact Message

By default, the 'contact administrators message' looks much like the highlighted area in the screenshot below, starting with 'Please enter information...'.

Screenshot: The default 'Contact Site Administrators' message

Contact Site Administrators

Please enter information about your request for the site administrators. If you are reporting an error please be sure you include information on what you were doing and the time the problem occurred.



To

From

Subject*

Request Details*

To restore the message to its default simply remove the custom message you entered when following the instructions [above](#), so that the 'Custom Contact Administrators Message' field is empty.

Disabling the Administrator Contact Form

If you prefer to disable the ability for users to send an email message to the site administrators, you can disable the form portion of this screen. You can only disable the form if you first provide a 'Custom Contact Administrators Message' as described [above](#).

To enable or disable the administrator contact form:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit** at the top of the 'Site Configuration' section.
4. Select **on** or **off** for the 'Contact Administrators Form'.
5. Choose **Save**.

Configuring Spam Prevention

You can configure Confluence to use Captcha to help prevent spam, including the spamming of Confluence administrators. The administrator contact form is covered by the site-wide Captcha settings as documented in [Configuring Captcha for Spam Prevention](#).

Configuring the Site Home Page

The dashboard is the default home page for your site, but you can choose to use a space homepage as the landing page for your site.

This can be useful if most people will be reading, rather than creating, pages in your site. However, for sites where you want to encourage teams to collaborate, the dashboard provides the best tools for resuming work in progress and keeping up with what is happening in the site.

Users can also choose to override the site homepage and use the dashboard or a different page as their landing page in their [personal settings](#).

Related pages:

- [Editing the Site Welcome Message](#)
- [Changing the Site Title](#)
- [Changing the Site Logo](#)

To use a page as your site home page:

1. Go to



> **General Configuration** > **Further Configuration**.

2. Choose **Edit**.
3. Select a space from the **Site Homepage** dropdown menu.
When users log in or click the site logo, Confluence will go to the home page of the space you choose here.
4. Choose **Save**.

Note about permissions

Before changing the site homepage you should check that the default 'confluence-users' or 'users' groups have permissions to view the space the page was created in, and that the page itself is not restricted to particular people or groups.

If your site is public, you'll also need to make sure anonymous users have permissions to view the space, otherwise anonymous users will be directed to the dashboard instead.

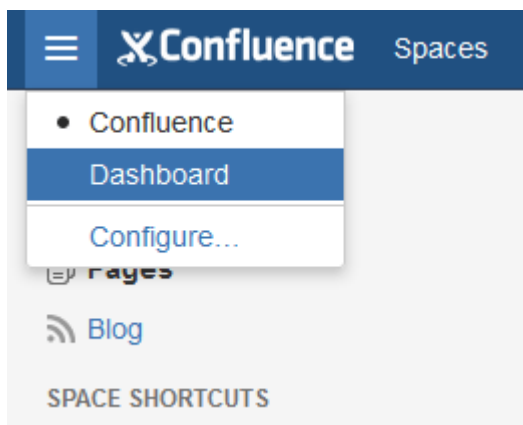
Accessing the dashboard with a site homepage set

If you choose to set a page as your site homepage but would like your users to still be able to access the Confluence dashboard, you can add a link to the Application Navigator.

To add the Confluence Dashboard to the Application Navigator:

1. Go to  > **General Configuration** > **Application Navigator**.
2. Enter the name for your link, for example, 'Dashboard'.
3. Enter the URL for your site dashboard, for example, `https://yoursite.com/wiki/dashboard.action`.
4. Choose **Add**.

A link to the dashboard will now appear in the Application Navigator.



Customising Default Space Content

[Confluence Administrators](#) can edit the template that is used to create the home page for new sites. This default content appears on the home page when a new space is created. There is a different template for site spaces, personal spaces and space blueprints.

The default content in the template only appears for new spaces (those that are created after you have defined the content). Changes to the template do not affect existing home pages.

Edit the default home page for a blank space

To edit the default (blank) space content template:

1. Choose the **cog icon**



- , then choose **General Configuration** under Confluence Administration
- Choose **Global Templates and Blueprints** in the left-hand panel.
- Choose **Edit** next to 'Default Space Content' or 'Default Personal Space Content' depending on whether you want to customise the content for new site space or personal space home pages.
- Enter the content that you want to appear on the home page for new blank spaces. you can add variables, macros and other content in the same way as edited a page template.
- Choose **Save**.

On this page:

- [Edit the default home page for a blank space](#)
- [Reset the original default content](#)

Related pages:

- [Spaces](#)
- [Page Templates](#)
- [Confluence Administrator's Guide](#)

The following variables are available to be added to the default space content templates.

- **\$spaceKey** - inserts the space key into the site space homepage
- **\$spaceName** - inserts the space name into the site space homepage
- **\$userFullName** - inserts the user (owner of the personal space) into the personal space homepage
- **\$userEmail** - inserts the email address of the user (owner of the personal space) into the personal space homepage.

Default space templates differ from ordinary page templates in that they do not present the user with a form to complete, so variables should be limited to those listed in the **Variables** menu.

Some macros, such as the Table of Contents macro, may not display correctly when you preview the template as they are designed to work on a page. The macros will display correctly on the home page when you create a new space. For more information on editing a template, including adding macros see - [Adding Content to a Template](#).

Reset the original default content**To reset the original default content:**

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Global Templates and Blueprints** in the left-hand panel.
3. Choose **Reset to default** next to the template you wish to reset.

From this point on, all new space home pages will be created with the original default content.

Editing the Site Welcome Message

Give your site's landing page some personality by editing the site welcome message.

The site welcome message appears on the right hand side of the dashboard and is perfect for adding announcements, useful links, or a fun photo from your last office party or team outing.

You'll need Confluence administrator permissions to edit the site welcome message.

To edit the site welcome message:

Confluence administrators can either click the **Edit** link below the site welcome message on the dashboard, or:

1. Go to

**On this page:**

- [Hints for using the template editor](#)
- [Allowing other people to edit the site welcome message](#)

Related pages:

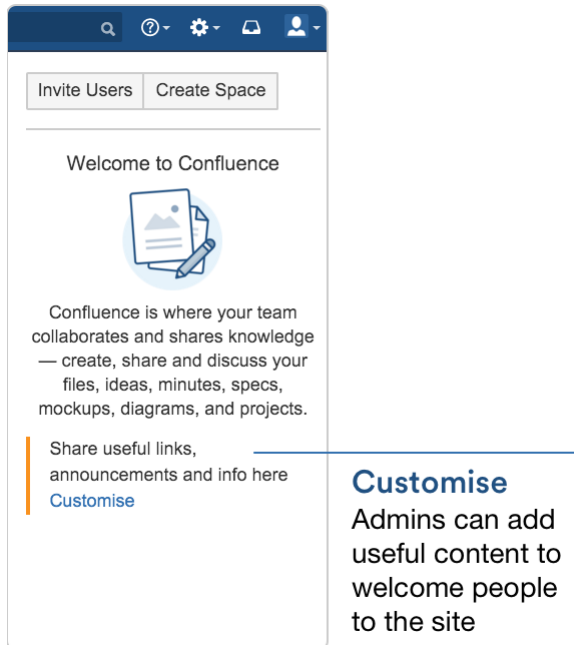
- [Configuring the Site Home Page](#)
- [Changing the Site Title](#)
- [Changing the Site Logo](#)

> **General Configuration > Global Templates and Blueprints.**

2. Scroll down to the System templates and choose **Edit** next to **Default Welcome Message**.
3. Add your content and choose **Save**.

You can go back to the original welcome message at any time - choose **Reset to Default** next to the **Default welcome message** template.

Screenshot: Default site welcome message



Hints for using the template editor

The site welcome message is a template, not a page, so you'll be using the template editor to make your changes.

You can add text, links and macros, as you would in any confluence page, but the process for adding files, including images is a little different.

You can't upload an image or other file into a template directly. First you'll need to upload the file to a page in your site, then in your template, choose **Insert > Files > Search on other pages** to embed the file or image.

You can't use template variables in the site welcome message.

Allowing other people to edit the site welcome message

You can allow people who are not Confluence administrators to edit the site welcome message by using the include [Include Page](#) macro to include content from elsewhere in your site, rather than adding content directly to the template.

To include content from a page in the site welcome message:

1. Create a new page in a space that is visible to all users. It's important that all users can see content in that space - if a person does not have permissions to view the space where you've created the page, they won't be able to see the page content on the dashboard.
2. Add some text, images or macros, then save the page.
3. [Restrict](#) who can edit the page (this is optional, but useful if you only want to allow some people to change the content).
4. Edit the site welcome message template (as described above) and use the [Include page](#) macro to include the contents of your newly created page.
5. Save the template.

People with permission to edit the page will now be able to make changes at any time, and their changes will be visible on the dashboard as soon as the page is saved.

Integrating Confluence with Other Applications

You can integrate Confluence with other applications using **Application Links**. The Application Links feature allows you to link Confluence to applications such as JIRA Software or JIRA Service Desk.

Linking two applications allows you to share information and access one application's functions from within the other. For example, you can display a list of issues on a Confluence page using the [JIRA Issues Macro](#).

Related Topics

- [Linking to Another Application](#)
- [Configuring Workbox Notifications](#)
- [Integrating JIRA and Confluence](#)
- [Registering External Gadgets](#)
- [Configuring the Office Connector](#)

Linking to Another Application

Application Links (sometimes called "AppLinks") is a bundled plugin that allows you to link Atlassian applications to each other. Linking two applications allows you to share information and access one application's functions and resources from within the other.

Atlassian only recommends using OAuth authentication for application links, because of the greater security inherent with that protocol. We no longer recommend the Trusted Applications and Basic authentication types.

Linking Confluence to other applications allows you to include information from those applications in pages or blogs that you create in Confluence. For example, you could link Confluence to JIRA Software and display issues on a Confluence page using the [JIRA Issues Macro](#).

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Application Links** in the left-hand panel. The Application Links configuration page appears and lists any links you already have set up.
3. Enter the URL of the application you want to link to, then click **Create new link**.
 - If you check **The servers have the same set of users...** then user impersonation with 2-Legged OAuth authentication will be configured for this link. You can change this later if necessary.
 - If you are *not* an admin on both servers you won't be able to set up a 2-way (reciprocal) application link. If you want to go ahead and create a 1-way link anyway, clear the **I am an administrator on both instances** checkbox.
4. Use the wizard to finish configuring the link. If the application you are linking to does not have the Application Links plugin, you must supply additional information to set up a link with OAuth authentication.

When you complete the wizard, the Application Links plugin will create the link between your applications using the most secure authentication method that is supported between the two applications. See the [Application Links User Guide](#) for more information.

The new link will appear on the "Configure Application Links" page, where you can:

- Edit the settings of the application link (for example, to [change the authentication type](#) of the link) using the **Edit** link.
- Specify the default instance if you have multiple links to the same type of application (for example, to multiple JIRA servers) using the **Make Primary** link. See [Making a primary link for links to the same application type](#) for more information.

Having trouble integrating your Atlassian products with Application Links?

We've developed a [guide to troubleshooting Application Links](#), to help you out. Take a look at it if you need a hand getting around any errors or roadblocks with setting up Application Links.

Configuring Workbox Notifications

You can view and manage in-app notifications and tasks in your Confluence workbox. In addition, you can receive notifications from JIRA applications and other Confluence servers in your Confluence workbox. To make this possible, your Confluence server must be linked to the other server(s) via [application links](#).

Possible configurations:

- Your Confluence server provides in-app notifications and displays them in its own workbox. There are two sub-configurations here:
 - This Confluence server is the only server involved.
 - Alternatively, this Confluence server displays its own in-app notifications, and also displays notifications from JIRA and/or other Confluence servers.
- Your Confluence server does not provide or display in-app notifications.
- Your Confluence server sends in-app notifications to another Confluence server.

Notes:

- **Workbox includes notifications and tasks:** When you enable in-app notifications, personal tasks are also enabled in the workbox. When you disable in-app notifications, the workbox no longer appears and personal tasks are therefore not available on this server.

On this page:

- [Which notifications are included?](#)
- [Enabling Confluence workbox and in-app notifications](#)
- [Configuring the polling intervals](#)
- [Including notifications from JIRA](#)
- [Stopping JIRA applications from sending notifications to Confluence](#)
- [Including notifications from another Confluence server](#)
- [Sending Confluence notifications to another Confluence server](#)
- [Disabling workbox and in-app notifications in Confluence](#)

Which notifications are included?

The workbox displays a notification when someone does one of the following in Confluence:

- [Shares](#) a page or blog post with you.
- [Mentions](#) you in a page, blog post, comment or task.
- Comments on a page or blog post that you are [watching](#).
- [Likes](#) a page or blog post that you are watching.

The workbox does **not** show notifications triggered because you are watching a space. Only watches on pages and blog posts are relevant here.

The notification in your workbox appears as 'read' if you have already viewed the page or blog post.

If your Confluence site is linked to a JIRA application, you will also see the following JIRA notifications in your workbox:

- Comments on issues that you are watching.
- Mentions.
- Shares of issues, filters and searches.

Enabling Confluence workbox and in-app notifications

Confluence workbox and in-app notifications are disabled by default.

To enable workbox and in-app notifications:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **In-app Notifications** in the left-hand panel

3. Choose **displays in-app notifications** (or **displays in-app notifications from other servers**)

The workbox icon



will appear in the Confluence top menu bar and will be visible to all users.

Screenshot: Simple configuration with Confluence workbox and in-app notifications enabled for this server only

This Confluence server:

☒ **displays in-app notifications**
In-app notifications are displayed for this Confluence instance.

Active polling interval seconds
Time to wait before checking for new notifications on the page the user is currently viewing.

Inactive polling interval seconds
Time to wait before checking for new notifications when the user isn't focused on a page.

☐ **does not provide in-app notifications**
In-app notifications are disabled on this server.

Configuring the polling intervals

The polling intervals are used by the Confluence server that displays in-app notifications and tasks in its workbox.

Option	Description
Active polling interval	This is the number of seconds that Confluence will wait before checking (polling) for new notifications relevant to the page that the user is currently viewing. This setting applies to the page open in the browser tab that currently has focus. It does not matter whether the user has the workbox open or not.
Inactive polling interval	<p>This is the number of seconds that Confluence will wait before checking (polling) for new notifications relevant to all pages that are not currently in focus. These pages may be on the Confluence server that displays the workbox, or on other Confluence or JIRA servers that send their notifications to this server.</p> <p>This setting defines an upper limit. For inactive pages, Confluence starts with a polling interval equal to the active polling interval, then gradually increases the interval between polls until it reaches the limit defined here.</p>

Including notifications from JIRA

If your Confluence site is connected to a JIRA application, you can include notifications from your JIRA application, for example JIRA Software or JIRA Service Desk .

To include notifications from a JIRA application:

Your JIRA application and Confluence must be connected via an application link to do this. See [Linking to Another Application](#).

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **In-app Notifications** in the left-hand panel of the Confluence administration console.
3. Choose **displays in-app notifications from other servers**.

- Your JIRA application will appear in the list of linked applications below this option.
- People will see JIRA notifications in their workbox, as described in [Workbox Notifications](#).

Notes:

- JIRA sends its notifications to the Confluence server that is configured as the **primary** application link.
- Your JIRA server must be running **JIRA 5.2 or later**.
- The following plugins must be present and enabled in JIRA. The plugins are shipped with JIRA 5.2 and later:
 - 'Workbox – Common Plugin'
 - 'Workbox – JIRA Provider Plugin'
- You do not need to configure JIRA. The plugins are enabled by default in JIRA, and JIRA will automatically send notifications to Confluence.
- Confluence can display notifications from more than one server.

Screenshot: This Confluence server displays in-app notifications from itself and from JIRA

This Confluence server:

☒ displays in-app notifications from other servers

Linked applications send notifications to this server. These are:

Application	URL
Your Company JIRA	http://pyko:2991/jira-client

[Modify linked applications](#)

Active polling interval seconds
Time to wait before checking for new notifications on the page the user is currently viewing.

Inactive polling interval seconds
Time to wait before checking for new notifications when the user isn't focused on a page.

☐ does not provide in-app notifications
In-app notifications are disabled on this server.

Stopping JIRA applications from sending notifications to Confluence

You may wish to configure Confluence to display its own notifications in its workbox, but prevent notifications from JIRA applications from appearing in the workbox, even when JIRA applications and Confluence are linked via application links.

The JIRA administration interface does not offer a way of disabling notifications sent to Confluence.

To stop JIRA applications from sending notifications to Confluence: Disable the following plugins in JIRA. (See the [Universal Plugin Manager guide to disabling plugins](#).)

- 'Workbox – Common Plugin'
- 'Workbox – JIRA Provider Plugin'

Including notifications from another Confluence server

Confluence workbox can include notifications from another Confluence server.

Let's assume that you have two Confluence servers, *ConfluenceChatty* and *ConfluenceQuiet*. Let's also assume that you want *ConfluenceChatty* to display a workbox, and to include notifications from *ConfluenceQuiet*.

To include notifications from other Confluence servers:

1. Connect *ConfluenceChatty* and *ConfluenceQuiet* via application links. In *ConfluenceChatty*:
 - Choose the **cog icon**



- , then choose **General Configuration** under Confluence Administration
 - Choose **Application Links** in the left-hand panel.
 - Set up the link as described in [Linking to Another Application](#).
2. Configure the notification settings in *ConfluenceChatty*:
 - Choose **In-app Notifications** in the left-hand panel of the Confluence administration console.
 - Choose **displays in-app notifications from other servers**.
 3. Configure the notification settings in *ConfluenceQuiet*:
 - Choose **In-app Notifications** in the left-hand panel of the Confluence administration console.
 - Choose **sends in-app notifications to another server**.
 - Select the Confluence server that will display the workbox – in our example, this is *Confluence Chatty*. (The entry for *ConfluenceChatty* will appear here only if you have already configured *ConfluenceChatty* to display in-app notifications.)

Notes:

- Your Confluence servers must be running **Confluence 4.3.3 or later**.
- Confluence can display notifications from more than one server.
- Confluence can send notifications to only one server.
- Only one of the linked Confluence servers can display the in-app notifications.

Screenshot: This Confluence server displays in-app notifications from itself, from JIRA, and from another Confluence server

This Confluence server:

☒ displays in-app notifications from other servers

Linked applications send notifications to this server. These are:

Application	URL
Your Company JIRA	http://pyko:2991/jira-client
pyko	http://pyko:1991/confluence-client

[Modify linked applications](#)

Active polling interval seconds
Time to wait before checking for new notifications on the page the user is currently viewing.

Inactive polling interval seconds
Time to wait before checking for new notifications when the user isn't focused on a page.

☐ sends in-app notifications to another server
Notifications from this Confluence server are sent to the specified server.

☐ does not provide in-app notifications
In-app notifications are disabled on this server.

Sending Confluence notifications to another Confluence server

You can configure Confluence to send all notifications to a different Confluence server. In this case, the current Confluence server will not display the workbox.

To send notifications to another Confluence server: Follow the instructions in our example for *ConfluenceQuiet* [above](#).

Screenshot: This Confluence server sends its in-app notifications to another Confluence server

This Confluence server:

☐ displays in-app notifications from other servers
Linked applications send notifications to this server.

☒ sends in-app notifications to another server
Notifications from this Confluence server are sent to the specified server.

Application	URL
<input checked="" type="radio"/> pyko	http://pyko:1991/confluence-client

☐ does not provide in-app notifications
In-app notifications are disabled on this server.

Disabling workbox and in-app notifications in Confluence

If you choose **does not provide in-app notifications**:

- The Confluence workbox icon will no longer be visible and people will be unable to access their workboxes on this server.
- This Confluence server will no longer send notifications to its workbox, and will not send notifications to any other Confluence server.

Integrating JIRA and Confluence

JIRA applications and Confluence complement each other. Collect your team's thoughts, plans and knowledge in Confluence, track your issues in your JIRA application, and let the two applications work together to help you get your job done.

Here's some ways you can get JIRA and Confluence working together.

On this page:

- [Installing JIRA and Confluence together](#)
- [Use JIRA and Confluence together](#)
- [Delegate user management to JIRA](#)
- [Connect JIRA and Confluence with an application link](#)

Installing JIRA and Confluence together

We recommend running JIRA and Confluence in separate stand-alone instances behind an Apache Web Server. The following documentation will guide you through the installation processes:

- [Installing Confluence](#)
- [Installing JIRA applications](#)
- [Running Confluence behind Apache](#)
- [Integrating JIRA with Apache](#)

We don't support deploying Confluence and any other application (including JIRA) in the same Tomcat container. See [Can Multiple Atlassian Products Be Deployed in a Single Tomcat Container?](#) for more information.

Use JIRA and Confluence together

This is the fun stuff. Check out [Use JIRA applications and Confluence together](#) to find out about all the integration points, great time saving features, and to check exactly which JIRA application and version you'll need.

Delegate user management to JIRA

If you already have a JIRA application you can choose to delegate user management to JIRA, and manage

all your users in one place. You can control which JIRA groups also have permissions to use Confluence. Your license tiers for each application do not need to be the same.

See [Configuring JIRA Integration in the Setup Wizard](#) to delegate user management to JIRA when installing Confluence for the first time.

See [Connecting to Crowd or JIRA for User Management](#) to delegate user management to JIRA for an existing Confluence site.

Connect JIRA and Confluence with an application link

See [Linking to Another Application](#) to find out how to connect Confluence to your JIRA application using an application link. This only needs to be done once.

If you delegated user management to JIRA as part of Confluence's setup process, an application link to JIRA (using Trusted Applications authentication) will be all set up and ready to go.

You can choose to change the authentication to OAuth for additional security. See [Configuring authentication for an application link](#) to find out about changing the authentication to OAuth for additional security.

Having trouble integrating your Atlassian products with Application Links?

We've developed a [guide to troubleshooting Application Links](#), to help you out. Take a look at it if you need a hand getting around any errors or roadblocks with setting up Application Links.

Registering External Gadgets

You can register gadgets from external sites (such as [JIRA](#) applications), so the gadgets appear in the [macro browser](#) and people can add them to Confluence pages using the [gadget macro](#).

There's two ways to register external gadgets:

- **Subscribe to all of the external application's gadgets:** You can add all the gadgets from your [JIRA](#) application, [Bamboo](#), [FishEye](#) or [Crucible](#) site – or from another Confluence site – to your Confluence gadget directory. People can then pick and choose the gadgets to add to their Confluence pages.
- **Register the external gadgets one by one:** If you cannot subscribe to an application's gadgets, you will need to add the gadgets one by one. This is necessary for applications and websites that do not support gadget subscription, and for applications where you cannot establish a trusted relationship via Application Links.

Both methods are described below. First, consider whether you need to set up a trust relationship between Confluence and the other application.

Setting up a trust relationship with the other application

In addition to registering the external gadgets, we recommend that you set up an OAuth or Trusted Application relationship between the application that serves the gadget (the service provider) and Confluence (the consumer). The trust relationship is required for gadgets that access restricted data from the external web application.

See how to configure OAuth or Trusted Applications Authentication, using [Application Links](#).

On this page:

- [Setting up a trust relationship with the other application](#)
- [Subscribing to all of the application's gadgets](#)
- [Registering individual gadgets](#)
- [Removing access to external gadgets](#)

Related pages:

- [Configuring the Whitelist](#)
- [The big list of Atlassian gadgets](#)
- [Linking to Another Application](#)


If the external web application provides anonymous access to all the data you need in the gadgets, then you do not need a trust relationship.

For example, if your gadgets will retrieve data from JIRA and your JIRA server includes projects and issues that are restricted to logged-in users, then you will need a trust relationship between Confluence and JIRA. If you do not set up the trust relationship, then the gadgets will show only the information that JIRA makes visible to anonymous users.

Subscribing to all of the application's gadgets

You can add all the gadgets from your JIRA, Bamboo, FishEye or Crucible site – or from another Confluence site – to your Confluence gadget directory. People can then pick and choose the gadgets to add to their Confluence pages.

To subscribe to another site's gadgets:

1. Go to  **> General Configuration > External Gadgets**
2. Choose the **Gadget Feeds** tab.
3. Enter the base URL of the application you want to subscribe to, for example, <http://example.com/jira> or <http://example.com/confluence>.
4. Choose **Add**. Confluence will convert the URL to a gadget feed and place it in the list of 'Added Gadget Feeds'.

Screenshot: Subscribing to a gadget feed

External Gadgets

Only add gadgets that you trust! Gadgets can allow unwanted or malicious code onto your web page.

Gadget Specifications
Gadget Feeds
Gadget Whitelist

Add a new Gadget Feed

Gadget Feed URL

Add

Added Gadget Feeds

Gadget Feed URL	Actions
https://pug.jira.com/rest/gadgets/1.0/g/feed	<div>Delete</div>
https://pug.jira.com/rest/gadgets/1.0/g/feed/	<div>Delete</div>

Registering individual gadgets

If you cannot subscribe to an application's gadgets, you will need to register the gadgets one by one. This is necessary for applications and websites that do not support gadget subscription, and for applications where you cannot establish a trusted relationship via Application Links.

First you will need to get the gadget URL and copy it to your clipboard.

Getting a gadget's URL from an Atlassian application

If your application is another Atlassian application:

A gadget's URL points to the gadget's XML specification file. In general, a gadget's URL looks something like

this:

```
http://example.com/my-gadget-location/my-gadget.xml
```

If the gadget is supplied by a plugin, the URL will have this format:

```
http://my-app.my-server.com:port/rest/gadgets/1.0/g/my-plugin.key:my-gadget/my-path/my-gadget.xml
```

For example:

```
http://mycompany.com/jira/rest/gadgets/1.0/g/com.atlassian.streams.streams-jira-plugin:activitystream-gadget/gadgets/activitystream-gadget.xml
```

To find a gadget's URL in JIRA:

- Go to your dashboard by clicking the **Dashboards** link at the top left of the screen.
- Click **Add Gadget** to see the list of gadgets in the directory.
- Find the gadget you want, using one or more of the following tools:
 - Use the scroll bar on the right to move up and down the list of gadgets.
 - Select a category in the left-hand panel to display only gadgets in that category.
 - Start typing a key word for your gadget in the **Search** textbox. The list of gadgets will change as you type, showing only gadgets that match your search term.
- Right-click the **Gadget URL** link for that gadget and copy the gadget's URL into your clipboard.

To find a gadget's URL in Confluence:

- Choose **Help > Confluence Gadgets** to see the list of available Confluence gadgets.
- Find the gadget you want.
- Right-click the **Gadget URL** link for that gadget and copy the gadget's URL into your clipboard.


Getting a gadget's URL from another application

If the gadget comes from a non-Atlassian web application or web site, please consult the relevant documentation for that application to get the gadget URL.

Registering the gadget for use in Confluence

Now that you have the gadget's URL, you can register it in Confluence, so that people can add it to their pages. You need system administrator permissions to register a gadget.

To register the gadget in Confluence:

1. Go to  **> General Configuration > External Gadgets**
2. Paste your gadget's URL into the **Gadget Specification URL** field in the 'Add a new Gadget' section.
3. Choose **Add**. Your gadget will be shown in the list of registered gadgets below and it will also become available in the [macro browser](#).

Screenshot: Registering external gadgets one by one

External Gadgets

Only add gadgets that you trust! Gadgets can allow unwanted or malicious code onto your web page.

Gadget Specifications Gadget Feeds Gadget Whitelist

You can add gadgets from Atlassian applications such as Confluence, JIRA and others. You can also add gadgets from other websites such as iGoogle. Many public gadgets will work on a Confluence page. Some gadgets may rely on specific OpenSocial features that will not work properly in Confluence.

If you are adding gadgets from another Atlassian Application you need to either setup the other application to trust Confluence using Trusted Applications (make sure you add the '/rest' path to the allowed URL paths) or add Confluence as an OAuth Consumer to the other application.

A gadget's URL looks something like this: <http://example.com/my-gadget-location/my-gadget.xml>

Add a new Gadget

Gadget Specification URL

Add

Added Gadgets

Gadget Specification URL	Actions
https://pug.jira.com/rest/gadgets/1.0/g/com.atlassian.streams.streams-jira-plugin:activitystream-gadget/gadgets/activitystream-gadget.xml	Delete
https://pug.jira.com/rest/gadgets/1.0/g/com.atlassian.jira.gadgets:created-vs-resolved-issues-chart-gadget/gadgets/createdvsresolved-gadget.xml	Delete

Removing access to external gadgets

To remove a single gadget from Confluence, click the **Delete** button next to the gadget URL.

If you have subscribed to an application's gadgets, you will need to remove the entire subscription. You cannot unregister a single gadget. Click the **Delete** button next to the gadget feed URL.

The gadget(s) will no longer be available in the macro browser, and people will not be able to add them using the Gadget macro. Any pages that already use the gadget will show a broken gadget link.

Configuring the Office Connector

The Office Connector allows Confluence users to view, edit and import content from Microsoft Office and Open Office files attached to a page.

The Office Connector add-on is bundled with Confluence, but a [System Administrator](#) can enable or disable parts of the Office Connector and can configure options.

Enabling and disabling the Office Connector

If you want to limit access to all or part of the Office Connector you can disable the add-on, or some modules in the add-on.

To enable or disable the Office Connector modules:

1. Go to

On this page:

- [Enabling and disabling the Office Connector](#)
- [Configuring the Office Connector Options](#)

Related pages:

- [Office Connector Prerequisites](#)
- [Office Connector Limitations and Known Issues](#)



> **Add-ons**

2. Choose **System** from the filter drop down and then search for **Office Connector**
3. Expand the Office Connector add-on listing. From here you can:
 - Choose **Configure** to specify preferences for the Office Connector (this opens the configuration screen described [below](#))
 - Click **Disable** to disable all modules of the add-on
 - Expand the **modules** list to enable or disable selected Office Connector modules

Note: only some Office Connector modules can be disabled. Modules that are integral to the operation of the add-on cannot be disabled, and do not have an **Enable** or **Disable** button. Modules that can be disabled include the button and provide a brief, on-screen description of the module.

Configuring the Office Connector Options

Users with System Administrator permissions can configure the behaviour of the Office Connector.

To set the configuration options for the Office Connector:

1. Go to



> **General Configuration** > **Office Connector**

Screenshot: Configuring the Office Connector options

Configure Office Connector plugin

User Interface

Edit in word button location: ☒ Page action icon
☐ View page tab

Importing From Word

Warnings: ☐ Show a warning before allowing a user to perform an import

Advanced Formatting Options: ☐ Use the footnote macro for Word footnotes

Maximum imported image size: height (pixels)
 width (pixels)

System Resource Usage

Temporary storage for viewfile macro: ☒ Confluence home directory: C:\Program Files\Atlassian\Application Data\Confluence\viewfile
☐ No directory specified in the properties file. Please set a correct value in resources/directories.properties.

Maximum file space for cache(MB):


Number of Conversion Queues: [Manage Queues](#)

Security

Authentication: ☐ Allow authentication tokens in the URL path

2. Set the configuration options as described in the table below

Option	Default Value	Description
Edit in word button location	Page action icon	Where the button for editing the content in Word is located. You can configure action icon or from the view page tab. Note: This setting has no effect in modern Confluence themes, including the C documentation theme.
Warnings: Show a warning before allowing a user to perform an import	Disabled	If this option is enabled, the user will receive a warning when importing a Word user when they are about to overwrite existing content.

Advanced Formatting Options: Use the footnote macro for Word footnotes	Disabled	<p> Note: This feature requires a third party add-on that is not supported for Confluence 5.9.</p> <p>If this option is enabled, a Confluence page created from an imported Word document will use the Footnote macro to render any footnotes contained in the document. Note that you need to add the Footnote add-on onto your Confluence site. For more information about this add-on and how to install it, see the Footnote add-on.</p>
Authentication: Allow authentication tokens in the URL path	Disabled	<p>If this option is enabled, the Office Connector will use authentication tokens in the URL path. This needs to be enabled to edit Office 2013 documents.</p>
Temporary storage for viewfile macro	The Confluence Home directory.	<p>The {viewfile} macro will cache data temporarily. This option allows you to set the following settings:</p> <ul style="list-style-type: none"> • Confluence home directory – The temporary file will be stored in your Confluence home directory. • A directory specified in the <code>directories.properties</code> file – You can specify a directory in the Office Connector's <code>directories.properties</code> file: <ol style="list-style-type: none"> 1. Locate the <code>OfficeConnector-x.xx.jar</code> file (where x.xx is the version number) in the Confluence Home directory and copy it to a temporary location 2. Unzip the JAR file and find the <code>resources/directories.properties</code> file. Edit the file like this: <pre>#Complete the following line to set a custom cache directory #If resetting to blank, don't delete anything before the '=' com.benryan.confluence.word.edit.cacheDir=</pre> <ol style="list-style-type: none"> 3. Edit the last line, adding the path to your required temporary location directory. For example: <ul style="list-style-type: none"> • On Windows: <pre>com.benryan.confluence.word.edit.cacheDir=c:\temp\</pre> • On Linux: <pre>com.benryan.confluence.word.edit.cacheDir=/home/temp/</pre> 4. Save the file, recreate the JAR and put it back in your Confluence Home directory.
Maximum file space for cache (MB)	500	This is the maximum size of the cache used by the {viewfile} macro. (See above for more information.)
Number of Conversion Queues	6	<p>This is the maximum number of threads used to convert PowerPoint, Excel files. This setting is used to manage Confluence performance, by limiting the number of threads. The default value of 6 does not consume too many resources.</p> <p>Click Manage Queues to view attachments that are still pending conversion.</p>

Managing your Confluence License

Your license entitles you to run Confluence and be

eligible for support and upgrades for a specified period. It also defines the number of users who are entitled to use Confluence.

To quickly check the status of your license you can go to




> **General Configuration > Support Tools.**

You'll need need Confluence Administrator or System Administrator permissions to view and edit your license.

Viewing your license details

To view your Confluence license:

1. Go to  > **General Configuration**.
2. Choose **License Details** in the left-hand panel.

The License Details page tells you:


- The type of license (for example: Commercial, Academic, Community, or Evaluation).
- Number of users you are licensed for, and how many are currently in use.
- Your license expiry date, for support and upgrade eligibility.
- Your server ID which is generated when you install Confluence for the first time and remains the same for the life of the installation (including after upgrades or changes to your license).
- Your support entitlement number (SEN).

Updating your license

Screenshot: License details

If you change your license (for example to a license with more users), or migrate from Confluence Cloud and you will need to update your license.

To update your Confluence license:

1. Go to  > **General Configuration > License Details**
2. Enter your new license in the **License** field.
3. Choose **Save**.

Understanding the user count for your license

The number of registered users allowed on your Confluence site may be limited, depending on your license type.

The License Details page will indicate the number of users currently signed up (your registered user count). It:

- includes only users who have the 'can use' [global permissions](#) for the Confluence site.
- does not include anonymous users, who may access your Confluence site if you have [allowed anonymous access](#).
- does not include [deactivated users](#).

On this page:

- [Viewing your license details](#)
- [Updating your license](#)
- [Understanding the user count for your license](#)
- [Exceeding your licensed user count](#)
- [Reducing your user count](#)
- [Downgrading your license](#)
- [Finding your Support Entitlement Number \(SEN\)](#)

Related pages:

- [Upgrading Beyond Current Licensed Period](#)
- [Confluence Installation and Upgrade Guide](#)
- [Confluence Administrator's Guide](#)

License Details

This page shows your current licensing information.

You can use the form below to update the license Confluence is running with.

Organisation	Atlassian Docs
Date Purchased	Dec 08, 2013
License Type	Confluence: Developer
Licensed Users	25 (5 signed up currently) Refresh
Support Period	Your commercial Confluence support and updates are available until Dec 08, 2014 .
Support Entitlement Number	SEN-L2180
Server ID	BC1S-1234-8T5O-ABCD (Atlassian sales or support may ask you to provide this ID)
License	<div></div> Save

Exceeding your licensed user count

If you exceed the number of users included in your license, your Confluence instance will become read-only, that means no users will be able to create or edit content until you reduce the number of users.

Reducing your user count

You can reduce your user count by removing or deactivating users who do not require access to Confluence. See [Delete or Disable Users](#).

If you have [connected Confluence to an LDAP directory](#), you may want configure Confluence to only synchronise a subset of users from LDAP rather than all users. See [How to change the number of users synchronized from LDAP to Confluence](#) in the Knowledge Base. This can be a complicated process and we recommend that you only use this method if necessary.


Downgrading your license

If you decide to downgrade your Confluence license to pay for fewer users you need to ensure that the number of users currently signed up (as shown on the License Details page) is lower than the number allowed by your new license before you apply the new license.

If you have more users than your new license allows you will need to reduce your user count before applying the new license.

Finding your Support Entitlement Number (SEN)

You can find your Support Entitlement Number (SEN) in three places:

- In Confluence - go to  **> General Configuration > License Details**)
- At my.atlassian.com
- On your Atlassian invoice.

See [Finding Your Support Entitlement Number](#) for more general information about how Atlassian Support uses this number.

Managing Confluence Data

This page is an overview of recommended techniques for managing the data on your Confluence site. This is of interest to Confluence administrators – people with System Administrator or Confluence Administrator permissions.

- [Database Configuration](#)
- [Site Backup and Restore](#)
- [Attachment Storage Configuration](#)
- [Confluence Data Model](#)
- [Finding Unused Spaces](#)
- [Data Import and Export](#)
- [Import a Text File](#)

Related pages:

- [Managing Add-ons or Plugins](#)
- [Integrating Confluence with Other Applications](#)
- [Getting Started as Confluence Administrator](#)
- [Confluence Administrator's Guide](#)

Database Configuration

This document provides information on connecting Confluence to an external database.

Choosing an external database

Note: Take time to choose your database wisely. The XML backup built into Confluence is not suited for migration or backup of large data sets. If you need to migrate later, you'll need to use a third party database migration tool.

Below is more information on selecting and migrating to an external database:

- [Migrating to a Different Database](#)
- [Supported Databases](#)
- [Database Troubleshooting](#)

About the embedded H2 database

Your Confluence installation includes an embedded H2 database, to enable you to try Confluence without setting up an external database. The embedded H2 database is **only** supported while you are evaluating Confluence. You must migrate to a [supported](#) external database before using Confluence as a production system.

To find out if you are still using the embedded database, go to



> **General Configuration > Support Tools.**

Database setup

Here are the setup instructions for the supported databases:

- [Database Setup for Oracle](#)
- [Database Setup For MySQL](#)
- [Database Setup for PostgreSQL](#)
- [Database Setup for SQL Server](#)

Database troubleshooting

For solving database-related problems:

- [Troubleshooting External Database Connections](#)
- [How to Interpret DB2 Error Codes](#)
- [Database Troubleshooting](#)

Obtain technical support from [Troubleshooting Problems and Requesting Technical Support](#).

Notes

Issue [CONF-12599](#) requests a more robust strategy for migrating large Confluence sites.

Database JDBC Drivers

This page provides the download links for the JDBC drivers for all databases currently supported for Confluence. You will need to make the driver available to your application server, as described in the appropriate [setup guide](#).

Note: We bundle some JDBC drivers with Confluence, as shown below. If you are using a direct JDBC connection, you do not need to download or install the drivers that are bundled. If you are connecting via a

On this page:

- [Choosing an external database](#)
- [About the embedded H2 database](#)
- [Database setup](#)
- [Database troubleshooting](#)
- [Notes](#)

Related pages:

- [Database JDBC Drivers](#)
- [Supported Platforms](#)
- [Embedded H2 Database](#)
- [Managing Confluence Data](#)

datasource, or if you are using a database whose driver is not bundled, you will need to download and install the drivers manually.

Related pages:

- [Database Configuration](#)
- [Supported Platforms](#)

Database	JDBC driver bundled with Confluence?	JDBC drivers	Notes	More information
PostgreSQL		9.2-1004 JDBC 4 driver download	The JDBC 4 driver will work under the 1.7 JVM. If you want to use other drivers, you can download it from the PostgreSQL website . However, we recommend that you use the bundled JDBC 4 driver.	Database Setup for PostgreSQL
Microsoft SQL Server		JTDS 1.2.2 driver download	The above version is the version bundled with Confluence. All our testing is done on that version. We do not know of any issues with later versions, so you are free to use them if you have tested them and find there are no issues in your environment. However, later versions are technically not supported. That means that if you do run into any problems, Atlassian Support may require you to move back to the above fully-tested version for troubleshooting.	Database setup for Microsoft SQL Server
MySQL		Connector/J 5.1.30 driver download	Note: In Confluence 5.1 and later, the MySQL drivers are no longer included in the Confluence distribution. For more information please refer to the Confluence 5.1 Upgrade Notes .	Database setup for MySQL
Oracle		JDBC driver downloads	<ul style="list-style-type: none"> • For Oracle 11.1, use the 11.1.0.7.0 driver (ojdbc6.jar). • For Oracle 11.2, use the 11.2.0.4.0 driver (ojdbc6.jar). • For Oracle 12c use the 12.1.0.1 driver (ojdbc7.jar) <p>We recommend using the thin drivers only.</p> <p>See the Oracle JDBC driver FAQ.</p>	Database setup for Oracle

Database Setup for Oracle

This page provides instructions for configuring Confluence to use an Oracle database. The setup process involves configuration of your Oracle server and your Confluence site.

Step 1. Check the prerequisites

Check the following before you start:

- Make sure your version of Oracle is

supported. See [Supported Platforms](#). If your version of Oracle is not supported, please upgrade to a supported version before installing Confluence.

- If you have been evaluating Confluence and wish to transfer your data to a new database, consult the following guide first: [Migrating to Another Database](#).
- If you are migrating from another database, consult the following guide first: [Migrating to Another Database](#).
- **Note:** This database can only be set up by an Oracle database administrator (DBA). Oracle is difficult to set up. If you are not a DBA, and you do not have access to an experienced Oracle DBA, we recommend that you choose an alternative database. For a list of supported databases, see [Supported Platforms](#). If you are evaluating Confluence, we recommend that you start with an alternative database and only consider migrating to Oracle after approval from your DBA. If you request Atlassian's technical support for difficulties with Oracle setup, we will assume you have the high level of skill required for an Oracle setup.

On this page:

- [Step 1. Check the prerequisites](#)
- [Step 2. Install your Oracle server](#)
- [Step 3. Set up your Oracle user with schema-creation privileges](#)
- [Step 4. Install Confluence](#)
- [Step 5. Determine your JDBC URL](#)
- [Step 6. Download and install the Oracle database driver](#)
- [Step 7. Set up your database connection in the Confluence Setup Wizard](#)
- [Troubleshooting](#)

Related pages:

- [Database Configuration](#)
- [Known Issues for Oracle](#)
- [Confluence Installation and Upgrade Guide](#)

Step 2. Install your Oracle server

If you do not already have an operational Oracle database server, download the installation package from the [Oracle download page](#) and follow the instructions in the [Oracle documentation](#).

Then follow the steps below, to deploy Confluence to a schema in your Oracle server.

Note: Your database should be configured to use the same character encoding as Confluence. The recommended encoding is AL32UTF8 (the Oracle equivalent of Unicode UTF-8). See [Configuring Database Character Encoding](#).

Step 3. Set up your Oracle user with schema-creation privileges

In this step you will create a Confluence user in Oracle and grant the appropriate roles to the user, so that the user can set up a connection, can create objects in its own schema, and can configure the schema.

To create the user and assign its privileges:

1. Access the command line interface to Oracle via the 'sqlplus' command.

```
sqlplus user/password <as sysdba|as sysoper>
```

You must add the 'as sysdba' or 'as sysoper' option if you are logging in with the user 'sys'. This determines which sys role you are using.

Once logged in, you can type arbitrary SQL commands.

2. Create a Confluence user (<user>) in Oracle, and grant the appropriate roles only to the user:
 - `connect` role is required to set up a connection.
 - `resource` role is required to allow the user to create objects in its own schema.
 - `Create table`, `sequence` and `trigger` are required to configure the schema.

```
create user <user> identified by <password> default tablespace
<tablespace_name> quota unlimited on <tablespace_name>;
grant connect to <user>;
grant resource to <user>;
grant create table to <user>;
grant create sequence to <user>;
grant create trigger to <user>;
```

Notes:

- Do not grant the user the `select any table` permission. That permission can cause problems with other schemas. See the bug report [CONF-3613](#).
 - When you create a user, specify the `tablespace` for the table objects as shown above.
3. Add a local `all_objects` view to the user's schema, to prevent a conflict that can occur when a table exists in another schema with the same name as one of the Confluence tables. This is a workaround for the bug [CONF-3613](#):

```
create view <user>.all_objects as
select *
from sys.all_objects
where owner = upper('<user>');
```

Step 4. Install Confluence

Install Confluence if you have not done so already. See the [Confluence Installation Guide](#). **Stop immediately after the installation, before opening the Confluence Setup Wizard in your browser**, and follow the steps below.

If you have already got part-way through the Confluence Setup Wizard, stop at the database setup step and follow the steps below. You will be able to restart the setup wizard at the same step later.

Step 5. Determine your JDBC URL

The JDBC thin driver for Oracle use three different styles of URL:

- New style:

New Style

```
jdbc:oracle:thin:@//[HOST][:PORT]/SERVICE
```

- Old style:

Old Style

```
jdbc:oracle:thin:@[HOST][:PORT]:SID
```

- 'tnsnames' style:

tnsnames

```
jdbc:oracle:thin:@(DESCRIPTION=
    (SDU=32768)
    (enable=broken)
    (LOAD_BALANCE=yes)
    (FAILOVER=yes)
    (ADDRESS=
        (PROTOCOL=TCP)
        (HOST=dbserver1.example.com)
        (PORT=1525))
    (ADDRESS=
        (PROTOCOL=TCP)
        (HOST=dbserver2.example.com)
        (PORT=1525))
    (CONNECT_DATA=
        (SERVICE_NAME=CONFDB)))
```

Notes:

- The `tnsnames` style is required for connecting to an Oracle RAC cluster. For easy reading, we have split the example above over multiple lines, but you should compact it into a single line. These values may need more analysis than documented here, so you should seek the assistance of an experienced DBA.
- If you use the new style URL, then `SERVICE` can be either an SID or Service Name.
- If you use the old style URL, then `SERVICE` can only be the SID.

To determine the host, port, service name, and/or SID, execute the following command as the user running Oracle. (By default, the user is "oracle"):

```
lsnrctl status
```

For reference, here is a sample output:

```

SNRCTL for Linux: Version 11.2.0.2.0 - Beta on 29-JUN-2012 15:20:59
Copyright (c) 1991, 2010, Oracle. All rights reserved.
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=EXTPROC_FOR_XE)))
STATUS of the LISTENER
-----
Alias                     LISTENER
Version                   TNSLSNR for Linux: Version 11.2.0.2.0 - Beta
Start Date                06-JUN-2012 08:36:34
Uptime                    23 days 6 hr. 44 min. 25 sec
Trace Level               off
Security                  ON: Local OS Authentication
SNMP                      OFF
Default Service           XE
Listener Parameter File   /u01/app/oracle/product/11.2.0/xe/network/admin/listener.ora
Listener Log File         /u01/app/oracle/diag/tnslsnr/<HOSTNAME>/listener/alert/log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC_FOR_XE)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=<HOSTNAME>)(PORT=1521)))

  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=<HOSTNAME>)(PORT=8080))(Presentation=HT
TP)(Session=RAW))
Services Summary...
Service "PLSExtProc" has 1 instance(s).
  Instance "PLSExtProc", status UNKNOWN, has 1 handler(s) for this service...
Service "XE" has 1 instance(s).
  Instance "XE", status READY, has 1 handler(s) for this service...
Service "XEXDB" has 1 instance(s).
  Instance "XE", status READY, has 1 handler(s) for this service...
The command completed successfully

```

Notes:

- The host and port are determined by the line containing `PROTOCOL=tcp`, **without** `Presentation=HT` `TTP`.
- Under `Services Summary`, each service which has an instance with `READY` status is a connectable service. The name following `Service` is a service name for connecting to the database name following `Instance` on the next line.
- The `SID` is the name of the database instance, as defined by the `$ORACLE_SID` variable when you have sourced the Oracle environment to your shell.

For example, assuming that you are running Confluence on the same server as the Oracle database, with the above `lsnrctl` status output, you would use one of the following URLs:

```

jdbc:oracle:thin:@//localhost:1521/XE
jdbc:oracle:thin:@localhost:1521:XE

```

The URL can be used in either a direct JDBC connection or using a Tomcat datasource.

For further information on Oracle JDBC URLs, see [the Oracle JDBC FAQ](#).

Step 6. Download and install the Oracle database driver

Decide whether you will set up a **direct JDBC connection** or a **datasource connection** to Oracle, to suit your environment. If unsure, choose direct JDBC.

To set up a direct JDBC connection:

If you plan to set up a direct JDBC connection to Oracle, you will need to copy the Oracle JDBC driver to your Confluence installation.

1. Download the latest compatible database driver. Links to the appropriate database drivers are available on this page: [Database JDBC Drivers](#).
2. Copy the driver JAR file to the `<Confluence installation>/confluence/WEB-INF/lib` folder in your new Confluence installation.

To set up a datasource connection:

If you plan to set up a datasource connection to Oracle, follow the steps described in [Configuring an Oracle Datasource in Apache Tomcat](#).

Step 7. Set up your database connection in the Confluence Setup Wizard

Start Confluence, and go to the Confluence Setup Wizard in your browser. Follow these steps to set up the new configuration:

1. Follow the initial steps in the [Confluence Setup Guide](#).
2. When prompted to choose an evaluation or production installation, choose **production installation**.
3. When prompted to choose an embedded or external database, select **Oracle xx** from the dropdown list, where 'xx' is your Oracle version, and choose **External Database**.
4. Choose either the direct JDBC or the datasource connection, to suit the choice you made earlier when setting up the Oracle database driver.
 - For the JDBC connection: Enter the **database URL** to match the JDBC URL you determined in the previous section. Enter the **user name** (for example, `confluenceuser`) and **password** you chose when adding the Confluence database to Oracle.
 - For a datasource connection: Set the JNDI name to `java:comp/env/jdbc/confluence`.
5. Wait a while, as Confluence will create the schema in Oracle.

Congratulations! Confluence is now using your Oracle database to store its data.

Troubleshooting

- The following page contains common issues encountered when setting up your Oracle database to work with Confluence: [Known Issues for Oracle](#).
- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If none of the above describes your issue, please create a support ticket at <http://support.atlassian.com> and be sure to include your logs (found in `<CONFLUENCE-INSTALLATION>/logs` and `<CONFLUENCE-HOME>/logs`).

Configuring an Oracle Datasource in Apache Tomcat

This page tells you how to set up an Oracle datasource connection for Confluence.

Step 1. Shut down Tomcat

1. Run `bin/shutdown.sh` or `bin/shutdown.bat` to bring Tomcat down while you are making these changes.
2. Make a backup of your `<CONFLUENCE_HOME>/confluence.cfg.xml` file and your `<CONFLUENCE_INSTALLATION>/conf/server.xml` file, so that you can easily revert if you have a problem.

Step 2. Install the Oracle database driver

1. Download the Oracle JDBC driver. Links are

On this page:

- [Step 1. Shut down Tomcat](#)
- [Step 2. Install the Oracle database driver](#)
- [Step 3. Configure Tomcat](#)
- [Step 4. Configure the Confluence web application](#)
- [Step 5. Restart Tomcat](#)

Related pages:

- [Database Setup For Oracle](#)
- [Confluence Installation and Upgrade Guide](#)

available on this page: [Database JDBC Drivers](#).

- Copy the driver JAR file into the `lib` folder of your Tomcat installation: `<TOMCAT-INSTALLATION>/lib`.

Step 3. Configure Tomcat

- If you are using the Confluence distribution, edit the `conf/server.xml` file in your Tomcat installation. If you are running your own Tomcat instance, edit the XML file where you declared the Confluence Context descriptor.
- Find the Context element in the Host element:

```
<Host name="localhost" debug="0" appBase="webapps" unpackWARs="true"
autoDeploy="false">
  <Context path="" docBase="../confluence" debug="0" reloadable="true">
    <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
    <Manager pathname="" />
  </Context>
</Host>
```

- Insert the DataSource Resource element into the Context element, directly after the opening `<Context.../>` line, before Manager, as shown here:

```
<Host name="localhost" debug="0" appBase="webapps" unpackWARs="true"
autoDeploy="false">
  <Context path="" docBase="../confluence" debug="0" reloadable="true">
    <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
    <!-- If you're using Confluence 5.7 or below; change maxTotal to
maxActive -->
    <Resource
      name="jdbc/confluence"
      auth="Container"
      type="javax.sql.DataSource"
      driverClassName="oracle.jdbc.OracleDriver"
      url="jdbc:oracle:thin:@hostname:port:sid"
      username="<username>"
      password="<password>"
      connectionProperties="SetBigStringTryClob=true"
      accessToUnderlyingConnectionAllowed="true"
      maxTotal="25"
      maxIdle="10"
      maxWaitMillis="10000"
    />
    <Manager pathname="" />
  </Context>
</Host>
```

- Change the username and password to match your Oracle login.
- Change the url to match the URL for your Oracle database. See [how to find your Oracle URL](#). For example:

```
jdbc:oracle:thin:@example.atlassian.com:1521:confluencedb
```

- If required, choose different maxTotal and maxIdle values. These define the number of database

connections that will be allowed at one time, and the number that will be kept open even when there is no database activity.

Step 4. Configure the Confluence web application

Configure Confluence to use this datasource:

1. Edit this file in your Confluence installation: `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`.
2. Insert the following element just before `</web-app>` near the end of the file:

```
<resource-ref>
  <description>Connection Pool</description>
  <res-ref-name>jdbc/confluence</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
```

Step 5. Restart Tomcat

Run `bin/startup.sh` or `bin/startup.bat` to start Tomcat with the new settings.

Database Setup for SQL Server

This page provides instructions for configuring Confluence to use the Microsoft SQL Server database.

Step 1. Check the prerequisites

Check the following before you start:

- Check that your version of SQL Server is supported. See [Supported Platforms](#). If your version is not supported, please upgrade to a supported version of SQL Server before installing Confluence.
- If you have been evaluating Confluence and wish to transfer your data to a new database, consult the following guide first: [Migrating to Another Database](#).
- If you are migrating from another database, consult the following guide first: [Migrating to Another Database](#).

Step 2. Install SQL Server

If you do not already have an operational SQL Server database, download the installation package from the [Microsoft SQL Server download page](#) and follow the instructions on [MSDN](#).

Step 3. Set up your SQL Server database and user

In this step you will create a database within SQL Server to hold your Confluence data, and a database user with authority to access that database. The database user should be in the `db_owner` role.

1. Identify which character encoding to use. To do this, check the encoding currently used by your [application server](#) and [Confluence](#). All three must use compatible encoding. For example, the default SQL Server encoding of UCS-2 is compatible with UTF-8.
2. Using your SQL administrator permissions, create a new database in SQL Server.
3. If you set your application server and Confluence to use an encoding incompatible with UCS-2, specify that character encoding for the database.
4. Set the default collation for the database to be `'SQL_Latin1_General_CP1_CS_AS'` (**case sensitive**).

On this page:

- [Step 1. Check the prerequisites](#)
- [Step 2. Install SQL Server](#)
- [Step 3. Set up your SQL Server database and user](#)
- [Step 4. Install Confluence and the SQL Server database driver](#)
- [Step 5. Set up your database connection in the Confluence Setup Wizard](#)
- [Troubleshooting](#)

Related pages:

- [Database Configuration](#)
- [Known issues for SQL Server](#)
- [Confluence Installation and Upgrade Guide](#)

You can do this by issuing the following SQL query:

```
ALTER DATABASE <database_name> COLLATE SQL_Latin1_General_CP1_CS_AS
```

Note: if you receive an error stating 'The database could not be exclusively locked to perform the operation', you may need to prevent other connections by setting the mode to single user for the transaction:

```
ALTER DATABASE <database_name> SET SINGLE_USER WITH ROLLBACK IMMEDIATE;  
<your ALTER DATABASE query>  
ALTER DATABASE <database_name> SET MULTI_USER;
```

5. Configure the database to use the isolation level, 'Read Committed with Row Versioning'. You can do this by issuing the following SQL query:

Determine if READ_COMMITTED_SNAPSHOT is enabled

```
SELECT is_read_committed_snapshot_on FROM  
sys.databases WHERE name= 'YourDatabase'
```

Return value:

1 = READ_COMMITTED_SNAPSHOT option is ON. Read operations under the read-committed isolation level are based on snapshot scans and do not acquire locks.

0 = READ_COMMITTED_SNAPSHOT option is OFF (default). Read operations under the read-committed isolation level use share locks.

```
ALTER DATABASE <database_name>  
SET READ_COMMITTED_SNAPSHOT ON  
WITH ROLLBACK IMMEDIATE;
```

6. Using your SQL administrator permissions, create a new SQL user account for Confluence (for example, `confluenceuser`). Give this user full create, read and write permissions for the database tables. Note that Confluence must be able to create its own schema.

Step 4. Install Confluence and the SQL Server database driver

Decide whether you will set up a **direct JDBC connection** or a **datasource connection** to SQL Server, to suit your environment. If unsure, choose direct JDBC.

Install Confluence if you have not done so already. See the [Confluence Installation Guide](#).

- If you plan to set up a **direct JDBC connection** to SQL Server, you can run the Confluence installation and move directly on to the Confluence Setup Wizard, as described [below](#). The SQL Server JDBC driver is bundled with Confluence, as documented on this page: [Database JDBC Drivers](#).
- If you plan to set up a **datasource connection** to SQL Server:
 - Stop immediately after the Confluence installation, before opening the Confluence Setup Wizard in your browser. If you have already got part-way through the Confluence Setup Wizard, stop at the database setup step. You will be able to restart the setup wizard at the same step later.
 - Follow the steps described in [Configuring a SQL Server Datasource in Apache Tomcat](#).

Step 5. Set up your database connection in the Confluence Setup Wizard

Start Confluence, and go to the Confluence Setup Wizard in your browser. Follow these steps to set up the new configuration:


1. Follow the initial steps in the [Confluence Setup Guide](#).
2. When prompted to choose an evaluation or production installation, choose **production installation**.
3. When prompted to choose an embedded or external database, select **Microsoft SQL Server** from the dropdown list and choose **External Database**.
4. Choose either the direct JDBC or the datasource connection, to suit the choice you made earlier when setting up the SQL Server database driver.
 - For the JDBC connection:

- When prompted for a **Driver Class Name**, enter the following:

```
net.sourceforge.jtds.jdbc.Driver
```

- When prompted for the **Database URL**, use this format:

```
jdbc:jtds:sqlserver://<server>:<port>/<database>
```

-  If MS SQL is clustered, use this format:

```
jdbc:jtds:sqlserver://<server>:<port>/<database>;instance=<instance>
```

- Enter the username (for example, `confluenceuser`) and password you chose earlier.
- For a datasource connection: Set the **Datasource Name** to `java:comp/env/jdbc/confluence`

Congratulations! Confluence is now using your SQL Server database to store its data.

Troubleshooting

- If you get the following error message, verify that you have given the `confluenceuser` user all the required database permissions when connecting from `localhost`.

```
Could not successfully test your database: : Server connection failure during transaction. Due to underlying exception: 'java.sql.SQLException: Access denied for user 'confluenceuser'@'localhost' (using password: YES)'
```

- The following page contains common issues encountered when setting up your SQL Server database to work with Confluence: [Known Issues for SQL Server](#).
- Additional ports may be required to be opened. See [this support document](#) from Microsoft about the ports required for SQL Server.
- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If none of the above describes your issue, please create a support ticket at <http://support.atlassian.com> and be sure to include your logs (found in `<CONFLUENCE-INSTALLATION>/logs` and `<CONFLUENCE-HOME>/logs`).

Configuring a SQL Server Datasource in Apache Tomcat

This page tells you how to set up a SQL Server

datasource connection for Confluence.

Step 1. Shut down Tomcat

1. Run `bin/shutdown.sh` or `bin/shutdown.bat` to bring Tomcat down while you are making these changes.
2. Make a backup of your `<CONFLUENCE_HOME>/confluence.cfg.xml` file and your `<CONFLUENCE_INSTALLATION>/conf/server.xml` file, so that you can easily revert if you have a problem.

Step 2. Install the SQL Server database driver

1. Download the SQL Server JDBC driver:
 - Links are available on this page: [Database JDBC Drivers](#).
 - Unpack the archive file you have downloaded, and find the JAR file called something like this: `jtds-x.x.x.jar`, where `x.x.x` is a version number.
2. Alternatively, you can get the driver from your Confluence installation: `/confluence/WEB-INF/lib/jtds-x.x.x.jar`.
3. Put the JAR file into the `lib` folder of your Tomcat installation: `<TOMCAT-INSTALLATION>/lib`.

Step 3. Configure Tomcat

1. Edit the `conf/server.xml` file in your Tomcat installation.
2. Find the following lines:

```
<Context path="" docBase="../confluence" debug="0" reloadable="true">
    <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
    Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
```

3. Insert the `DataSource` Resource element inside the `Context` element, directly after the opening `<Context.../>` line, before `Manager`:

```
<!-- If you're using Confluence 5.7 or below; change maxTotal to maxActive -->
<Resource name="jdbc/confluence" auth="Container"
type="javax.sql.DataSource"
    username="yourDatabaseUser"
    password="yourDatabasePassword"
    driverClassName="net.sourceforge.jtds.jdbc.Driver"
    url="jdbc:jtds:sqlserver://localhost:1433/yourDatabaseName"
    maxTotal="25"
    maxIdle="10"
    validationQuery="select 1" />
```

- Replace the `username` and `password` parameters with the correct values for your database
- In the `url` parameter, replace the word `'yourDatabaseName'` with the name of the database your Confluence data will be stored in.

Notes:

- If switching from a direct JDBC connection to datasource, you can find the above details in your `<CON`

On this page:

- [Step 1. Shut down Tomcat](#)
- [Step 2. Install the SQL Server database driver](#)
- [Step 3. Configure Tomcat](#)
- [Step 4. Configure the Confluence web application](#)
- [Step 5. Restart Tomcat](#)

Related pages:

- [Database Setup for SQL Server](#)
- [Confluence Home and other important directories](#)
- [Confluence Installation and Upgrade Guide](#)

FLUENCE_HOME>/confluence.cfg.xml file.

- Here are the configuration properties for Tomcat's standard data source resource factory (`org.apache.tomcat.dbcp.dbcp.BasicDataSourceFactory`):
 - `driverClassName` — Fully qualified Java class name of the JDBC driver to be used.
 - `maxTotal` — The maximum number of active instances that can be allocated from this pool at the same time.
 - `maxIdle` — The maximum number of connections that can sit idle in this pool at the same time.
 - `maxWaitMillis` — The maximum number of milliseconds that the pool will wait (when there are no available connections) for a connection to be returned before throwing an exception.
 - `password` — Database password to be passed to our JDBC driver.
 - `url` — Connection URL to be passed to our JDBC driver. (For backwards compatibility, the property `driverName` is also recognized.)
 - `user` — Database username to be passed to our JDBC driver.
 - `validationQuery` — SQL query that can be used by the pool to validate connections before they are returned to the application. If specified, this query MUST be an SQL SELECT statement that returns at least one row.
- Why is the `validationQuery` element needed? When a database server reboots, or there is a network failure, all the connections in the connection pool are broken and this normally requires a Application Server reboot. However, the Commons DBCP (Database Connection Pool) which is used by the Tomcat application server can validate connections before issuing them by running a simple SQL query, and if a broken connection is detected, a new one is created to replace it. To do this, you will need to set the "validationQuery" option on the database connection pool.

Step 4. Configure the Confluence web application

1. Edit this file in your Confluence installation: `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`.
2. Insert the following element just before `</web-app>` near the end of the file:

```
<resource-ref>
  <description>Connection Pool</description>
  <res-ref-name>jdbc/confluence</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
```

If you are changing an existing Confluence installation over to using a Tomcat datasource:

1. Edit the `<CONFLUENCE_HOME>/confluence.cfg.xml` file.
2. Delete any line that contains a property that begins with `hibernate`.
3. Insert the following at the start of the `<properties>` section.

```
<property name="hibernate.setup">true</property>
<property
name="hibernate.dialect">net.sf.hibernate.dialect.SQLServerIntJ dialect</pro
perty>
<property
name="hibernate.connection.datasource">java:comp/env/jdbc/confluence</prope
rty>
```

Step 5. Restart Tomcat



Run `bin/startup.sh` or `bin/startup.bat` to start Tomcat with the new settings.

Database Setup For MySQL

This page provides instructions for configuring Confluence to use the [MySQL database](#).

Step 1. Check the prerequisites

Check the following before you start:

- Check that your version of MySQL is supported. See [Supported Platforms](#).
 - ▼ [Note about MariaDB and Percona Server...](#)
We do not currently support MariaDB or Percona Server. Both are known to have issues with Confluence. See  and  for more information.
- If you have been evaluating Confluence using the embedded database and wish to transfer your data to a new database, consult the following guide first: [Migrating to Another Database](#).
- If you are migrating from another database, consult the following guide first: [Migrating to Another Database](#).

On this page:

- [Step 1. Check the prerequisites](#)
- [Step 2. Install MySQL Server](#)
- [Step 3. Configure MySQL Server](#)
- [Step 4. Set up your MySQL database and user](#)
- [Step 5. Install Confluence](#)
- [Step 6. Download and install the MySQL database driver](#)
- [Step 7. Check settings for internationalisation](#)
- [Step 8. Set up your database connection in the Confluence Setup Wizard](#)
- [Troubleshooting](#)

Related pages:

- [Configuring Database Character Encoding](#)
- [Database Configuration](#)
- [Known issues for MySQL](#)
- [Confluence Installation and Upgrade Guide](#)

Step 2. Install MySQL Server

If you do not already have an operational MySQL database server, install 'MySQL Community Edition'. Download the installation package from the [MySQL download page](#) and follow the instructions in the [MySQL documentation](#).

Step 3. Configure MySQL Server

In this step, you will configure your MySQL database server.

Note: *If you intend to connect Confluence to an existing MySQL database server, we strongly recommend that you reconfigure this database server by running through the configuration steps in the MySQL installation wizard as described below.*

To configure MySQL Server:

- Run the MySQL installation wizard:
 - If you are connecting Confluence to your existing MySQL server, choose **Reconfigure Instance**.
 - Choose **Advanced Configuration**.
 - Choose the **type of MySQL Server** that best suits your hardware requirements. This will affect the MySQL Server's usage of memory, disk and CPU resources. Refer to the [MySQL documentation](#) for further information.
 - Choose **Transactional Database Only** to ensure that your MySQL database will use **InnoDB** as its default storage engine.
*It is **highly recommended** that you only use the InnoDB storage engine with Confluence. Avoid using the MyISAM storage engine as this can lead to data corruption.*
 - Set the **InnoDB Tablespace** settings to your requirements. (The default settings are acceptable.)
 - Set the approximate **number of concurrent connections** permitted to suit your Confluence usage requirements. You can use one of the presets or enter a number manually. Refer to the [MySQL documentation](#) for further information.
 - For the **networking options**, ensure the **Enable TCP/IP Networking** and **Enable Strict Mode** options are selected (default). Refer to the MySQL documentation on setting the [networking](#) and [server SQL modes](#) for further information.
 - For the MySQL server's **default character set**, choose **Best Support For Multilingualism** (in other words, UTF-8). This will ensure Confluence's support for internationalisation. For more information, see [Configuring Database Character Encoding](#).

- i. For the Windows configuration option, choose whether or not to install the MySQL Server as a Windows service. If your hardware is going to be used as a dedicated MySQL Server, you may wish to choose the options to **Install As Windows Service** (and **Launch the MySQL Server automatically**). Refer to the [MySQL documentation](#) for further information.
Note: If you choose not to install the MySQL Server as a Windows Service, you will need to ensure that the database service has been started before running Confluence.
 - j. Select **Modify Security Settings** to enter and set your MySQL Server (root) access password.
2. Edit the `my.cnf` file (`my.ini` on Windows operating systems) in your MySQL server. Locate the `[mysqld]` section in the file, and add or modify the following parameters:
(Refer to [MySQL Option Files](#) for detailed instructions on editing `my.cnf` and `my.ini`.)
Locate the `[mysqld]` section in the file, and add or modify the following parameters:
- Specify the default character set to be UTF-8:

```
[mysqld]
...
character-set-server=utf8
collation-server=utf8_bin
...
```

- Set the default storage engine to InnoDB:

```
[mysqld]
...
default-storage-engine=INNODB
...
```

- Specify the value of `max_allowed_packet` to be at least 256M:

```
[mysqld]
...
max_allowed_packet=256M
...
```

- Specify the value of `innodb_log_file_size` to be at least 256M for MySQL 5.5 and below:

```
[mysqld]
...
innodb_log_file_size=256M
...
```

NB: This should be set to at least 2G for MySQL 5.6 and above.

- Ensure the `sql_mode` parameter does not specify `NO_AUTO_VALUE_ON_ZERO`

```
// remove this if it exists
sql_mode = NO_AUTO_VALUE_ON_ZERO
```

3. Restart your MySQL server for the changes to take effect:
- On Windows, use the Windows Services manager to restart the service.
 - On Linux:
 - Run one of the following commands, depending on your setup: `/etc/init.d/mysqld stop` or `/etc/init.d/mysql stop` or `service mysqld stop`.
 - Then run the same command again, replacing `'stop'` with `'start'`.
 - On Mac OS X, run `'sudo /Library/StartupItems/MySQLCOM/MySQLCOM restart'`.

Step 4. Set up your MySQL database and user

In this step you will create a database within MySQL to hold your Confluence data, and a database user with authority to access that database.

To create the database and user privileges:

1. Run the 'mysql' command as a MySQL super user. The default user is 'root' with a blank password.
2. Create an empty Confluence database schema by running this command:

```
CREATE DATABASE confluence CHARACTER SET utf8 COLLATE utf8_bin;
```

3. Create the Confluence database user by running this command. Replace 'confluenceuser' and 'confluencepass' with a username and password of your choice. If Confluence is not running on the same server as your MySQL database server, replace 'localhost' with the hostname or IP address of the Confluence server:

```
GRANT ALL PRIVILEGES ON confluence.* TO 'confluenceuser'@'localhost'  
IDENTIFIED BY 'confluencepass';
```

Step 5. Install Confluence

Install Confluence if you have not done so already. See the [Confluence Installation Guide](#). **Stop immediately after the installation, before opening the Confluence Setup Wizard in your browser**, and follow the steps below.

If you have already got part-way through the Confluence Setup Wizard, stop at the database setup step and follow the steps below. You will be able to restart the setup wizard at the same step later.

Step 6. Download and install the MySQL database driver

If you are **upgrading Confluence to a later version, and you are already using the recommended MySQL driver** (JDBC Connector/J 5.1), you can skip the instructions in this section. The Confluence upgrade task will automatically copy over your existing driver to the upgraded installation.

If you are installing Confluence, or you are upgrading Confluence and not using the recommended MySQL driver (JDBC Connector/J 5.1), follow the steps below.

Choose whether you will set up a **direct JDBC connection or a datasource connection** to MySQL, to suit your environment. If unsure, choose direct JDBC.

To set up a direct JDBC connection:

If you plan to set up a direct JDBC connection to MySQL, you will need to copy the MySQL JDBC driver to your Confluence installation.

1. Get the MySQL driver:
 - If you are **installing Confluence**, download the recommended MySQL driver . Links to the appropriate database drivers are available on this page: [Database JDBC Drivers](#). You can download either the .tar.gz or the .zip archive. Extract the driver JAR file (for example, mysql-connector-java-x.x.x-bin.jar, where x.x.x is a version number) from the archive.
 - If you are **upgrading Confluence to a later version, and you are not using the recommended MySQL driver** (JDBC Connector/J 5.1), copy the driver JAR file from your existing Confluence installation before you upgrade. The driver will be in the <Confluence installation> /confluence/WEB-INF/lib folder.
2. Copy the driver JAR file to the <Confluence installation> /confluence/WEB-INF/lib folder in your new or upgraded Confluence installation.

To set up a datasource connection:

If you plan to set up a datasource connection to MySQL, follow the steps described in [Configuring a MySQL Datasource in Apache Tomcat](#).

Step 7. Check settings for internationalisation

If you are using a existing database, use the `status` command to verify database character encoding information. The results should be UTF-8. See [Configuring Database Character Encoding](#).

Step 8. Set up your database connection in the Confluence Setup Wizard

Start Confluence, and go to the Confluence Setup Wizard in your browser. Follow these steps to set up the new configuration:

1. Follow the initial steps in the [Confluence Setup Guide](#).
2. When prompted to choose an evaluation or production installation, choose **production installation**.
3. When prompted to choose an embedded or external database, select **MySQL** from the dropdown list and choose **External Database**.
(Note: you'll see a warning that a driver is required. You downloaded or copied the driver in step 6 so you can ignore this warning)
Choose either the direct JDBC or the datasource connection, to suit the choice you made earlier when setting up the MySQL database driver.
 - For the JDBC connection: Enter the username (for example, `confluenceuser`) and password you chose earlier.
 - For a datasource connection: Set the JNDI name to `java:comp/env/jdbc/confluence`

Congratulations! Confluence is now using your MySQL database to store its data.

Troubleshooting

- If you get the following error message, verify that you have given the `confluenceuser` user all the required database permissions when connecting from `localhost`.

```
Could not successfully test your database: : Server connection
failure during transaction. Due to underlying exception:
'java.sql.SQLException: Access denied for user
'confluenceuser'@'localhost' (using password: YES)'
```

- The following page contains common issues encountered when setting up your MySQL database to work with Confluence: [Known issues for MySQL](#)
- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If none of the above describes your issue, please create a support ticket at <http://support.atlassian.com> and be sure to include your logs (found in `<CONFLUENCE-INSTALLATION>/logs` and `<CONFLUENCE-HOME>/logs`).
- There's currently a known issue for MySQL 5.5

CONF-37829 - Improve select query that used in Confluence for copying a page
RESOLVED

Configuring a MySQL Datasource in Apache Tomcat

This page tells you how to set up a MySQL datasource connection for Confluence.

Step 1. Shut down Tomcat

1. Run `bin/shutdown.sh` or `bin/shutdown.bat` to bring Tomcat down while you are making these changes.
2. Make a backup of your `<CONFLUENCE_HOME>/confluence.cfg.xml` file and your `<CON`

FLUENCE_INSTALLATION>/conf/server.xml file, so that you can easily revert if you have a problem.

Step 2. Install the MySQL database driver

1. Download the MySQL JDBC driver. Links are available on this page: [Database JDBC Drivers](#).
2. Unpack the archive file you have downloaded, and find the JAR file called something like this: mysql-connector-java-x.x.x-bin.jar, where x.x.x is a version number.
3. Copy the JAR file into the lib folder of your Confluence installation: <CONFLUENCE_INSTALLATION>/lib.

Step 3. Configure Tomcat

1. Edit the conf/server.xml file in your Tomcat installation.
2. Find the following lines:

```
<Context path="" docBase="../confluence" debug="0" reloadable="true">
  <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
  Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
```

3. Insert the DataSource Resource element within the Context element, directly after the opening <Context.../> line, before Manager:

```
<!-- If you're using Confluence 5.7 or below; change maxTotal to maxActive -->
<Resource name="jdbc/confluence" auth="Container"
type="javax.sql.DataSource"
    username="yourusername"
    password="yourpassword"
    driverClassName="com.mysql.jdbc.Driver"

url="jdbc:mysql://localhost:3306/confluence?useUnicode=true&characterEncoding=utf8"
    maxTotal="25"
    maxIdle="10"
    defaultTransactionIsolation="READ_COMMITTED"
    validationQuery="Select 1" />
```

- Replace the username and password parameters with the correct values for your database.
- In the url parameter, replace the word 'confluence' with the name of the database your Confluence data will be stored in.

Notes

- If switching from a direct JDBC connection to a datasource connection, you can find the above details in your <CONFLUENCE_HOME>/confluence.cfg.xml file.
- The configuration properties for Tomcat's standard datasource resource factory (org.apache.tomcat.dbcp.dbcp.BasicDataSourceFactory) are as follows:
 - driverClassName – Fully qualified Java class name of the JDBC driver to be used.
 - maxTotal – The maximum number of database connections in the pool at the same time.
 - maxIdle – The maximum number of connections that can sit idle in this pool at the same time.
 - maxWaitMillis – The maximum number of milliseconds that the pool will wait (when there are no available connections) for a connection to be returned before throwing an exception.
 - password – Database password to be passed to your JDBC driver.

On this page:

- [Step 1. Shut down Tomcat](#)
- [Step 2. Install the MySQL database driver](#)
- [Step 3. Configure Tomcat](#)
- [Step 4. Configure the Confluence web application](#)
- [Step 5. Restart Tomcat](#)

Related pages:

- [Database Setup For MySQL](#)
- [Confluence Installation and Upgrade Guide](#)

- `url` – Connection URL to be passed to your JDBC driver. (For backwards compatibility, the property `driverName` is also recognised.)
- `user` – Database username to be passed to your JDBC driver.
- `validationQuery` – SQL query that can be used by the pool to validate connections before they are returned to the application. If specified, this query must be an SQL SELECT statement that returns at least one row.
- Why is the `validationQuery` element needed? When a database server reboots, or there is a network failure, all the connections in the connection pool are broken and this normally requires an application server reboot. However, the Commons DBCP (Database Connection Pool) which is used by the Tomcat application server can validate connections before issuing them by running a simple SQL query, and if a broken connection is detected, a new one is created to replace it. To do this, you will need to set the `validationQuery` option on the database connection pool.

Step 4. Configure the Confluence web application

1. Edit this file in your Confluence installation: `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`.
2. Insert the following element just before `</web-app>` near the end of the file:

```
<resource-ref>
  <description>Connection Pool</description>
  <res-ref-name>jdbc/confluence</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
```

If you are changing an existing Confluence installation over to using a Tomcat datasource:

1. Edit the `<CONFLUENCE_HOME>/confluence.cfg.xml` file.
2. Delete any line that contains a property that begins with `hibernate`.
3. Insert the following at the start of the `<properties>` section.

```
<property name="hibernate.setup"><![CDATA[true]]></property>
<property
name="hibernate.dialect"><![CDATA[com.atlassian.hibernate.dialect.MySQLDialect]]></property>
<property
name="hibernate.connection.datasource"><![CDATA[java:comp/env/jdbc/confluence]]></property>
```

Step 5. Restart Tomcat

Run `bin/startup.sh` or `bin/startup.bat` to start Tomcat with the new settings.

Database Setup for PostgreSQL

This page provides instructions for configuring Confluence to use a PostgreSQL database.

Step 1. Check the prerequisites

Check the following before you start:

- Check that your version of PostgreSQL is supported. See [Supported Platforms](#). If your version is not supported, please upgrade to a supported version of PostgreSQL before installing Confluence.
- If you have been evaluating Confluence and wish to transfer your data to a new database, consult the following guide first: [Migrating to Another Database](#).
- If you are migrating from another database, consult the following guide first: [Migrating to Another Database](#).

Step 2. Install PostgreSQL

If you do not already have an operational PostgreSQL database, install it now.

On this page:

- [Step 1. Check the prerequisites](#)
- [Step 2. Install PostgreSQL](#)
- [Step 3. Set up your PostgreSQL database and user](#)
- [Step 4. Install Confluence and the PostgreSQL database driver](#)
- [Step 5. Set up your database connection in the Confluence Setup Wizard](#)
- [Notes](#)
- [Troubleshooting](#)

Related pages:

- [Database Configuration](#)
- [Known issues for PostgreSQL](#)
- [Confluence Installation and Upgrade Guide](#)

1. Download and install [PostgreSQL](#). Please note the following information when installing PostgreSQL:
 - The **password** that you are prompted to provide during the installation process is for the 'postgres' account, which is the database root-level account, sometimes called the super user ('postgres'). Remember this username and password. You will need it each time you log in to the database.
 - The **default port** for PostgreSQL is 5432. If you decide to change the default port, please ensure that your new port number does not conflict with any services running on that port. You will also need to remember to update all further mentions of the database port.
 - Choose the **locale** that best fits your geographic location.
 - Do **not** launch **Stack Builder** at the completion of the installer.

Step 3. Set up your PostgreSQL database and user

Next you need to create a database within PostgreSQL to hold your Confluence data, and a database user with authority to access that database.

1. Create a database user (for example `confluenceuser`)
 - Your new user must be able to **create database objects** and **create roles**.
2. Create a database (for example `confluence`)
 - **Owner** is your new database user (for example `confluenceuser`)
 - **Character encoding** should be `utf8` encoding.

You can use [pgAdmin](#) as an alternative to the command line to complete this step. If you used the graphical installer when installing PostgreSQL, pgAdmin will be already installed on your computer.

Step 4. Install Confluence and the PostgreSQL database driver

Decide whether you will set up a **direct JDBC connection** or a **datasource connection** to PostgreSQL, to suit your environment. If unsure, choose direct JDBC.

Install Confluence if you have not done so already. See the [Confluence Installation Guide](#).

- If you plan to set up a **direct JDBC connection** to PostgreSQL, you can run the Confluence installation and move directly on to the Confluence Setup Wizard, as described [below](#). The PostgreSQL JDBC driver is bundled with Confluence, as documented on this page: [Database JDBC Drivers](#).
- If you plan to set up a **datasource connection** to PostgreSQL:
 - Stop immediately after the Confluence installation, before opening the Confluence Setup Wizard in your browser. If you have already got part-way through the Confluence Setup Wizard, stop at the database setup step. You will be able to restart the setup wizard at the same step later.
 - Follow the steps described in [Configuring a PostgreSQL Datasource in Apache Tomcat](#).

Step 5. Set up your database connection in the Confluence Setup Wizard

Start Confluence, go to the [Confluence Setup Wizard](#) in your browser, and follow these steps:

1. When prompted to choose an evaluation or production installation, choose **production installation**.
2. When prompted to choose an embedded or external database, select **PostgreSQL** from the dropdown list and choose **External Database**.
3. Choose either the **direct JDBC** or the **datasource connection**, to suit the choice you made earlier.

- **For the JDBC connection:**

- When prompted for a **Driver Class Name**, enter:

```
org.postgresql.Driver
```

- When prompted for the **Database URL**, use this format:

```
jdbc:postgresql://<server>:<port>/<database>
```

For example: `jdbc:postgresql://localhost:5432/confluence`

Note: If you need to connect to an SSL database, add the `ssl=true` parameter in the database URL. For example: `jdbc:postgresql://localhost:5432/confluence?ssl=true`

- Enter the username (for example `confluenceuser`) and password you chose earlier when setting up your Confluence database.
 - **For a datasource connection:** Set the **Datasource Name** to the following: `java:comp/env/jdbc/confluence`

That's it - Confluence is now using your PostgreSQL database to store its data.

Screenshot: Setting up the PostgreSQL JDBC connection in the Confluence Setup Wizard

Notes

- If the server that is hosting the PostgreSQL database is not the same server as Confluence, then please ensure that the Confluence server can contact the database server. Please also refer to the [PostgreSQL documentation on how to set up pg_hba.conf](#). If the `pg_hba.conf` file is not set properly, remote communication to the PostgreSQL server will fail.
- Running SQL queries: For ongoing maintenance of your server, you can continue to use PGAdmin III as your SQL browser.

Troubleshooting

- If you get the following error message, verify that you have given the `confluenceuser` user all the required database permissions when connecting from `localhost`.

```
Could not successfully test your database: : Server connection
failure during transaction. Due to underlying exception:
'java.sql.SQLException: Access denied for user
'confluenceuser'@'localhost' (using password: YES)'
```

- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If you are unable to connect to the database from Confluence and they are on different machines, most likely you have a firewall in between the two machines or your `pg_hba.conf` file is misconfigured. Verify that your firewall is set to allow connections through 5432 or double check your [hba configuration](#).
- The following page contains common issues encountered when setting up your PostgreSQL database to work with Confluence: [Known issues for PostgreSQL](#).
- If none of the above describes your issue, please create a support ticket at <http://support.atlassian.com> and be sure to include your logs (found in `<CONFLUENCE-INSTALLATION>/logs` and `<CONFLUENCE-HOME>/logs`).

Configuring a PostgreSQL Datasource in Apache Tomcat

This page tells you how to set up a PostgreSQL datasource connection for Confluence.

Step 1. Shut down Tomcat

1. Run `bin/shutdown.sh` or `bin/shutdown.bat` to bring Tomcat down while you are making these changes.
2. Make a backup of your `<CONFLUENCE_HOME>/confluence.cfg.xml` file and your `<CONFLUENCE_INSTALLATION>/conf/server.xml` file, so that you can easily revert if you have a problem.

Step 2. Install the PostgreSQL Server database driver

1. Download the PostgreSQL Server JDBC driver JAR file.
 - Links are available on this page: [Database JDBC Drivers](#).
 - Alternatively, you can get the driver from your Confluence installation: `/confluence/WEB-INF/lib/postgresql-x.x-x.jdbcx.jar`, where 'x' represents a version number.
2. Copy the JAR file into the `lib` folder of your Tomcat installation: `<TOMCAT-INSTALLATION>/lib`.

Step 3. Configure Tomcat

On this page:

- [Step 1. Shut down Tomcat](#)
- [Step 2. Install the PostgreSQL Server database driver](#)
- [Step 3. Configure Tomcat](#)
- [Step 4. Configure the Confluence web application](#)
- [Step 5. Restart Tomcat](#)

Related pages:

- [Database Setup for PostgreSQL](#)

1. Edit the `conf/server.xml` file in your Tomcat installation.
2. Find the following lines:

```
<Context path="" docBase="../../confluence" debug="0" reloadable="true">
  <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
  Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
```

3. Insert the `DataSource` Resource element inside the `Context` element, directly after the opening `<Context.../>` line, before `Manager`:

```
<!-- If you're using Confluence 5.7 or below; change maxTotal to maxActive
-->
<Resource name="jdbc/confluence" auth="Container"
type="javax.sql.DataSource"
    username="postgres"
    password="postgres"
    driverClassName="org.postgresql.Driver"
    url="jdbc:postgresql://localhost:5432/yourDatabaseName"
    maxTotal="25"
    maxIdle="10"
    validationQuery="select 1" />
```

- Replace the `username` and `password` parameters with the correct values for your database
- In the `url` parameter, replace the word `'yourDatabaseName'` with the name of the database your Confluence data will be stored in.

Notes:

- If switching from a direct JDBC connection to `datasource`, you can find the above details in your `<CONFLUENCE_HOME>/confluence.cfg.xml` file.
- Here are the configuration properties for Tomcat's standard data source resource factory (`org.apache.tomcat.dbcp.dbcp.BasicDataSourceFactory`):
 - `driverClassName` — Fully qualified Java class name of the JDBC driver to be used.
 - `maxTotal` — The maximum number of database connections in the pool at the same time.
 - `maxIdle` — The maximum number of connections that can sit idle in this pool at the same time.
 - `maxWaitMillis` — The maximum number of milliseconds that the pool will wait (when there are no available connections) for a connection to be returned before throwing an exception.
 - `password` — Database password to be passed to our JDBC driver.
 - `url` — Connection URL to be passed to our JDBC driver. (For backwards compatibility, the property `driverName` is also recognized.)
 - `user` — Database username to be passed to our JDBC driver.
 - `validationQuery` — SQL query that can be used by the pool to validate connections before they are returned to the application. If specified, this query MUST be an SQL SELECT statement that returns at least one row.
- Why is the `validationQuery` element needed? When a database server reboots, or there is a network failure, all the connections in the connection pool are broken and this normally requires a Application Server reboot. However, the Commons DBCP (Database Connection Pool) which is used by the Tomcat application server can validate connections before issuing them by running a simple SQL query, and if a broken connection is detected, a new one is created to replace it. To do this, you will need to set the `"validationQuery"` option on the database connection pool.

Step 4. Configure the Confluence web application

1. Edit this file in your Confluence installation: `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`.
2. Insert the following element just before `</web-app>` near the end of the file:

```
<resource-ref>
  <description>Connection Pool</description>
  <res-ref-name>jdbc/confluence</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
```

If you are changing an existing Confluence installation over to using a Tomcat datasource:

1. Edit the `<CONFLUENCE_HOME>/confluence.cfg.xml` file.
2. Delete any line that contains a property that begins with `hibernate`.
3. Insert the following at the start of the `<properties>` section:

```
<property name="hibernate.setup"><![CDATA[true]]></property>
<property
name="hibernate.dialect"><![CDATA[net.sf.hibernate.dialect.PostgreSQLDialect]]></property>
<property
name="hibernate.connection.datasource"><![CDATA[java:comp/env/jdbc/confluence]]></property>
```

Step 5. Restart Tomcat

Run `bin/startup.sh` or `bin/startup.bat` to start Tomcat with the new settings.

Embedded H2 Database

To enable you to try Confluence without setting up an external database, your Confluence installation includes an embedded H2 database.

The embedded H2 database is used by default when you choose the **Trial** installation path.

The embedded database files are stored in your Confluence home directory `<confluence-home>/database`.

The embedded H2 database is **only** supported while you are evaluating Confluence. You must migrate to a [supported](#) external database before using Confluence as a production system.

To find out if you are still using the embedded database, go to



> **General Configuration > Support Tools.**

Related pages:

- [Confluence Home and other important directories](#)
- [Database Configuration](#)

Connect to the embedded H2 database using DB Visualizer

[DBVisualizer](#) is just one database administration tool. You can use any administration tool that supports embedded H2 databases. The steps will be similar.

1. Shut down Confluence.
2. Back up your `<confluence-home>/database` directory.
3. Launch [DBVisualizer](#).
4. Choose **Create new database connection** and follow the prompts to set up the connection.

The information you'll need is:

- **Database driver:** H2 embedded
- **Database Userid:** sa
- **Database password:** leave this field blank
- **Database filename:** `<confluence-home>/database/h2db`
⚠ leave off the `.h2.db` file extension.

5. Connect to the database.

Refer to the [DBVisualizer documentation](#) for help using DBVisualizer.

Connect to the embedded H2 database using the H2 console

Alternatively you can connect using the browser based H2 console. The easiest way to access the console is to double click the H2 database jar file at `<installation-directory>\confluence\WEB-INF\lib\h2-x.x.x.jar`.

Migrating to Another Database

This document describes how to migrate your Confluence data from your existing database to another database. The instructions are designed primarily for migrating from an evaluation to a production database. Large data sets will require third party database migration tools.

This page covers the following scenarios:

- Moving from the embedded, trial database to a supported external database.
- Moving from one external database to another, for example from Oracle to PostgreSQL (provided your dataset is not large)
- Upgrading to a new version of the same external database. Note: you don't need to migrate your data if you're upgrading the database in place.

Note: If you are moving your database from one server to another you can change the JDBC URL in `<confluence.home>/confluence.cfg.xml` (if you are using a direct JDBC connection) or in the definition of your datasource (if you are connecting via a datasource).

On this page:

- [Limitations of database migration](#)
- [Database migration](#)
- [Method one – standard procedure](#)
 - [Step 1: Take note of your add-ons](#)
 - [Step 2: Back up your data](#)
 - [Step 3: Set up the new database](#)
 - [Step 4. Install Confluence \(same version number\) in a new location](#)
 - [Step 5. Download and install the database driver if necessary](#)
 - [Step 6. Run the Confluence setup wizard and copy your data to your new database](#)
 - [Step 7. Re-install your add-ons](#)
 - [Step 8. Check settings for new machine](#)
- [Method two – for installations with a large volume of attachments](#)
 - [Before you start](#)
 - [Step 1: Take note of your add-ons](#)
 - [Step 2: Back up your data](#)
 - [Step 3: Set up the new database](#)
 - [Step 4. Install Confluence \(same version number\) in a new location](#)
 - [Step 5. Download and install the database driver if necessary](#)
 - [Step 6. Run the Confluence setup wizard and copy your data to your new database](#)
 - [Step 7: Copy your attachments across](#)
 - [Step 8. Re-install your add-ons](#)
 - [Step 9. Check settings for new machine](#)
- [A note about case sensitivity in your database](#)
 - [Setting up a New Confluence Instance](#)
 - [Migrating an Existing Confluence Instance to a Different Database](#)
- [Troubleshooting](#)

Related pages:

- [Database Configuration](#)
- [Confluence Home and other important directories](#)

Note: The XML export built into Confluence is not suited for the backup or migration of large data sets. There are a number of third party tools that may be able to assist you with the data migration. If you would like help in selecting the right tool, or help with the migration itself, we can put you in touch with one of the [Atlassian Experts](#).

Database migration

There are two ways you can perform the migration, both described on this page:

1. **Method one** is the standard procedure.
2. Use **method two** if the total size of attachments in your installation exceeds 500MB.

Method one – standard procedure

Step 1: Take note of your add-ons

Take note of the add-ons (plugins) currently installed and enabled in Confluence, so that you can reinstate them later. Make a note of the following for each add-on:

- Add-on name
- Version
- Enabled or disabled status. This is useful if you have enabled or disabled modules yourself, making your configuration differ from the default.

Step 2: Back up your data

1. Create an XML backup of your existing data, via the **Confluence administration console**. See [Manually Backing Up the Site](#). Make a note of the location where you put the XML file. You will need it later to import your Confluence data into your new database.
2. Shut down Confluence.
3. Make a copy of the [Confluence Home](#) directory. This is a precautionary measure, to ensure you can recover your data if it is mistakenly overwritten.
4. If you are using an external database, make a separate backup using the utilities that were installed with that database. This also is a precautionary measure.

Step 3: Set up the new database

Choose the [database setup instructions](#) for your new database, and follow those instructions to do the following:

- Install the database server.
- Perform any required configuration of the database server, as instructed.
- Add the Confluence database and user. Make a note of the username and password that you define in this step. You will need them later, when running the Confluence Setup Wizard.

Step 4. Install Confluence (same version number) in a new location

Now you will install Confluence again, with a different home directory path and installation path.

Note: You must use the same version of Confluence as the existing installation. (If you want to upgrade Confluence, you must do it as a separate step.) For example, if your current site is running Confluence 5.1.2, your new installation must also be Confluence 5.1.2.

When running the [Confluence installer](#):

- Choose **Custom Install**. (Do **not** choose to upgrade your existing installation.)
- Choose a **new destination directory**. This is the installation directory for your new Confluence. It must not be the same as the existing Confluence installation.
- Choose a **new home directory**. This is the data directory for your new Confluence. It must not be the same as the existing Confluence installation.

Step 5. Download and install the database driver if necessary

Note that Confluence bundles some database drivers, but you'll need to install the driver yourself if it is not bundled. Follow the [database setup instructions](#) for your new database, to download and install the database driver if necessary.

Step 6. Run the Confluence setup wizard and copy your data to your new database

When running the [Confluence setup wizard](#):

- Enter your license key, as usual.
- Choose **Production Installation** as the installation type.
- In the database configuration step, choose your new database type from the dropdown menu, then choose **External Database**.
- Choose the connection type: **Direct JDBC** or **Datasource**. If you are not sure which, choose 'Direct JDBC'. This is the most common connection type.
- When prompted for the database **user and password**, supply the credentials you defined earlier when adding the Confluence database to your database server.
- On the load content step, choose **Restore From Backup**. This is where you will import the data from your XML backup. There are two options for accessing the XML file:
 - Browse to the location of your XML backup on your network, and choose **Upload and Restore**.
 - Alternatively, put the XML file in the Confluence home directory of the new site (`<CONFLUENCE-HOME-DIRECTORY>\restore`) then choose **Restore**.

Note: If you choose not to restore during the Confluence setup wizard, you can do the import later. Go to the Confluence administration console and choose to restore an XML backup. See [Site Backup and Restore](#).

Step 7. Re-install your add-ons

Re-install any add-ons (plugins) that are not bundled with Confluence.

- Use the same version of the add-on as on your old Confluence site.
- The data created by the add-ons will already exist in your new Confluence site, because it is included in the XML backup.

Step 8. Check settings for new machine

If you are moving Confluence to a different machine, you need to check the following settings:

- Configure your new base URL. See [Configuring the Server Base URL](#).
- Check your application links. See [Linking to Another Application](#).
- Update any gadget subscriptions from external sites pointing to this Confluence site. For example, if your JIRA site subscribes to Confluence gadgets, you will need to update your JIRA site.
- Review any other resources that other systems are consuming from Confluence.

Method two – for installations with a large volume of attachments**Before you start**

These instructions only apply to attachments stored in the file system. If you store attachments in the database see [Attachment Storage Configuration](#) to find out how to migrate between different attachment storage methods.

Step 1: Take note of your add-ons

Take note of the add-ons (plugins) currently installed and enabled in Confluence, so that you can reinstate them later. Make a note of the following for each add-on:

- Add-on name
- Version
- Enabled or disabled status. This is useful if you have enabled or disabled modules yourself, making your configuration differ from the default.

Step 2: Back up your data

1. Create an XML backup of your existing data, via the **Confluence administration console**. See [Manually Backing Up the Site](#). Make a note of the location where you put the XML file. You will need it later to import your Confluence data into your new database.
2. Shut down Confluence.
3. Make a copy of the attachments directory (`<CONFLUENCE-HOME-DIRECTORY>\attachments`) in your Confluence Home directory. You will need it later to copy your Confluence attachments data into your new Confluence installation.

4. If you are using an external database, make a separate backup using the utilities that were installed with that database. This also is a precautionary measure.

Step 3: Set up the new database

Choose the [database setup instructions](#) for your new database, and follow those instructions to do the following:

- Install the database server.
- Perform any required configuration of the database server, as instructed.
- Add the Confluence database and user. Make a note of the username and password that you define in this step. You will need them later, when running the Confluence Setup Wizard.

Step 4. Install Confluence (same version number) in a new location

Now you will install Confluence again, with a different home directory path and installation path.

Note: You must use the same version of Confluence as the existing installation. (If you want to upgrade Confluence, you must do it as a separate step.) For example, if your current site is running Confluence 5.1.2, your new installation must also be Confluence 5.1.2.

When running the [Confluence installer](#):

- Choose **Custom Install**. (Do **not** choose to upgrade your existing installation.)
- Choose a **new destination directory**. This is the installation directory for your new Confluence. It must not be the same as the existing Confluence installation.
- Choose a **new home directory**. This is the data directory for your new Confluence. It must not be the same as the existing Confluence installation.

Step 5. Download and install the database driver if necessary

Note that Confluence bundles some database drivers, but you'll need to install the driver yourself if it is not bundled. Follow the [database setup instructions](#) for your new database, to download and install the database driver if necessary.

Step 6. Run the Confluence setup wizard and copy your data to your new database

When running the [Confluence setup wizard](#):

- Enter your license key, as usual.
- Choose **Production Installation** as the installation type.
- In the database configuration step, choose your new database type from the dropdown menu, then choose **External Database**.
- Choose the connection type: **Direct JDBC** or **Datasource**. If you are not sure which, choose 'Direct JDBC'. This is the most common connection type.
- When prompted for the database **user and password**, supply the credentials you defined earlier when adding the Confluence database to your database server.
- On the load content step, choose **Restore From Backup**. This is where you will import the data from your XML backup. There are two options for accessing the XML file:
 - Browse to the location of your XML backup on your network, and choose **Upload and Restore**.
 - Alternatively, put the XML file in the Confluence home directory of the new site (<CONFLUENCE-HOME-DIRECTORY>\restore) then choose **Restore**.

Note: If you choose not to restore during the Confluence setup wizard, you can do the import later. Go to the Confluence administration console and choose to restore an XML backup. See [Site Backup and Restore](#).

Step 7: Copy your attachments across

Copy the contents of the attachments directory (<CONFLUENCE-HOME-DIRECTORY>\attachments) from your old Confluence Home directory to your new Confluence Home directory.

Step 8. Re-install your add-ons

Re-install any add-ons (plugins) that are not bundled with Confluence.

- Use the same version of the add-on as on your old Confluence site.
- The data created by the add-ons will already exist in your new Confluence site, because it is included

in the XML backup.

Step 9. Check settings for new machine

If you are moving Confluence to a different machine, you need to check the following settings:

- Configure your new base URL. See [Configuring the Server Base URL](#).
- Check your application links. See [Linking to Another Application](#).
- Update any gadget subscriptions from external sites pointing to this Confluence site. For example, if your JIRA site subscribes to Confluence gadgets, you will need to update your JIRA site.
- Review any other resources that other systems are consuming from Confluence.

A note about case sensitivity in your database

'Collation' refers to a set of rules that determine how data is sorted and compared. Case sensitivity is one aspect of collation. Other aspects include sensitivity to kana (Japanese script) and to width (single versus double byte characters).

Case sensitive or case insensitive collation – how should you create your Confluence database? What about when you are migrating your existing Confluence instance from one database to another?

Setting up a New Confluence Instance

For new Confluence instances, we recommend using case sensitive collation for your Confluence database. This is the default collation type used by many database systems.

Note: Even if the database is configured for case sensitive collation, Confluence reduces all usernames to lower case characters before storing them in the database. For example, this means that 'joebloggs', 'joeBloggs' and 'JoeBloggs' will be treated as the same username.

Migrating an Existing Confluence Instance to a Different Database

The default Confluence configuration uses case sensitive database collation. This is typical of databases created under default conditions. If you are migrating from this type of configuration to a new database, we recommend that the new database uses case sensitive collation. If you use case insensitive collation, you may encounter data integrity problems after migration (for example, via an XML import) if data stored within your original Confluence site required case sensitive distinctions.

Troubleshooting

See our [troubleshooting guide](#) if you're unable to restore your XML backup.

Configuring Database Character Encoding

The database used with Confluence should be configured to use the same character encoding as Confluence. **The recommended encoding is Unicode UTF-8 (the equivalent for Oracle databases is AL32UTF8).**

There are two places where character encoding may need to be configured:

- when creating the database
- when connecting to the database (JDBC connection URL or properties).

The configuration details for each type of database are different. Some examples are below.

JDBC connection settings

MySQL

Append "useUnicode=true to your JDBC URL:

On this page:

- [JDBC connection settings](#)
- [Creating a UTF-8 database](#)
- [Updating existing database to UTF-8](#)

Related pages:

- [Configuring Character Encoding](#)
- [Known issues for MySQL](#)

```
jdbc:mysql://hostname:port/database?useUnicode=true&characterEncoding=utf8
```


If you are modifying `confluence.cfg.xml` directly rather than via the Confluence Installation GUI, you'll need to escape out the `&` in the URL string as this is a reserved XML token and will break the syntax when the XML is parsed. An effective URL could be similar to:

```
<property
name="hibernate.connection.url">jdbc:mysql://hostname:port/data
base?useUnicode=true&amp;characterEncoding=utf8</property>
```

Creating a UTF-8 database

MySQL

1. Create a UTF-8 database with binary UTF-8 collation.

 Binary UTF-8 provides case-sensitive collation.

```
CREATE DATABASE confluence CHARACTER SET utf8 COLLATE utf8_bin;
```

2. You will also need to set the Server Character set to utf8. This can be done by adding the following in `my.ini` for Windows or `my.cnf` for other OS. It has to be declared in the Server section, which is the section after `[mysqld]`:

```
[mysqld]
default-character-set=utf8
```

If the above option does not work, try using `character_set_server=utf8` in lieu of `default-character-set=utf8`

3. Use the **status** command to verify database character encoding information.

Screenshot: Using the Status Command to Verify Database Character Encoding

```
mysql> CREATE DATABASE confluence CHARACTER SET utf8 COLLATE utf8_bin;
Query OK, 1 row affected (0.02 sec)

mysql> show databases;
+-----+
| Database |
+-----+
| information_schema |
| confluence |
| confluencedb |
| mysql |
| test |
+-----+
5 rows in set (0.02 sec)

mysql> use confluence;
Database changed
mysql> status;
-----
C:\Program Files\MySQL\MySQL Server 5.0\bin\mysql.exe  Ver 14.12 Distrib 5.0.83,
for Win32 (ia32)

Connection id:          1800
Current database:       confluence
Current user:           root@localhost
SSL:                    Not in use
Using delimiter:        ;
Server version:         5.0.83-community-nt MySQL Community Edition (GPL)
Protocol version:       10
Connection:             localhost via TCP/IP
Server characterset:    utf8
Db characterset:        utf8
Client characterset:    utf8
Conn. characterset:     utf8
TCP port:               3306
Uptime:                 20 hours 56 min 23 sec
```

4. In some cases, the individual tables collation and character encoding may differ from the one that the database as a whole has been configured to use. Please use the command below to ensure all tables within your Confluence database are correctly configured to use UTF-8 character encoding and binary UTF-8 collation:

```
use confluence;
show table status;
```

Check for the value listed under the **Collation** column, to ensure it has been set to `utf8_bin` (that is, case-sensitive) collation for all tables.

If not, then this can be changed by the following command, executed for each table in the Confluence database:

```
ALTER TABLE tablename CONVERT TO CHARACTER SET utf8 COLLATE utf8_bin;
```

Please substitute the <tablename> above, with each table within the confluence database.

Relevant MySQL manual for more detailed explanation:

- [Specifying Character Sets and Collations documentation.](#)
- [Connection Character Sets and Collations.](#)
- [SHOW TABLE STATUS Syntax.](#)
- [ALTER TABLE Syntax.](#)

PostgreSQL

```
CREATE DATABASE confluence WITH ENCODING 'UNICODE';
```

Or from the command-line:

```
$ createdb -E UNICODE confluence
```

For more information see the [PostgreSQL documentation](#).

For PostgreSQL running under Windows

Please note that international characters sets are only fully supported and functional when using PostgreSQL 8.1 and above under Microsoft Windows.

For PostgreSQL running under Linux

Please make sure you check the following to ensure proper handling of international characters in your database

When PostgreSQL creates an initial database cluster, it sets certain important configuration options based on the host environment. The command responsible for creating the PostgreSQL environment `initdb` will check environment variables such as `LC_CTYPE` and `LC_COLLATE` (or the more general `LC_ALL`) for settings to use as database defaults related to international string handling. As such it is important to make sure that your PostgreSQL environment is configured correctly **before you install Confluence**.

To do this, connect to your PostgreSQL instance using `psql` and issue the following command:

```
SHOW LC_CTYPE;
```

If `LC_CTYPE` is set to either "C" or "POSIX" then certain string functions such as converting to and from upper and lower case will not work correctly with international characters. Correct settings for this value take the form `<LOCALE> . <ENCODING>` (`en_AU.UTF8` for example).

If your `LC_CTYPE` is incorrect please check the [PostgreSQL documentation for information on configuring database localisation](#). It is not easy to change these settings with a database that already contains data.

Updating existing database to UTF-8

MySQL database with existing data

For an existing database

If you're using an existing database, confirm the Character Encoding by executing the query: `SHOW VARIABLES LIKE 'character%';` and `SHOW VARIABLES LIKE 'collation%';`. The results should be UTF-8.

Before proceeding with the following changes, please **backup your database**.

This example shows how to change your database from latin1 to utf8, where your database is named "confluence".

1. Dump the database (except the plugindata table) to a text file using the `mysqldump` tool from the command-line :

```
mysqldump -p --default-character-set=latin1 -u <username>  
--skip-set-charset --ignore-table='`confluence`.plugindata'  
--ignore-table='`confluence`.attachmentdata' confluence >  
confluence_database.sql
```
2. Dump the plugindata and attachmentdata tables to a text file using `mysqldump` separately. This is done separately as the recode step below can corrupt the binary data in these tables :

```
mysqldump -p --default-character-set=latin1 -u <username>
```

- ```
--skip-set-charset confluence attachmentdata plugindata >
confluence_blobtables.sql
```
3. copy confluence\_database.sql to confluence\_utf8.sql
  4. Open confluence\_utf8.sql in a text editor and change all character sets from 'latin1' to 'utf8'
  5. Encode all the latin1 characters as UTF-8:  
`recode latin1..utf8 confluence_utf8.sql` (the recode utility is described at <http://directory.fsf.org/recode.html>; it can actually be downloaded from <http://recode.progiciels-bpi.ca/>, and is available for Ubuntu via `apt-get`)

In MySQL:

1. `DROP DATABASE confluence;`
2. `CREATE DATABASE confluence CHARACTER SET utf8 COLLATE utf8_bin;`

Reimport the UTF-8 text file, and also the plugindata and attachmentdata dumps:

1. `mysql -u <username> -p --default-character-set=utf8 --max_allowed_packet=64M confluence < /home/confluence/confluence_utf8.sql`
2. `mysql -u <username> -p --default-character-set=latin1 --max_allowed_packet=64M confluence < /home/confluence/confluence_blobtables.sql`

To support large imports, the parameter '`--max_allowed_packet=64M`' used above sets the maximum size of an SQL statement to be very large. In some circumstances, you may need to increase it further, especially if attachments are stored in the database.

Finally, since the plugindata and attachmentdata tables were not actually converted before, you need to instruct MySQL to convert all the text fields manually:

1. `ALTER TABLE plugindata CONVERT TO CHARACTER SET utf8;`
2. `ALTER TABLE attachmentdata CONVERT TO CHARACTER SET utf8;`

### Testing database encoding

See [Troubleshooting Character Encodings](#) for a number of tests you can run to ensure your database encoding is correct.

## Configuring database query timeout

If database queries are taking too long to perform, and your application is becoming unresponsive, you can configure a timeout for database queries. There is no default timeout in Confluence. To configure a database query timeout, do the following on your test server:


1. Shut down Confluence.
2. Extract `databaseSubsystemContext.xml` from the `confluence-x.x.x.jar` that is in `confluence/WEB-INF/lib/`, and put a copy in `confluence/WEB-INF/classes/`.
3. Edit `confluence/WEB-INF/classes/databaseSubsystemContext.xml` to add the `defaultTimeout` property to the "transactionManager" bean:

```
<bean id="tenantedTransactionManager"
class="org.springframework.orm.hibernate.HibernateTransactionManager"
plugin:available="true">
 <property name="sessionFactory" ref="sessionFactory"/>
 <property name="defaultTimeout" value="120"/>
</bean>
```

The timeout is measured in seconds and will forcibly abort queries that take longer than this. In some cases, these errors are not handled gracefully by Confluence and will result in the user seeing the Confluence error page.

4. Start Confluence.

Once the timeout is working properly in your test environment, migrate the configuration change to Confluence.

 You will need to reapply these changes when upgrading Confluence, as the original `databaseSubsystemContext.xml` file changes from version to version.

## Troubleshooting External Database Connections

A common administration issue when configuring Confluence is identifying database connectivity problems. This page tells you about a helper utility, in the form of a JSP page, that can help you to isolate database connectivity issues. It checks whether you can connect to a database with your application server. If your application server cannot connect to the database, Confluence will not be able to connect to the database either.

### Introduction to the Atlassian Database Check Utility

You can use this utility to:

- Check that your application server can successfully query your database, either via immediate JDBC connectivity or a datasource in the context of your application server.
- Pinpoint problems in your configuration which may occur if the above is failing.

This is what the utility does:

- Check that a JDBC driver can be loaded into memory and view what is already loaded.
- Connect to a JDBC URL and do a 'select 1' from the database.
- Find a DataSource in the JNDI environment and do the above.
- View the System classpath (to ensure that the JDBC JAR file is there).

### Using the Utility

If you have already set up Confluence completely


1. Download the attached [testdatabase.jsp](#) to your `<confluence-install>\confluence` directory.
2. Restart Confluence
3. Go to <http://MY-CONF-SERVER:MY-CONF-PORT/testdatabase.jsp>.
4. Check that your database driver is loaded into memory. If not, check the system classpath for the JDBC driver file, and that the driver is in the `<confluence-install>\lib` directory (for Confluence version 2.10 onwards) or `<confluence-install>\common\lib` (for earlier versions). Here are some [instructions](#).
5. Enter the DB settings Confluence is using and test the database. If an error appears, check that the db service is running, the location matches, and that any users specified actually exist with the right login and permissions. You may be able to find a workaround by Googling the error.

If you cannot set up Confluence because of an error in 'Configuring Database'

1. Record the DB settings you are using for your direct JDBC or datasource connection in the 'Configure Database' step of your setup.
2. Download the attached [testdatabase.jsp](#) to your `<confluence-install>\confluence` directory.
3. Rename your `<confluence-install>\confluence\WEB-INF\web.xml` file to `backup web.xml`. This disables redirection.
4. Restart Confluence.
5. Go to <http://MY-CONF-SERVER:MY-CONF-PORT/testdatabase.jsp>.
6. Check that your database driver is loaded into memory. If not, check the system classpath for the JDBC driver file, and that the driver is in the `<confluence-install>\common\lib` directory as described in [these instructions](#).
7. Enter the DB settings you recorded and test the database. If an error appears, check that the db service is running, the location matches, and that any users specified actually exist with the right login and permissions. You may be able to find a workaround by Googling the error.
8. After correcting the error, rename `<confluence-install>\confluence\WEB-INF\backup web.xml` back to `web.xml`.

### Notes

If you use this utility, please let us know ways in which we could improve it or leave helpful hints for others here.

 For a comprehensive set of database instructions that might be helpful for troubleshooting, please refer to the following links:

- [PostgreSQL](#)

- [MySQL](#)

### Requesting Technical Support

If you are still stuck after attempting the suggestions above, [lodge a free technical support request](#) with information on your database setup.

### Surviving Database Connection Closures

When a database server reboots or a network failure has occurred, all connections in the database connection pool are broken. To overcome this issue, Confluence would normally need to be restarted.

However, database connections in the database connection pool can be validated by running a simple SQL query. If a broken database connection is detected in the pool, a new one is created to replace it.

To do this, you can specify an optional validation query for your database connection. Depending on whether you are using a direct JDBC URL, or a data source, this is configured differently.

### Determining the validation query SQL for your database type

Different database types have slightly different SQL syntax requirements for their validation query. The validation query should be as simple as possible, as this is run every time a connection is retrieved from the pool.

The following validation queries are recommended for the following types of databases:

| Database Type        | Validation Query                |
|----------------------|---------------------------------|
| MySQL                | <code>select 1</code>           |
| Microsoft SQL Server | <code>select 1</code>           |
| Oracle               | <code>select 1 from dual</code> |
| PostgreSQL           | <code>select 1</code>           |

### Enabling validation query using direct JDBC

**To ensure Confluence validates database connections in the database connection pool:**

1. Shut down Confluence
2. Edit the `confluence.cfg.xml` file at the root of your [Confluence Home and other important directories](#)
3. Add the property `"hibernate.c3p0.validate"` and set it to `"true"`, and add the property `"hibernate.c3p0.preferredTestQuery"` and set it to the value of the query you determined above for your database type. See this excerpt of the file with the two added properties for details:

#### confluence.cfg.xml (excerpt)

```
...
<property name="hibernate.c3p0.acquire_increment">1</property>
<property name="hibernate.c3p0.idle_test_period">100</property>
<property name="hibernate.c3p0.max_size">30</property>
<property name="hibernate.c3p0.max_statements">0</property>
<property name="hibernate.c3p0.min_size">0</property>
<property name="hibernate.c3p0.timeout">30</property>
<property name="hibernate.c3p0.validate">true</property>
<property name="hibernate.c3p0.preferredTestQuery">select 1</property>
...
```

4. Save `confluence.cfg.xml`
5. Restart Confluence

## Ensuring validation query using a data source

To ensure Confluence validates database connections in the database connection pool:

1. Shut down Confluence (or the Tomcat installation running Confluence).
2. Edit the `conf/server.xml` file in your Confluence Install Directory, or in the Tomcat installation's CATALINA\_HOME directory.
3. Find the Resource element for your data source, and add the "validationQuery" field, with the value of the query you determined above for your database type. See this excerpt of the file with this added for details:

### server.xml (excerpt)

```
...
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
 username="postgres"
 password="postgres"
 driverClassName="org.postgresql.Driver"
 url="jdbc:postgresql://localhost:5432/yourDatabaseName"
 maxTotal="25"
 maxIdle="10"
 validationQuery="select 1" />
...
```

4. Save `conf/server.xml`
5. Restart Confluence (or the Tomcat installation running Confluence).

## Results and Considerations

You should now be able to recover from a complete loss of all connections in the database connection pool without the need to restart Confluence or the application server running Confluence.

### Performance Considerations:

- Setting this option has a performance impact. The overall decrease in performance should be minimal, as the query itself is quick to run. In addition, the query will only execute when you make a connection. Thus, if the connection is kept for the duration of a request, the query will only occur once per request.
- If you are running a large Confluence installation, you may wish to assess the performance impact of this change before implementing it.

## Site Backup and Restore


Atlassian recommends establishing a [backup strategy using a native database tool](#) for production installations of Confluence.

By default, Confluence backs up all data and attachments once a day to an XML backup file. These files are called XML site backups, and are stored in the `backups` directory of Confluence home. You can also create XML site backups manually. This mechanism is intended for small to medium-sized deployments of Confluence. It is not intended for use with large deployments with lots of pages and attachments (see below).

- [Restore your site from an XML site backup](#)
- [Manually create an XML site backup](#)
- [Configuring Backups](#)
- [User Submitted Backup & Restore Scripts](#)

XML site backups are fine for most small to medium-sized instances of Confluence, containing a few thousand pages and attachments. However, large instances of Confluence may find that backups become slow to create and use large amounts of disk space.

**Note:** Plugins are not included in the XML backup. After importing your backup into a new Confluence site, you will need to re-install all plugins (add-ons) that are not bundled with Confluence. (The `plugindata` table is not backed up in a manual backup.)

 The information on this page *does not apply to* Confluence Cloud.

## Backups for large installations

XML site backups are unsuitable for installations of Confluence that contain thousands of pages, as XML backups take progressively longer to complete as the amount of text increases. Another issue with XML site backups is that Confluence instances with gigabytes of attachments will consume disk space rapidly. This is because each site backup contains all content needed for a site restore. For example, if a 1 GB instance of Confluence is backed up daily, it will create 30 GB of backups per month if left unattended. When administering a large instance, you can reduce disk space by setting XML site backups to exclude attachments, then manually scheduling a backup of your attachments from the Confluence [home directory](#) or database. The backup manager can save space by saving changed files instead of all content.

| Creation Delay | Disk Usage   | Recommended Backup Method                                            |
|----------------|--------------|----------------------------------------------------------------------|
| Acceptable     | Acceptable   | XML site backup with attachments                                     |
| Acceptable     | Unacceptable | XML site backup minus attachments, plus manual backup of attachments |
| Unacceptable   | Unacceptable | Manual backup of database and attachments                            |

**Creation Delay** is the time it takes to create an XML site backup *minus attachments*.

**Disk Usage** can be estimated by multiplying the frequency of your XML site backups by their current size.

## Manual backups

Confluence's [Attachment Storage Configuration](#) can be set to store attachments in the Confluence [home directory](#), or in the database.

### Database backup

Use your Database Administration Tool to create a backup of your Confluence database. If your database is storing your attachments, importing this later will restore all content. For instances with big attachments, please note that currently Confluence migrate attachments in a single transaction: [CONF-9888](#).

### Attachment backup

If stored on the filesystem, attachments are placed under the `attachments` directory of your Confluence home directory. Copy this directory to create a backup of all attachments.

To restore from these backups, please refer to [Restoring Data from other Backups](#).

## Production Backup Strategy

Confluence automatic daily XML backup is suitable if you:

- are evaluating Confluence
- do not have database administration familiarity, and your Confluence installation is small

Once your Confluence installation reaches more than a few thousand pages, the XML backup facility can be inefficient compared to your database's own backup tools. The built in backup functionality requires a lot of memory to run and is less reliable when restoring data.

### Establishing a production system backup solution

Atlassian recommends establishing an alternative


#### On this page:

- [Establishing a production system backup solution](#)
- [Which files need to be backed up?](#)
- [How do I back up?](#)
- [How do I restore?](#)
- [Other processes](#)

#### Related pages:

- [Site Backup and Restore](#)

database backup strategy:

- Create a backup or dump of your database using tools provided by your database  
 To avoid any data inconsistency and corruption, it is recommended to shut down Confluence before creating a database backup or dump.
- Create a file system backup of your Confluence home directory
- Create a file system backup of your shared home directory (Clustered instances only)

Once this is in place, disable the daily backups through the [scheduled jobs](#) feature via '**Administration Console > Administration > Scheduled Jobs**'.

We want to stress that creating these two backups is *better* than having a Confluence XML backup. It is more robust and far more reliable for large production instances. You will be able to restore your whole site, including all data, attachments and configuration information intact with these two backups. See [Restoring Data from other Backups](#).

#### Which files need to be backed up?

Backing up the whole home directory is the safest option, however most files and directories are populated on startup and can be ignored. At minimum, these files/directories *must* be backed up:

- **<conf-home>/confluence.cfg.xml**
- **<conf-home>/attachments**

The rest of the directories will be auto-populated on start up. You may also like to backup these directories:

- **<conf-home>/config** – if you have modified your ehcache.xml file.
- **<conf-home>/index** – if your site is large or reindexing takes a long time – this will avoid the need for a full reindex when restoring.

The location of the home directory is configured on installation and is specified in the `confluence.init.properties` file. For installation created with the automatic installer the default locations are:

- Windows `C:\Program Files\Atlassian\Application Data\Confluence`
- Linux `/var/atlassian/application-data/confluence`

**For Clustered instances only:** Backing up the whole shared home directory is the safest option, however some files and directories are populated at runtime and can be ignored:

- **<conf-home>/thumbnails**
- **<conf-home>/viewfile.**

#### How do I back up?

The commands to back up your database will vary depending on your database vendor, for example the command for PostgreSQL is `pg_dump dbname > outfile`.

You should refer to the documentation for your particular database to find out more.

#### How do I restore?

Our guide on [Migrating Confluence Between Servers](#) has instructions on restoring a backup using this technique.

#### Other processes

XML backups are described and used for other processes in Confluence, like upgrading and moving servers. Using the backup strategy described above will work for those processes too.

- Our [upgrade guide](#) does not require the use of an XML backup (although the earlier Confluence upgrade procedure, and the JIRA upgrade guide, do use XML backups).
- Our [migrate server procedure](#)– used to set up a test server – can use a SQL dump as well.
- The [database migration](#) procedure uses the XML backup for small data sets. Large data sets will require third party database migration tools.

**Note:** The XML export built into Confluence is not suited for the backup or migration of large data sets. There are a number of third party tools that may be able to assist you with the data migration. If you would like help in selecting the right tool, or help with the migration itself, we can put you in touch with one of the [Atlassian Experts](#).

## Configuring Backups

Confluence backs up your data regularly into a zipped XML file. By default, this backup is performed at 2.00 a.m. each day and the backup files are stored in the `backups` folder under the [Confluence Home directory](#). The default naming convention for the backup files is `'backup-yyyy_MM_dd'`. Confluence can write backups to both local and mapped network drives.

### On this page:

- [Configuring Confluence Backups](#)
- [Enabling Backup Path Configuration](#)
- [Notes](#)

From the **Backup Administration** section of Confluence's administration console, you can:

- Include or exclude attachments in backups.
- Configure a different path to store backup files. (By default, this option is not available. See [below](#) for information about enabling the configuration option.)
- Change the naming format used for the files.

✔ You can also change the schedule of this backup using Confluence's [scheduled jobs](#) feature.

i You need to have [System Administrator](#) permissions in order to configure these options.

## Configuring Confluence Backups

### To configure Confluence backups:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Click **'Backup Administration'** in the **'Configuration'** section.
3. Click the **'Edit'** button on the **'Backup Administration'** screen.
4. Now you can do the following:
  - To use a different naming prefix format — Enter the new format in the **'Backup File Prefix'** input field.
  - To use a different date format — Enter the date format in the **'Backup File Date Pattern'** input field using the syntax described in [this document from Sun](#).
  - To exclude attachments from backups — Deselect **'Backup Attachments'**. By default, this feature is 'On'.
  - To specify an alternate path to store backup files (if enabled) — Select **'Custom'** and then enter the path. The directory must be on either a local drive or a [mounted network drive](#).
5. **'Save'** your changes.



Notes:

- By default, this option is not available. See [below](#) for information about enabling the configuration option.
- Please ensure the mapped drive is on a physical server, not a Virtual Machine image.

✔ You can disable Confluence backups through the [scheduled jobs](#) feature.

## Backup Administration

Perform a backup of your site daily to a chosen directory on your filesystem.

### Backup Settings

Backup File Prefix

Backup File Date Pattern

Backup Path

Custom backup paths are not enabled. [More about custom backup paths](#)

☒ Backup Attachments

Screenshot above: Editing the Backup Configuration

### Enabling Backup Path Configuration

By default, it is not possible to specify a backup path via the Confluence Administration Console. This feature is disabled by default for security reasons. Administrators can restore this functionality by updating the relevant configuration property as described below. However, we recommend that you turn the feature **off** in production environments. For production environments, please review our [Production Backup Strategy](#).

#### To enable the configuration option:

1. Edit the `confluence.cfg.xml` file found in the [Confluence Home and other important directories](#).
2. Set the value of property `admin.ui.allow.daily.backup.custom.location` to 'true' (without the quotation marks).

confluence.cfg.xml

```
<property
name="admin.ui.allow.daily.backup.custom.location">true</property>
```

3. Restart Confluence.


If the value of the above configuration property is 'true', it will be possible to specify a backup path via the Confluence Administration Console. If the value of this property is 'false' or the property is not present in the configuration file, the backup path is not configurable.

### Notes

#### *Time is derived from the Confluence server*

The time zone is taken from the server on which Confluence is running.

#### To check the time according to the server, do the following:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration
2. Click '**System Information**' in the left-hand panel and look at the '**System Time**'.

#### **Backup strategy for large Confluence sites**

Consider using the [production backup strategy](#) if your Confluence site is large or you are encountering problems with your automated backup.

### User Submitted Backup & Restore Scripts

These scripts are user-submitted and should be used with caution as they are not covered by Atlassian technical support. If you have questions on how to use or modify these scripts, please [post them to Atlassian Answers](#).

#### Delete Old Backups - Wscript Script On Windows

This script examines backup filename and deletes them if necessary, it may need to be edited.

```
'If you want 3 day old files to be deleted then insert 3 next to Date - "your
number here"
'This script will search out and delete files with this string in them
".2005-12-04-" This of course depends on the number you enter.
'You can always do a wscript.echo strYesterday or strFileName to see what the
script thinks you are searching for.

dtmYesterday = Date - 3

strYear = Year(dtmYesterday)

strMonth = Month(dtmYesterday)
If Len(strMonth) = 1 Then
 strMonth = "0" & strMonth
End If

strDay = Day(dtmYesterday)
If Len(strDay) = 1 Then
 strDay = "0" & strDay
End If

strYesterday = strYear & "-" & strMonth & "-" & strDay

strFileName = "C:\test*." & strYesterday & "-*"

Set objFSO = CreateObject("Scripting.FileSystemObject")
objFSO.DeleteFile(strFileName)
```

#### Delete Old Backups - Basic Bash Script For Linux

Old XML backups can be deleted automatically by inserting a nightly or weekly automation script or cron similar to the following:

```
ls -t <path to your backup dir>/* | tail -n +6 | xargs -i rm {}
```

Or, using the older form of the `tail` command if your system does not support the standard form:

```
ls -t <path to your backup dir>/* | tail +6 | xargs -i rm {}
```

#### Delete Old Backups - Advanced Bash Script For Linux

Old XML backups can be deleted automatically by inserting a nightly or weekly automation script or cron similar to the following. Set the `BACKUP_DIR` and `DAYS_TO_RETAIN` variables to appropriate values for your site. Between runs, more files than `DAYS_TO_RETAIN` builds up.

```
#!/bin/sh

Script to remove the older Confluence backup files.
Currently we retain at least the last two weeks worth
of backup files in order to restore if needed.

BACKUP_DIR="/data/web/confluence/backups"
DAYS_TO_RETAIN=14

find $BACKUP_DIR -maxdepth 1 -type f -ctime +$DAYS_TO_RETAIN -delete
```

#### Manual Database & Home Backup - Bash Script For Linux

This backs up a mySQL database and the Confluence home directory.

```
#!/bin/bash
CNFL=/var/confluence
CNFL_BACKUP=/backup/cnflBackup/`date +%Y%m%d-%H%M%S`

rm -rf $CNFL/temp/*
mkdir $CNFL_BACKUP
mysqldump -uroot -p<password> confluence|gzip >
$CNFL_BACKUP/confluence.mysql.data.gz
tar -cjvf $CNFL_BACKUP/data.bzip $CNFL > $CNFL_BACKUP/homedir.status
```

#### Backup by Date - Postgres

```
export d=`date +%u`
mkdir -p /home/backup/postgres/$d

sudo -u postgres pg_dumpall | bzip2 > /home/backup/postgres/$d/sql.bz2
```

#### Related Topics

- [Site Backup and Restore](#)
- [Backup FAQ](#)

#### Manually Backing Up the Site

Confluence is configured to back up its data automatically, as a zipped XML file. You can also manually perform this backup from the **Administrati on Console**.

You need to have [System Administrator](#) permissions in order to perform this function.

#### Note:

- Atlassian recommends that you follow the [Pro duction backup strategy](#) if your Confluence site is large or you are encountering problems with your automated backup.
- Plugins are not included in the XML backup. After importing your backup into a new Confluence site, you will need to re-install all plugins (add-ons) that are not bundled with Confluence. (The plugindata table is not backed up in a manual backup.)

#### Related pages:

- [Restoring a Site](#)
- [Configuring Backups](#)
- [Production Backup Strategy](#)

## Creating the site backup

### To manually back up your site:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Backup & Restore** in the left-hand panel.
3. Choose **Archive to backups folder** to store a copy of the backup in the same folder as Confluence's backups.  
If you do not archive the backup it will be made available for you to download, and then deleted from the server after 24 hours.
4. Choose **Backup attachments** to include attachments in your backup.
5. Choose **Backup**.

The process will take a few minutes.

## Retrieving the Backup File

Confluence stores the backup as a zipped XML file in the 'backups' directory under the Confluence Home directory on your Confluence server. To find your Confluence Home directory, see the [documentation](#). You will need access to the Confluence server in order to retrieve this file.

## Enabling the download of the backup file via the administration console

By default, it is not possible to retrieve the backup file via the Confluence Administration Console. This feature is disabled for security reasons.

Administrators can enable this functionality by updating the relevant configuration property as described below. If this functionality is enabled, Confluence will prompt you to download the backup file when the backup process finished. However, we recommend that you turn the feature **off** in production environments.

### To enable download of the backup file from the Administration Console:

1. Edit the `confluence.cfg.xml` file found in the [Confluence Home and other important directories](#).
2. Set the value of property `admin.ui.allow.manual.backup.download` to 'true' (without the quotation marks).
3. Restart Confluence.

If the value of the above configuration property is 'true', it will be possible to download the backup file after manually backing up the site via the Confluence Administration Console. If the value of this property is 'false' or the property is not present in the configuration file, you will need to retrieve the backup file from the file system on the Confluence server. By default, the value is 'false'.

## Notes

If you experience timeout errors, please consider bypassing Apache and creating the export directly from Tomcat. This will speed up the process and prevent timeouts. For example, your URL might be something like `http://<domain>.com`. To bypass this and access Tomcat directly, use this URL: `http://localhost:8080/confluence/admin/backup.action`.

## Restoring a Site

**CAUTION:** Restoring a backup of an entire confluence site will:

- Wipe out all Confluence content in the database. Please ensure that your database is backed up before you start.
- Log you out after the restoration process. Please make sure you know the login details contained in the data

### Related pages:

- [Production Backup Strategy](#)
- [Manually Backing Up the Site](#)

that you are about to restore.

This page describes how to restore data from an XML backup file into an existing Confluence installation. If you want to restore data into a new site, follow the instructions on [restoring from backup during setup](#).

You need [System Administrator](#) permissions in order to perform this function.

#### Notes before you start:

- **All content replaced.** Restoring a site from backup will replace all your content, as described in the warning above.
- **Selective space restoration not possible.** You cannot select a single space to restore from the entire site backup when the backup contains more than one space.
- **Version compatibility.** Confluence supports importing site and [space](#) backups into the same major version only. In Confluence, the major version is up to the second number, eg 5.4.3's major version is 5.4. Importing an old XML backup file to a new major version (for example, Confluence 5.5 to Confluence 5.6) is not supported.

|                         | Import into 5.4.x | Import into 5.5.x | Import into 5.6.x |
|-------------------------|-------------------|-------------------|-------------------|
| Site backup from 5.5.1  |                   |                   |                   |
| Space backup from 5.4.2 |                   |                   |                   |

- **XML backups should not be used to upgrade Confluence.** Upgrade Confluence by following: [Upgrading Confluence](#).

#### Restoring data from an XML backup

You can restore data from an XML backup file located somewhere on your local computer or a shared drive, or you can copy the XML file into the Confluence installation and restore it from there. The second option is recommended for large backup files. Both options are described below.

#### To restore data from an XML backup located outside Confluence:


1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Backup and Restore** in the left-hand panel.
3. Choose **Choose File** and browse for the backup file.
4. Uncheck **Build Index** if you want to create the index at a later stage.
5. Choose **Upload and Restore**.

#### To restore data from an XML backup located in your Confluence installation:

1. Copy your XML backup zip file into the `restore` directory in your Confluence home directory. For example:
  - On UNIX: `/opt/java/src/confluence/deployments/conf.atlassian.com/home/restore`
  - On Windows: `C:\Program Files\Atlassian\Application Data\Confluence x.x\restore`
2. Choose the **cog icon**


, then choose **General Configuration** under Confluence Administration
3. Choose **Backup and Restore** in the left-hand panel.
4. The zip file that you copied in step 1 will appear in the list of files under the heading **Restore a backup from the Confluence Home Directory** on your Confluence Administration Console. Select the zip file.
5. Uncheck **Build Index** if you want to create the index at a later stage.
6. Choose **Restore**.

## Notes

- **Production backup strategy preferred.** Atlassian recommends that you follow the [Production Backup Strategy](#) for your production Confluence site, because Confluence XML backups are not recommended for non-evaluation sites.
- **Restoring from other backups.** If your daily backup zip files cannot be restored for some reason, but you have backups of both your database and your Confluence home directory, then it is still possible to [restore from these backups](#).

## Restoring a Space

You can [export a space](#), including pages, comments and attachments, to a zip that contains an XML file and, optionally, all the attachments in the space. To import the space to another Confluence site, restore the zip as described below.

You need to have [System Administrator](#) permissions in order to restore a space from an XML zip file.

## Export and import compatibility

Confluence only accepts space exports from compatible versions. To find out which versions your current Confluence version can accept space exports from, go to



> **General Configuration** > **Backup and Restore**.

You may not be able to import spaces into earlier versions of Confluence; the **Backup and Restore** page in the later Confluence version will give you some indication of whether it is likely to be compatible.

If you need to import a space from an incompatible version, see the workaround described [below](#).

## On this page:

- [Export and import compatibility](#)
- [Restoring a space from an XML backup](#)
- [Workaround for restoring spaces from Confluence 5.3 and below](#)

## Related pages:

- [Restoring a Site](#)
- [Confluence Administrator's Guide](#)

Also, Confluence won't allow you to import a space if there's already a space with the same space key. If you already have a space with the same key, you'll need to delete the existing space before importing the new one.

## Restoring a space from an XML backup

You can restore data from an XML backup file located somewhere on your local computer or a shared drive, or you can copy the XML file into the Confluence restore directory and restore it from there. We recommend the second option for large backup files.

## To restore data from an XML backup located outside Confluence:

1. Choose the **cog icon**




, then choose **General Configuration** under Confluence Administration

2. Choose **Backup and Restore** in the left-hand panel
3. Choose **Choose File** and browse for the backup file
4. Uncheck **Build Index** if you want to create the index at a later stage
5. Choose **Upload and Restore**

## To restore data from an XML backup located in your Confluence installation:

1. Copy your XML backup zip file into the `restore` directory in your Confluence home directory. For example:
  - On UNIX: `/opt/java/src/confluence/deployments/conf.atlassian.com/home/r`

- On Windows: C:\Program Files\Atlassian\Application Data\Confluence x.x\restore
2. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration
  3. Choose **Backup and Restore** in the left-hand panel
  4. Select the zip file that you copied in step 1 from the list of files under the heading **Restore a backup from the Confluence Home Directory**
  5. Uncheck **Build Index** if you want to create the index at a later stage
  6. Choose **Restore**

#### Workaround for restoring spaces from Confluence 5.3 and below

If you need to import a space from a version prior to Confluence 5.3, you can use a temporary Confluence installation to upgrade the space export to the right version number:


1. Download the same version of Confluence as the version you exported the space from (you can get older versions of Confluence at the [Confluence Downloads Archive](#))
2. Install that version of Confluence on a temporary server
3. Import the space into this temporary Confluence site
4. Upgrade Confluence on your temporary site to the same version as the site where you want to import the space (see [Upgrading Confluence](#) for instructions)
5. Export the space from your temporary Confluence site (it'll now have the right version number)
6. Import the space into your production Confluence site

#### Restoring a Test Instance from Production

See [Migrating Confluence Between Servers](#) for a more comprehensive explanation.

Many Confluence administrators will have a production instance running the "live" version of Confluence, as well as a test instance for testing upgrades and so on. In this situation, it's quite common that the two instances are running different versions of Confluence. This document describes how to copy the data from a production instance to a test instance, where the production version may be different to the test version.

Before proceeding with this guide, ensure you have read and understood the [normal procedure for upgrading Confluence](#).

 *The information on this page **does not apply** to Confluence Cloud.*

#### Upgrading a test Confluence instance with production data

Essentially, we are copying both the production home directory and database to the test instance. We then update the database details on the test instance to point to the test database, leaving all other instance metadata (most importantly the Confluence *build number*) the same as production.

1. Shut down your test instance.
2. Restore the production database to the test database server.
3. Create a backup of the `confluence.cfg.xml` file found in the home directory of the test instance.
4. Copy the production confluence-home directory to the test application server.
5. Open the `confluence.cfg.xml` which has been copied in a text editor. Change the database settings to match the test database server. **Ensure you do not point to your production database.** (You can compare with the backup you made in Step 3 if you need to get the database settings. Don't just copy this file – you need the build number unchanged from production to indicate the database is from an older version of Confluence.)

Before starting your test instance, you need to do the following steps to ensure no contact with production systems.

### Ensuring no contact with production systems

To ensure no contact with external systems, you will need to disable both inbound and outbound mail services.

1. Disable global outbound mail by running the following database query:

```
SELECT * FROM BANDANA WHERE BANDANAKEY =
'atlassian.confluence.smtp.mail.accounts';
```

2. Disable space-level mail archiving by running the following database query:

```
SELECT * FROM BANDANA WHERE BANDANAKEY =
'atlassian.confluence.space.mailaccounts';
```

Change the 'SELECT \*' to a 'DELETE' in the above queries once you are sure you want to remove the specified accounts.

Once this is done, you can start your test instance without any mails being sent or retrieved. Think carefully about other plugins which may access production systems (SQL macro, etc.). These should be disabled promptly after starting the test instance.

You can create a [developer](#) license for this server and update the License Details after starting up.

### See also

[Upgrading Confluence](#)

[Migrating Confluence Between Servers](#)

[Restoring to a Test Instance of Confluence from Production](#)

### Restoring Data from other Backups

Typically, Confluence data is restored from the [Administration Console](#) or from the [Confluence Setup Wizard](#).

If you are experiencing problems restoring from a zipped XML backup file, it is still possible to restore provided you have:

1. A backup of your home directory.
2. A backup of your database (if you're using an external database).

Instructions for this method of restoring differ depending on whether you are using the embedded database or an external database (like Oracle, MS SQL Server, MySQL or Postgres).

### Embedded Database

If you are running against the embedded database, the database is located inside the `database` folder of your Confluence Home Directory. Hence, all you need to do is:

1. Retrieve the most recent backup of your home directory.
2. Unpack the Confluence distribution and point the `confluence-init.properties` file to this directory.

### External Database

*If you're using an external database, you need to do the following.*


1. Prepare backups of your home directory and database (preferably backups that are dated the same). That is, make sure the home directory is accessible on the filesystem and the database available to be connected to.
2. If this database happens to have a different name, or is on a different server, you need to modify the `jdbc` url in the `confluence.cfg.xml` file inside the Confluence Home Directory. The value of this property is specified as `hibernate.connection.url`.
3. Unpack the Confluence distribution and point the `confluence-init.properties` file to the home directory.

### Retrieving File Attachments from a Backup

File attachments on pages can be retrieved from a backup without needing to import the backup into Confluence. This is useful for recovering attachments that have been deleted by users.

Both [automated](#) and [manual](#) backups allow this, as long as the 'Include attachments' property was set. If you want to restore pages, spaces or sites, see the [Confluence Administrator's Guide](#) instead.

Before following the instructions for recovering attachments below, we will review how backups store file and page information.

 *The information on this page **does not apply to** Confluence Cloud.*

### How Backups Store File and Page Information

The backup zip file contains `entities.xml`, an XML file containing the Confluence content, and a directory for storing attachments.

#### Backup Zip File Structure

Page attachments are stored under the attachments directory by page and attachment id. Here is an example listing:

```
Listing for test-2006033012_00_00.zip
\attachments\98\10001
\attachments\98\10002
\attachments\99\10001
entities.xml
```

Inside the attachment directory, each numbered directory inside is one page, and the numbered file inside is one attachment. The directory number is the page id, and the file number is the attachment id. For example, the file `\attachments\98\10001` is an attachment with page id 98 and attachment id 10001. You can read `entities.xml` to link those numbers to the original filename. `Entities.xml` also links each page id to the page title.

#### Entities.xml Attachment Object

Inside the `entities.xml` is an Attachment object written in XML. In this example, the page id is [98](#), the attachment id is [10001](#) and the filename is [myimportantfile.doc](#). The rest of the XML can be ignored:

```
<object class="Attachment" package="com.atlassian.confluence.pages">
 <id name="id">98</id>
 <property name="fileName"><![CDATA[myimportantfile.doc]]></property>
 ...
 <property name="content" class="Page" package="com.atlassian.confluence.pages"><id
name="id">10001</id>
</property>
 ...
</object>
```

#### Entities.xml Page Object

This XML describes a page. In this example, the page id is [98](#) and the title is [Editing Your Files](#). The rest of the XML can be ignored:

```
<object class="Page" package="com.atlassian.confluence.pages">
<id name="id">98</id>
<property name="title"><![CDATA[Editing Your Files]]></property>
...
</object>
```

### Instructions for Recovering Attachments

Each file must be individually renamed and re-uploaded back into Confluence by following the instructions below. Choose one of the three methods:

#### Choice A - Recover Attachments By Filename

Best if you know each filename you need to restore, especially if you want just a few files:

1. Unzip the backup directory and open entities.xml.
2. Search entities.xml for the filename and find the attachment object with that filename. Locate its page and attachment id.
3. Using the page and attachment id from entities.xml, go to the attachments directory and open that directory with that page id. Locate the file with the attachment id.
4. Rename the file to the original filename and test it.
5. Repeat for each file.
6. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

#### Choice B - Restore Files By Page

Best if you only want to restore attachments for certain pages:

1. Unzip the backup directory and open entities.xml.
2. Search entities.xml for the page title and find the page object with that title. Locate its page id.
3. Go to the attachments directory and open that directory with that page id. Each of the files in the directory is an attachment that must be renamed.
4. Search entities.xml for attachment objects with that page id. Every attachment object for the page will have an attachment id and filename.
5. Rename the file with that attachment id to the original filename and test it.
6. Repeat for each page.
7. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

#### Choice C - Restore All Files

Best if you have a small backup but want to restore many or all the attachments inside:

Following process is applicable to **space** export only. Site xml backups do not require page id to be updated manually due to the nature of persistent page\_id's.

1. Unzip the backup directory and open entities.xml.
2. Go to the attachments directory and open any directory. The directory name is a page id. Each of the files in the directory is an attachment that must be renamed.
3. Search entities.xml for attachment objects with that page id. When one is found, locate the attachment id and filename.
4. Rename the file with that attachment id to the original filename and test it.
5. Find the next attachment id and rename it. Repeat for each file in the directory.
6. Once all files in the current directory are renamed to their original filenames, search entities.xml for the page id, eg directory name. Find the page object with that page id and locate its page title.
7. Rename the directory to the page title and move on to the next directory. Repeat for each un-renamed directory in the attachments directory.
8. To import each file back into Confluence, upload to the original page by attaching the file from within

Confluence.

## Troubleshooting failed XML site backups

XML site backups are only necessary for migrating to a new database. [Setting up a test server](#) or [Establishing a reliable backup strategy](#) is better done with an SQL dump.

### Related pages:

- [Enabling detailed SQL logging](#)

Seeing an error when creating or importing a backup?

Problem	Solution
Exception while creating backup	Follow instructions below
Exception while importing backup	Follow <a href="#">Troubleshooting XML backups that fail on restore</a> instead

## Resolve Errors With Creating An XML Backup

The errors may be caused by a slightly corrupt database. If you're seeing errors such as '[Couldn't backup database data](#)' in your logs, this guide will help you correct the error on your own. We strongly recommend that you backup your database and your Confluence home directory beforehand, so that you can [restore your site from those](#) if required. If you are unfamiliar with SQL, we suggest you contact your database administrator for assistance.

### Preferable solution

The [Production Backup Strategy](#) is a very reliable and more efficient way to do backups. If you are running into problems with XML backups - whether memory related or because of problems like the one described here - use the native backup tool as an alternate solution.

### To Identify And Correct The Problem

To work out where the data corruption or problems are, increase the status information reported during backup, then edit the invalid database entry:

1. Stop Confluence.
2. If you have an external database, use a database administration tool to create a manual database backup.
3. Backup your Confluence home directory. You will be able to [restore your whole site](#) using this and the database backup.
4. Open the `my_confluence_install/confluence/WEB-INF/classes/log4j.properties` and add this to the bottom and save:

```
log4j.logger.com.atlassian.confluence.importexport.impl.XMLDatabinder=DEBUG
, confluencelog
log4j.additivity.com.atlassian.confluence.importexport.impl.XMLDatabinder=false
```

5. Find your [atlassian-confluence.log](#). Move or delete all existing Confluence logs to make it easier to find the relevant logging output.
6. Restart Confluence and login.
7. Begin a backup so that the error reoccurs.
8. You must now check your log files to find out what object could not be converted into XML format. Open `confluence-home/logs/atlassian-confluence.log`. Scroll to the bottom of the file.
9. Do a search for 'ObjectNotFoundException'. You should see an error similar to this:

```

01 2005-08-24 00:00:33,743 DEBUG
[DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
object: com.atlassian.confluence.core.ContentPermission with ID:
5 to XML.
02 2005-08-24 00:00:33,743 DEBUG
[DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: type
03 2005-08-24 00:00:33,743 DEBUG
[DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: group
04 2005-08-24 00:00:33,743 DEBUG
[DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: expiry
05 2005-08-24 00:00:33,743 DEBUG
[DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: content
06 [DOCPRIV2:ERROR] LazyInitializer - Exception initializing
proxy <net.sf.hibernate.ObjectNotFoundException: No row with the
given identifier exists: 2535,
07 of class:
com.atlassian.confluence.core.ContentEntityObject>net.sf.hiberna
te.ObjectNotFoundException:
08 No row with the given identifier exists: 2535, of class:
com.atlassian.confluence.core.ContentEntityObject
09 at
net.sf.hibernate.ObjectNotFoundException.throwIfNull(ObjectNotFo
undException.java:24)
10 at
net.sf.hibernate.impl.SessionImpl.immediateLoad(SessionImpl.java
:1946)
11 at
net.sf.hibernate.proxy.LazyInitializer.initialize(LazyInitializ
er.java:53)
12 at
net.sf.hibernate.proxy.LazyInitializer.initializeWrapExceptions(
LazyInitializer.java:60)
13 at
net.sf.hibernate.proxy.LazyInitializer.getImplementation(LazyIni
tializer.java:164)
14 at
net.sf.hibernate.proxy.CGLIBLazyInitializer.intercept(CGLIBLazyI
nitializer.java:108)
15 at
com.atlassian.confluence.core.ContentEntityObject$$EnhancerByCGL
IB$$cc2f5557.hashCode(<generated>)
16 at java.util.HashMap.hash(HashMap.java:261)
17 at java.util.HashMap.containsKey(HashMap.java:339)
18 at
com.atlassian.confluence.importexport.impl.XMLDatabinder.toGener
icXML(XMLDatabinder.java:155)

```

10. Open a DBA tool such as DbVisualizer and connect to your database instance. Scan the table names in the schema. You will have to modify a row in one of these tables.
11. To work out which table, open `catalina.out`, check the first line of the exception. This says there

was an error writing the `ContentPermission` object with id 5 into XML. This translates as *the row with primary key 5 in the CONTENTLOCK table* needs fixing. To work out what table an object maps to in the database, here's a rough guide:

- Pages, blogposts, comments --> CONTENT table
  - attachments --> ATTACHMENTS table
  - More information can be found in the [schema documentation](#)
12. Now you must find the primary key of the incorrect row in this table. In this case, you can check the first line and see that the row has a primary key of 5.
  13. Each property is written to a column, so the last property that was being written has the incorrect value. The row being written to when the exception was thrown was `CONTENT` (line 5) with a value of 2535 (line 6). Now you know the column and value. This value 2535 is the id of an entry that no longer exists.
  14. Using a database administrative tool, login to the Confluence database. Locate the row in the relevant table and correct the entry. Check other rows in the table for the default column value, which may be null, 0 or blank. Overwrite the invalid row value with the default.
  15. Restart Confluence.
  16. Attempt the backup again. If the backup fails and you are stuck, please [lodge a support request](#) with your latest logs.

### Troubleshooting "Duplicate Key" related problems

If you are encountering an error message such as:

```
could not insert:
[bucket.user.propertyset.BucketPropertysetItem#bucket.user.propertyse
t.BucketPropertysetItem@a70067d3]; SQL []; Violation of PRIMARY KEY
constraint 'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate key
in object 'OS_PROPERTYENTRY'.; nested exception is
java.sql.SQLException: Violation of PRIMARY KEY constraint
'PKOS_PROPERTYENTRY_314D4EA8'. Cannot insert duplicate key in object
'OS_PROPERTYENTRY'.
```

this indicates that the Primary Key constraint 'PK\_OS\_PROPERTYENTRY\_314D4EA8' has duplicate entries in table 'OS\_PROPERTYENTRY'.

You can locate the constraint key referring to 'PK\_OS\_PROPERTYENTRY\_314D4EA8' in your table 'OS\_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS\_PROPERTYENTRY' table is:

```
SELECT ENTITY_NAME, ENTITY_ID, ENTITY_KEY, COUNT(*) FROM
OS_PROPERTYENTRY GROUP BY ENTITY_NAME, ENTITY_ID, ENTITY_KEY HAVING
COUNT(*) > 1
```

### To Help Prevent This Issue From Reoccurring

1. If you are using the embedded database, be aware that it is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss, which is why an external database is recommended for production use. You should [migrate to an external database](#).
2. If you are using an older version of Confluence than [the latest](#), you should consider [upgrading](#) at this point.

### Troubleshooting XML backups that fail on restore

XML site backups are only necessary for migrating to a new database. [Upgrading Confluence](#), [Setting up a test server](#) or [Production Backup Strategy](#) is better done with an SQL dump.

Seeing an error when creating or importing a site or space backup?

Problem	Solution
Exception while creating backup	See <a href="#">Troubleshooting failed XML site backups</a>
Exception while importing backup	See instructions below

#### On this page:

- [Resolve Errors When Attempting To Restore An XML Backup](#)
  - [Troubleshooting "Duplicate Entry" for key "cp\\_" or "cps\\_"](#)
  - [Troubleshooting "Duplicate Key" related problems](#)
  - [Troubleshooting "net.sf.hibernate.PropertyValueException: not-null" related problems](#)
  - [To Help Prevent this Issue from Recurring](#)

#### Related Topics:

- [Troubleshooting failed XML site backups](#)

### Resolve Errors When Attempting To Restore An XML Backup

The errors may be caused by a slightly corrupt database. You will need to find the XML backup file entry that is violating the DB rules, modify the entry and recreate the XML backup:

1. On the instance being restored, follow the instructions to disable batched updates (for simpler debugging), log SQL queries and log SQL queries **with parameters** at [Enabling Detailed SQL Logging](#).
2. Once all three changes have been made, restart Confluence.
3. Attempt another restore.
4. Once the restore fails, check your log files to find out what object could not be converted into XML format. For Confluence distribution users, check your Confluence install directory under the `/logs/` and check both `atlassian-confluence.log` and `catalina.out` file. The correct file will contain SQL debug output.
5. Scroll to the bottom of the file and identify the last error relating to a violation of the database constraint. For example:

```
2006-07-13 09:32:33,372 ERROR
[confluence.importexport.impl.ReverseDatabinder] endElement
net.sf.hibernate.exception.ConstraintViolationException:
 could not insert:
[com.atlassian.confluence.pages.Attachment#38]
net.sf.hibernate.exception.ConstraintViolationException: could
not insert: [com.atlassian.confluence.pages.Attachment#38]
...
Caused by: java.sql.SQLException: ORA-01400: cannot insert NULL
into ("CONFUSER"."ATTACHMENTS"."TITLE")
at
oracle.jdbc.driver.DatabaseError.throwSQLException(DatabaseError
.java:112)
at oracle.jdbc.driver.T4CTTIoer.processError(T4CTTIoer.java:331)
at oracle.jdbc.driver.T4CTTIoer.processError(T4CTTIoer.java:288)
```

This example indicates a row in your attachment table with ID = 38 that has a null title.

6. Go to the server that the backup was created on. You must have a copy of the database from which the backup was created. If you do not have this, use a DBA tool to restore a manual backup of the database.
7. Open a DBA tool and connect to the original database instance and scan the table names in the schema. You will have to modify a row in one of these tables.

8. To work out which table, open `catalina.out`, check the first line of the exception. To work out what table an object maps to in the database, here's a rough guide:
  - Pages, blogposts, comments --> CONTENT table.
  - attachments --> ATTACHMENTS table.
9. To correct the example error, go to the attachment table and find that attachment object with id 38. This will have a null title. Give a title using the other attachments titles as a guide. You may have a different error and should modify the database accordingly.
10. Once the entry has been corrected, create the XML backup again.
11. Import the backup into the new version.
12. If the import succeeds, revert the changes made in your [SQL logging](#) to re-enable disabled batched updates and turn off log SQL queries and log SQL queries with parameters.
13. Restart Confluence.

#### Troubleshooting "Duplicate Entry" for key "cp\_" or "cps\_"

If you are encountering an error message such as:

```
com.atlassian.confluence.importexport.ImportExportException: Unable
to complete import because the data does not match the constraints in
the Confluence schema. Cause:
MySQLIntegrityConstraintViolationException: Duplicate entry
'1475804-Edit' for key 'cps_unique_type'
```

This indicates that the XML export came from a version of Confluence with a corrupt permissions database, caused by some 3rd party plugin. This is an issue that was fixed when [CONF-22123](#) was implemented in Confluence 3.5.2. The simplest workaround is to export the space again after upgrading the instance to 3.5.2 or above. If that is not an option, then either the export will need to be edited manually to remove the duplicate permission entries or the source instance will need to have the offending entries removed. The following SQL queries can be used to look for such entries:

```

SELECT * FROM CONTENT_PERM WHERE USERNAME IS NULL AND GROUPNAME IS NULL;

SELECT cp.ID, cp.CP_TYPE, cp.USERNAME, cp.GROUPNAME, cp.CPS_ID, cp.CREATOR,
cp.CREATIONDATE, cp.LASTMODIFIER, cp.LASTMODDATE
FROM CONTENT_PERM cp
WHERE cp.USERNAME IS NOT NULL AND cp.GROUPNAME IS NOT NULL;

SELECT cps1.ID, cps1.CONTENT_ID, cps1.CONT_PERM_TYPE FROM CONTENT_PERM_SET cps1,
CONTENT_PERM_SET cps2
WHERE cps1.ID <> cps2.ID AND
cps1.CONTENT_ID = cps2.CONTENT_ID AND
cps1.CONT_PERM_TYPE = cps2.CONT_PERM_TYPE
ORDER BY cps1.CONTENT_ID, cps1.CONT_PERM_TYPE, cps1.CREATIONDATE ASC;

SELECT cp.ID, cp.CP_TYPE, cps.CONTENT_ID,
(SELECT scps.ID FROM CONTENT_PERM_SET scps WHERE scps.CONTENT_ID = cps.CONTENT_ID
AND scps.CONT_PERM_TYPE = cp.CP_TYPE) AS suggested_cps_id
FROM CONTENT_PERM cp, CONTENT_PERM_SET cps
WHERE cp.CPS_ID = cps.ID AND
cp.CP_TYPE <> cps.CONT_PERM_TYPE;

SELECT DISTINCT cp1.ID, cp1.CP_TYPE, cp1.USERNAME, cp1.GROUPNAME, cp1.CPS_ID,
cp1.CREATOR, cp1.CREATIONDATE, cp1.LASTMODIFIER, cp1.LASTMODDATE
FROM CONTENT_PERM cp1, CONTENT_PERM_SET cps1, CONTENT_PERM cp2, CONTENT_PERM_SET
cps2
WHERE
cp1.CPS_ID = cps1.ID AND
cp2.CPS_ID = cps2.ID AND
cp1.ID <> cp2.ID AND
cps1.CONTENT_ID = cps2.CONTENT_ID AND
cp1.CP_TYPE = cp2.CP_TYPE AND
cp1.USERNAME = cp2.USERNAME
ORDER BY cp1.CPS_ID, cp1.CP_TYPE, cp1.USERNAME, cp1.CREATIONDATE;

SELECT DISTINCT cp1.ID, cp1.CP_TYPE, cp1.USERNAME, cp1.GROUPNAME, cp1.CPS_ID,
cp1.CREATOR, cp1.CREATIONDATE, cp1.LASTMODIFIER, cp1.LASTMODDATE
FROM CONTENT_PERM cp1, CONTENT_PERM_SET cps1, CONTENT_PERM cp2, CONTENT_PERM_SET
cps2
WHERE
cp1.CPS_ID = cps1.ID AND
cp2.CPS_ID = cps2.ID AND
cp1.ID <> cp2.ID AND
cps1.CONTENT_ID = cps2.CONTENT_ID AND
cp1.CP_TYPE = cp2.CP_TYPE AND
cp1.GROUPNAME = cp2.GROUPNAME
ORDER BY cp1.CPS_ID, cp1.CP_TYPE, cp1.GROUPNAME, cp1.CREATIONDATE;

SELECT * FROM CONTENT_PERM_SET
WHERE ID NOT IN (SELECT DISTINCT CPS_ID FROM CONTENT_PERM);

```

Remove all matching entries and perform the export again.

#### Troubleshooting "Duplicate Key" related problems

If you are encountering an error message such as:

```
could not insert:
[bucket.user.propertyset.BucketPropertysetItem#bucket.user.propertyse
t.BucketPropertysetItem@a70067d3]; SQL []; Violation of PRIMARY KEY
constraint 'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate key
in object 'OS_PROPERTYENTRY'.; nested exception is
java.sql.SQLException: Violation of PRIMARY KEY constraint
'PKOS_PROPERTYENTRY_314D4EA8'. Cannot insert duplicate key in object
'OS_PROPERTYENTRY'.
```

This indicates that the Primary Key constraint 'PK\_OS\_PROPERTYENTRY\_314D4EA8' has duplicate entries in table 'OS\_PROPERTYENTRY'.

You can locate the constraint key referring to 'PK\_OS\_PROPERTYENTRY\_314D4EA8' in your table 'OS\_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS\_PROPERTYENTRY' table is:

```
SELECT ENTITY_NAME, ENTITY_ID, ENTITY_KEY, COUNT(*) FROM
OS_PROPERTYENTRY GROUP BY ENTITY_NAME, ENTITY_ID, ENTITY_KEY HAVING
COUNT(*) > 1
```

#### Troubleshooting "net.sf.hibernate.PropertyValueException: not-null" related problems

If you're receiving a message like:

```
ERROR [Importing data task]
[confluence.importexport.impl.ReverseDatabinder] endElement
net.sf.hibernate.PropertyValueException: not-null property references
a null or transient value:
com.atlassian.user.impl.hibernate.DefaultHibernateUser.name
```

This means there's an unexpected null value in a table. In the above example, the error is in the name column in the USERS table. We've also seen them in the ATTACHMENTS table.

Remove the row with the null value, redo the xml export, and reimport.

#### To Help Prevent this Issue from Recurring

1. If you are using the embedded database, be aware that it is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss, which is why an external database is recommended for production use. You should [migrate to an external database](#).
2. If you are using an older version of Confluence than [the latest](#), you should consider [upgrading](#) at this point.

The problem with different settings for case sensitivity varies between databases. The case sensitivity of the database is usually set through the collation that it uses. Please vote on the existing [issue](#)

## Attachment Storage Configuration

Database attachment storage was deprecated in Confluence 5.5. If you currently store attachments in the database you will be able to continue to do so, but you will be unable to switch to storing attachments in the database in Confluence 5.5 or later.

System Administrators can configure where

Confluence stores attachments. Attachments can be stored in a:

- File system - locally in the Confluence home directory, or
- Database - in Confluence's configured database (deprecated)

To configure Confluence attachment storage:

- Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

- Choose **Attachment Storage**.

#### On this page:

- [Attachment Storage Options](#)
  - [Local File System](#)
  - [Database \(deprecated\)](#)
- [Migrating between attachment storage systems](#)

#### Related pages:

- [Working with Confluence Logs](#)

## Attachment Storage Options

### Local File System

By default, Confluence stores attachments in the `attachments` directory within the configured Confluence home folder.

### Database (deprecated)

Confluence 5.4 and earlier gives administrators the option to store attachments in the database that Confluence is configured to use.

While storing attachments in the database can offer some advantages (such as ease of backup, and avoiding issues with some characters in attachment filenames), please be aware that the amount of space used by the database will increase because of greater storage requirements.

## Migrating between attachment storage systems

You can migrate your attachments from one storage system to another. All existing attachments will be moved over to the new attachment storage system.

When the migration occurs, all other users will be locked out of the Confluence instance. This is to prevent modification of attachments while the migration occurs. Access will be restored as soon as the migration is complete.

When migrating attachments from your database to a filesystem, the attachments are removed from the database after migration. However, when migrating attachments from a filesystem to your database, the attachments remain on the filesystem after migration.

To improve logging during the migration, add the package `com.atlassian.confluence.pages.persis`  
`tence.dao` with level `DEBUG`. See [Configuring Logging](#) for more information.

**To migrate, follow the steps below:**

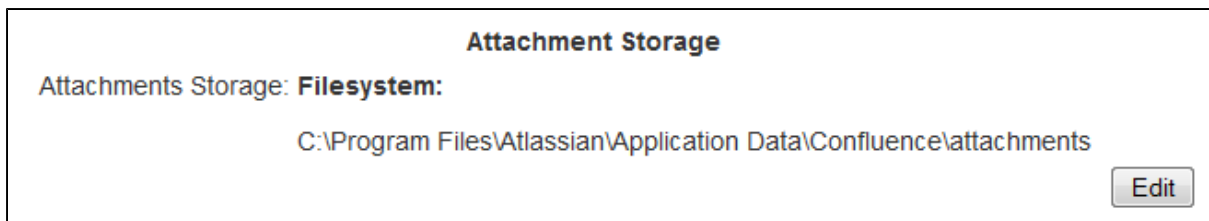
1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Click '**Attachment Storage**' in the left-hand panel. The current configuration will be displayed.

*Screenshot: Attachment storage configuration*



**Attachment Storage**

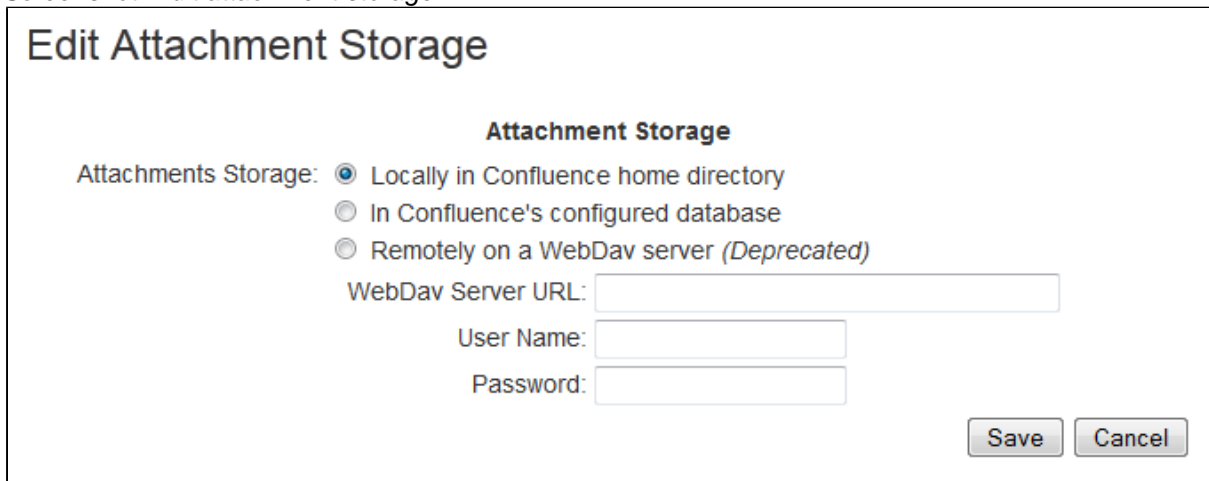
Attachments Storage: **Filesystem:**

C:\Program Files\Atlassian\Application Data\Confluence\attachments

Edit

3. Click the '**Edit**' button to modify the configuration.
4. Select the storage system you desire.

*Screenshot: Edit attachment storage*



**Edit Attachment Storage**

**Attachment Storage**

Attachments Storage: ☒ Locally in Confluence home directory  
☐ In Confluence's configured database  
☐ Remotely on a WebDav server (*Deprecated*)

WebDav Server URL:

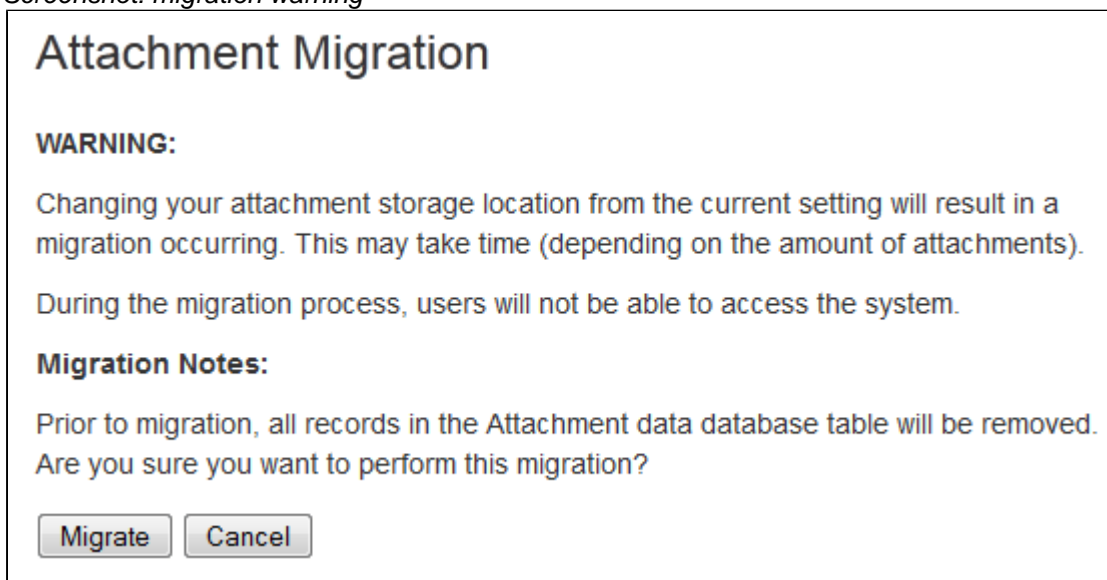
User Name:

Password:

Save Cancel

5. Click the '**Save**' button to save the changes.
6. A screen will appear, asking you to confirm your changes. Clicking 'Migrate' will take you to a screen that displays the progress of the migration.

*Screenshot: migration warning*



**Attachment Migration**

**WARNING:**

Changing your attachment storage location from the current setting will result in a migration occurring. This may take time (depending on the amount of attachments).

During the migration process, users will not be able to access the system.

**Migration Notes:**

Prior to migration, all records in the Attachment data database table will be removed.  
Are you sure you want to perform this migration?

Migrate Cancel

The following external website provides further information on migrating attachments from database to file system storage that you might find helpful - <http://www.scandio.de/2013/05/confluence-attachment-migration-the-safe-way-2/>.

### Configuring Attachment Size

Confluence gives you the option of limiting the maximum size of a single file attachment.

**To configure the maximum size allowed for an attachment:**

1. Go to



> **General Configuration.**

#### Related pages:

- [Recognised System Properties](#)
- [Files](#)

2. Choose **Edit**.
3. Enter the maximum size next to **Attachment Maximum Size**.  
The default is 100mb.
4. Choose **Save**.

### To configure the maximum 'index-able size of attachments':

By default, large attachment is defined as greater than 1 MB. The threshold for attachments that won't get excerpts can be modified using the system property `atlassian.indexing.contentbody.maxsize`, which takes a size in bytes.

For example to specify 250 kb you would use the following JVM parameter:

```
-Datlassian.indexing.contentbody.maxsize=256000
```

Limiting the size of attachment indexing has the following effects:

- Decreases the size of the index when large attachments are present.
- Prevent excerpts of large attachments being displayed in search results.

## Hierarchical File System Attachment Storage

The way attachments are stored changed significantly in Confluence 3.0. If you are upgrading from Confluence 2.10 or earlier see [Upgrading Confluence](#) for recommended upgrade paths, and read the version of the [Hierarchical File System Attachment Storage](#) page in our Confluence 3.0 documentation which provides more detail about migrating to the new file system structure.

Confluence stores attachments, such as files and images, in a file system. Confluence's attachment storage layout is designed to:

1. Limit the number of entries at any single level in a directory structure (as some file systems have a limit on the number of files that can be stored in a directory).
2. Partition attachments per space making it possible for a system admin to selectively back up attachments from particular spaces.

Attachments in Confluence have a number of identifying attributes: *content id of the file itself*, the *space id* and *content id of the page the file is attached to*. This means the file logically belongs to a piece of content which logically belongs in a space (not all content belongs to a space). For files within a space in Confluence, the directory structure is typically 8 levels, with the name of each directory level based on the following algorithm:

level	Derived From
1 (top)	Always 'ver003' indicating the Confluence version 3 storage format
2	The least significant 3 digits of the <i>space id</i> , modulo 250
3	The next 3 least significant digits of the <i>space id</i> , modulo 250
4	The full <i>space id</i>
5	The least significant 3 digits of the <i>content id</i> of the page the file is attached to, modulo 250
6	The next 3 least significant digits of the <i>content id</i> of the page the file is attached to, modulo 250
7	The full <i>content id</i> of the page the file is attached to
8	The full <i>content id</i> of the attached file

Within the 8th level will be a file for each version of that attachment, named to match the version number e.g. 1

An example:

A	B	C	D
file content id: 745644 space id: 800432 page content id: 632780	file content id: 782234 space id: 800432 page content id: 620002	file content id: 771250 space id: 810032 page content id: 603101	file content id: 771250 space id: 810032 page content id: 603101

```

graph TD
 Root[Attachments directory] --> ver003[ver003]
 ver003 --> nonspaced[nonspaced]
 ver003 --> 182[182]
 ver003 --> 32[32]
 nonspaced --> 242[242]
 242 --> 11[11]
 11 --> 511242[511242]
 511242 --> 701002[701002]
 182 --> 50[50]
 50 --> 800432[800432]
 800432 --> 30[30]
 800432 --> 2[2]
 30 --> 132[132]
 132 --> 632780[632780]
 632780 --> 745644[745644]
 2 --> 120[120]
 120 --> 620002[620002]
 620002 --> 782234[782234]
 32 --> 60[60]
 60 --> 810032[810032]
 810032 --> 101[101]
 101 --> 103[103]
 103 --> 603101[603101]
 603101 --> 771250[771250]

```

Attachments directory

ver003

nonspaced

182

32

242

11

511242

701002

50

800432

30

132

632780

745644

2

120

620002

782234

60

810032

101

103

603101

771250

top level directory partitioni

3 least significant digits of s

next 3 least significant digits

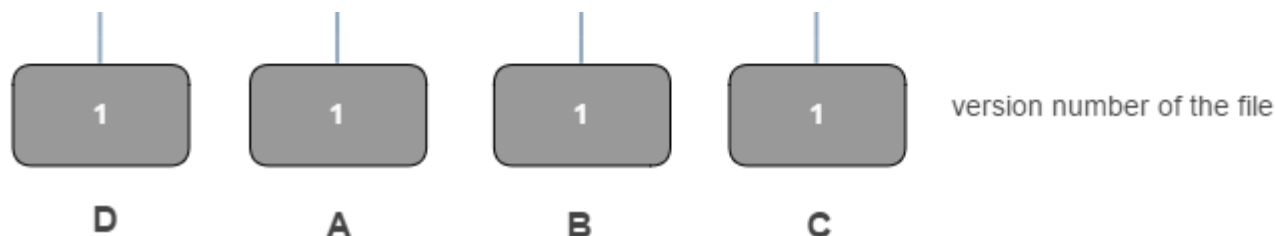
space id

3 least significant digits of ti  
content id % 250

next 3 least significant digits /  
page's content id % 250

content id of the page the fi

content id of the attached fi



To find the directory where attachments for a particular space are stored, go to `<confluence url>/admin/findspaceattachments.jsp` and enter a space key. It will return the directory on the file system where attachments for that space are stored.

File D in the above diagram is stored in a slightly different structure. Files that are not conceptually within a space replace the level 2 - 4 directories with a single directory called 'nonspaced'. Examples of such files are the global site logo and attachments on unsaved content.

## Confluence Data Model

This document provides a diagram of the Confluence schema and a conceptual overview of the data model.

### Notes:

- The Hibernate mapping files are the authoritative reference for the Confluence data model. These are the \*.hbm.xml files which you will find in the main Confluence JAR file (`<CONFLUENCE-INSTALLATION>\confluence\WEB-INF\lib\confluence-5.1.1.jar`).
- The tables, columns and other attributes are likely to change with each major release of Confluence. To find the exact DDL of your Confluence site, please run a query after installation.

### On this page:

- [Database diagrams](#)
- [Database tables and references](#)
- [Authentication](#)
- [Content](#)
- [Clustering](#)
- [System information](#)
- [Spaces](#)
- [Appearance](#)
- [Miscellaneous](#)

## Database diagrams

### Detailed diagrams

The following SVG images (Scalable Vector Graphics) include all the tables in the Confluence database. Click the links below to open the images in your browser, or download the SVG files for later use. You can use the browser's zoom (Ctrl++ or Cmd++) to see more detail in the diagrams:

- [ConfluenceTables-KeysOnly.svg](#) – Shows all tables, with primary keys only for each table.
- [ConfluenceTables-AllColumns.svg](#) – Shows all tables, and all columns for each table.

### Overview diagram

This image shows the core tables. Note that the image is very large. You may need to download it (right-click on the image) and view it in an image viewer. Alternatively, use the SVG images linked in the previous section.

▼ [Click here to show/hide the image...](#)



Expand the link below to see a table of the primary and foreign keys for each table.

▼ [Click here to show/hide the table...](#)

Primary key table name	Primary key column name	Foreign key table name	Foreign key column name	Foreign
AO_9412A1_AOUSER	ID	AO_9412A1_USER_APP_LINK	USER_ID	fk_ao_9412A1_
attachments	attachmentid	attachmentdata	attachmentid	fk9dc3e3f6
attachments	attachmentid	attachments	prevver	fk54475fb8
attachments	attachmentid	content_label	attachmentid	fkf0e743d1
attachments	attachmentid	imagedetails	attachmentid	fka76804c7
content	contentid	attachments	pageid	fk54475fb8
content	contentid	bodycontent	contentid	fka898d403
content	contentid	confancestors	ancestorid	fk9494e2d1

content	contentid	confancestors	descendentid	fk9494e2
content	contentid	content	prevver	fk6382c0
content	contentid	content	parentid	fk6382c0
content	contentid	content	parentcommentid	fk6382c0
content	contentid	content	pageid	fk6382c0
content	contentid	content_label	contentid	fkf0e743
content	contentid	content_perm_set	content_id	fkbf45a7
content	contentid	contentproperties	contentid	fk984c5e
content	contentid	extrnlInks	contentid	fk97c10f
content	contentid	likes	contentid	fk4514b5
content	contentid	links	contentid	fk451575
content	contentid	notifications	pageid	fk594acc
content	contentid	spaces	homepage	fk922824
content	contentid	spaces	spacedescid	fk922824
content	contentid	trackbacklinks	contentid	fkf6977a
content_perm_set	id	content_perm	cps_id	fkbd74b3
cwd_app_dir_mapping	id	cwd_app_dir_group_mapping	app_dir_mapping_id	fk_app_c
cwd_app_dir_mapping	id	cwd_app_dir_operation	app_dir_mapping_id	fk_app_c
cwd_application	id	cwd_app_dir_group_mapping	application_id	fk_app_c
cwd_application	id	cwd_app_dir_mapping	application_id	fk52050e
cwd_application	id	cwd_application_address	application_id	fk_applic
cwd_application	id	cwd_application_attribute	application_id	fk_applic
cwd_directory	id	cwd_app_dir_group_mapping	directory_id	fk_app_c
cwd_directory	id	cwd_app_dir_mapping	directory_id	fk_app_c
cwd_directory	id	cwd_directory_attribute	directory_id	fk_direct
cwd_directory	id	cwd_directory_operation	directory_id	fk_direct
cwd_directory	id	cwd_group	directory_id	fk_direct
cwd_directory	id	cwd_group_attribute	directory_id	fk_group
cwd_directory	id	cwd_user	directory_id	fk_user_
cwd_directory	id	cwd_user_attribute	directory_id	fk_user_
cwd_group	id	cwd_group_attribute	group_id	fk_group
cwd_group	id	cwd_membership	parent_id	fk_paren
cwd_group	id	cwd_membership	child_group_id	fk_child_
cwd_user	id	cwd_membership	child_user_id	fk_child_

cwd_user	id	cwd_user_attribute	user_id	fk_user_
cwd_user	id	cwd_user_credential_record	user_id	fk76f874
external_entities	id	external_members	extentityid	fk8c8d8
groups	id	external_members	groupid	fk8c8d8
groups	id	local_members	groupid	fk6b8fb4
keystore	keyid	trustedapp	public_key_id	fkddb119
label	labelid	content_label	labelid	fkf0e743
os_group	id	os_user_group	group_id	fk932472
os_user	id	os_user_group	user_id	fk932472
pagetemplates	templateid	content_label	pagetemplateid	fkf0e743
pagetemplates	templateid	pagetemplates	prevver	fkbc7ce9
spacegroups	spacegroupid	spacegrouppermissions	spacegroupid	fk487d9f
spacegroups	spacegroupid	spaces	spacegroupid	fk922824
spaces	spaceid	content	spaceid	fk6382cc
spaces	spaceid	notifications	spaceid	fk594acc
spaces	spaceid	pagetemplates	spaceid	fkbc7ce9
spaces	spaceid	spacepermissions	spaceid	fk433f23
trustedapp	trustedappid	trustedapprestriction	trustedappid	fke8496t
users	id	local_members	userid	fk6b8fb4

The following sections describe the principal tables involved in each logical area of Confluence – authentication, content, system information, and so on.

## Authentication

This section describes the tables involved in user authentication, which is implemented via the Atlassian Crowd framework embedded in Confluence.

Table	Description
cwd_user	Information for each user in Confluence.
cwd_group	The groups to which users can belong.
cwd_membership	Mapping the membership of users to groups.
cwd_directory	The user directories in your Confluence site. Examples of directories are the Confluence internal directory, or an LDAP directory.
cwd_application	The applications (JIRA, Confluence, and so on) defined in the authentication framework.

## Content

This section describes the tables involved in storing content. Content is the information that Confluence users are storing and sharing.

Table	Description
attachmentdata	The binary data for attached files. This table is only used when Confluence is configured to store attachments in the database. Otherwise, attachments are stored in the local file system.
attachments	Metadata for the files attached to Confluence pages.
bodycontent	The content of Confluence pages. No version information or other metadata is stored here. That is all in the <code>content</code> table.
content	A persistence table for the <code>ContentEntityObject</code> class of objects. The subclass is indicated by the <code>contenttype</code> column.
content_label	Arbitrary text labels for content.
label	The other half of the <code>content_label</code> system.
content_perm	Content-level permissions objects.
content_perm_set	A one-to-many mapping for content items and their permissions, with added metadata.
pagetemplates	The back end of the templates feature.
likes	The pages and other content liked by a particular user.
follow_connections	A mapping of users who are following other users.

## Clustering

The following table contains information about clustered Confluence sites.

Table	Description
clustersafety	Normally, this table only contains one row. The value of the <code>safetynumber</code> is what Confluence uses to find out whether another Confluence site is sharing its database without being part of the cluster.

## System information

These tables store data related to the status and configuration of the Confluence site.

Table	Description
confversion	Used by the upgrade system to determine what to expect from the database, so as to negotiate upgrades.
plugindata	A record of the plugins that have been installed, and when. data is a blob of the actual plugin JAR file. This is principally cluster-related.

## Spaces

This table is related to the management of spaces.

Table	Description
spaces	Information about the spaces themselves: key, human-friendly name and numeric ID.

## Appearance

The following table contains information about the look and feel of your Confluence site.

Table	Description
decorator	The custom display templates used to customise Velocity layouts.

## Miscellaneous

This section includes other tables worth commenting on.

Table	Description
os_propertyentry	Arbitrary association of entities and properties.
bandana	A catch-all persistence layer. This table contains things like user settings and space- and global-level configuration data, and is used as storage by plugins such as the Dynamic Task List plugin. Essentially, for storing arbitrary data that doesn't fit anywhere else.
extrnlinks	Referral links.
hibernate_unique_key	Used by the high/low ID generator – the subsystem which generates our primary keys.  If you interfere with this table, you may not be able to create objects in Confluence.
indexqueueentries	Manages full-content indexing across the system. The table generally contains the last 12 hours (approximately) of updates, to allow re-syncing of cluster nodes after restarts.
keystore	Used by the trusted apps framework to store the server's private key, and other servers' public keys.
links	Tracks links within the server (that is, across and within spaces).
notifications	Stores page- and space-level watches.
trackbacklinks	Trackback links.
confancestors	Used to speed up permissions checks, by allowing quick lookup of all a page's ancestors.

## Finding Unused Spaces

Sometimes, you want to know what is *not* being used. It's great to know what's getting most attention, but what about stagnant pages, or even entire spaces that are no longer active?

While [viewing space activity](#) can provide hints, it doesn't always provide enough detail. It is possible to find out this information directly from the database.

The following query identifies the last date on which content was modified in each space within a single Confluence instance:

```
SELECT spaces.spacename, MAX(content.lastmoddate)
FROM content, spaces
WHERE content.spaceid = spaces.spaceid
GROUP BY spaces.spacename;
```

It returns a list of space names, and the last date and time at which any content was added or changed.

Alternatively, this query identifies spaces where the content hasn't changed since a specified date:

```
SELECT spaces.spacename
FROM content, spaces
WHERE content.spaceid = spaces.spaceid
GROUP BY spaces.spacename
HAVING MAX(content.lastmoddate) < '2006-10-10';
```

The result is a simple list of space names.

## Data Import and Export

Confluence administrators and users can import data into Confluence from a number of sources. The permissions required differ, depending on the scope of the import. See [Import Content Into Confluence](#).

You can also export Confluence content to various formats. See [Export Content to Word, PDF, HTML and XML](#).

### Related pages:

- [Managing Confluence Data](#)
- [Confluence Administrator's Guide](#)

## Import a Text File

Confluence allows you to import text files from a directory on the Confluence server, and convert them into Confluence pages. Each file is imported as a separate Confluence page with the same name as the file.

- The text file may contain plain text, HTML or [Confluence Storage Format](#)
- You need to be part of the `confluence-administrators` group or a System Administrator to import text files
- You can import pages from disk into site spaces, but not into personal spaces  
Please see [Spaces](#) for information about differences between site spaces and personal spaces.

### Related pages:

- [Import Content Into Confluence](#)
- [Site Backup and Restore](#)

To make sure Confluence maintains the formatting of the text document, add `<pre>` to the beginning and `</pre>` to the end. This will let Confluence know that it should treat the text as pre-formatted.

If you're working in a Unix-like environment, you can add the opening and closing tags to all files in a particular directory by following these steps:

1. Go to the directory containing the files
2. Run the following command in the terminal:

```
for i in $(ls); do echo "<pre>" >> m$i; cat $i >> m$i; echo "</pre>" >> m$i; mv m$i $i; done
```

### To import text files:

1. Go to the space and choose **Space tools > Content Tools** from the bottom of the sidebar
2. Choose **Import**
3. Type the directory path into the **Import directory** box
4. Select **Trim file extensions** to remove file extensions from the page titles when converting the files to Confluence pages  
The Confluence pages will take their titles from the files' names (including their extensions). To avoid having page titles with a suffix like `.txt` check this box.
5. Select **Overwrite existing pages** if you want to replace existing Confluence pages with the same title

- with the one you're importing
6. Choose **Import**

If your space uses the Documentation theme:

1. Choose **Browse > Space Admin** from the header

**Note:** The **Space Admin** option appears only if you have [space admin permissions](#), or if you're part of the 'confluence-administrators' group.

2. Choose **Import** from the space administration options
3. Follow the steps above

Confluence will place the new pages at the root of the space.

*Screenshot: Importing text files*

You can use this action to import text files from a directory on the Confluence server.

These text files become pages in Confluence, with the following features:

- The page title is taken from the filename
- The content is the entire page body

Import directory

Trim file extensions ☐

Overwrite existing pages ☐

## Configuring a Confluence Environment

This section describes the external setup of your Confluence installation. It includes information on configuring the web server, application server, directories and files – everything to do with the environment that Confluence runs in. For guidelines on modifying settings inside the application, see [Configuring Confluence](#) instead.

Confluence is a J2EE web application. On the client side, users access Confluence primarily via a web browser.

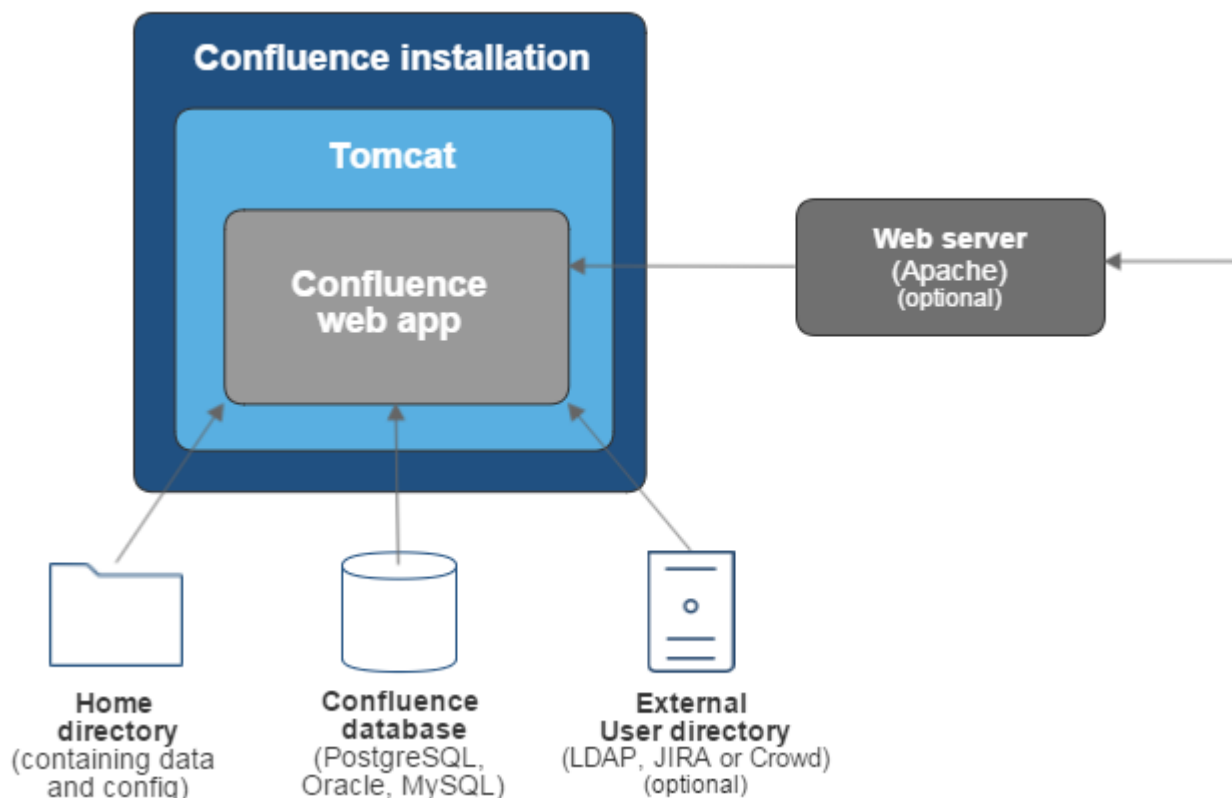
This section contains the following guidelines:

- [Confluence Home and other important directories](#)
- [Application Server Configuration](#)
- [Web Server Configuration](#)
- [Starting Confluence Automatically on System Startup](#)

*Diagram: A Confluence installation*

### Related pages:

- [Getting Started as Confluence Administrator](#)
- [Supported Platforms](#)



## Confluence Home and other important directories

### Confluence installation directory

The 'Confluence Installation directory' is the directory where Confluence was installed. This directory is also sometimes called the 'Confluence Install directory'.

Important files in the installation directory:

- `bin/setenv.bat` or `bin/setenv.sh`  
This file is used to edit `CATALINA_OPTS` memory and garbage collection settings and define system properties.
- `confluence/WEB-INF/classes/confluence-init.properties`  
This file contains the location of the Confluence Home directory.

### Confluence home directory

The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. Another term for 'Home directory' would be 'data directory'.

### Finding the home directory

The location of the Confluence home directory is defined when you install Confluence. This location is stored in the `confluence-init.properties` file, which is located in the `confluence/WEB-INF/classes` directory of your Confluence Installation directory.

When Confluence is running you can find the location of the home directory in

#### On this page:

- [Confluence installation directory](#)
- [Confluence home directory](#)
- [Changing the location of the home directory](#)
- [Database](#)
- [Temp directory](#)



## > General Configuration > System Information > Confluence Information - Confluence Home.

If your Confluence instance is clustered, you will also have a shared home directory which will contain some data (such as attachments and backups) that would otherwise reside in the home directory.

### Contents of the home directory

The Confluence home directory contains some of the configuration data used by Confluence. This section outlines the purpose of the files and directories in the Confluence home directory.

File or directory	Purpose
<code>confluence.cfg.xml</code>	This file contains all of the information necessary for Confluence to start up, such as: <ul style="list-style-type: none"> <li>• Product license</li> <li>• Context path</li> <li>• Database details, such as location and connection pool settings</li> <li>• Paths to important directories</li> </ul>
<code>attachments/</code>	This directory contains every version of each attachment stored in Confluence. You can specify an alternative directory for attachment storage by setting the <code>attachments.dir</code> property in <code>confluence.cfg.xml</code> .
<code>backups/</code>	Confluence will place its daily backup archives in this directory, as well as any manually generated backups. Backup files in this directory take the following form <code>daily-backup-YYYY_MM_DD.zip</code> . You can specify an alternative directory for backups by setting the <code>daily.backup.dir</code> property in <code>confluence.cfg.xml</code> .
<code>bundled-plugins/</code>	Confluence includes a set of <i>bundled</i> plugins. The <code>bundled-plugins</code> directory is where Confluence will unpack its bundled plugins when it starts up. This directory is refreshed on every restart, so removing a plugin from this directory will not uninstall the plugin, as it will be replaced the next time Confluence starts up.
<code>database/</code>	This is where Confluence stores its database when configured to run with the <a href="#">Embedded H2 Database</a> . In such cases this directory contains all Confluence runtime data. Installations configured to run using an external database such as MySQL will not use this directory. The H2 database is provided for evaluating Confluence and is not supported as a production database.
<code>index/</code>	The Confluence index is heavily used by the application for content searching and recently updated lists and is critical for a running Confluence instance. If data in this directory is lost or corrupted, it can be restored by running a full reindex from within Confluence. This process can take a long time depending on how much data is stored Confluence's database. An alternative directory may be specified for the index by setting the <code>lucene.index.dir</code> property in <code>confluence.cfg.xml</code> .
<code>journal/</code>	Entries are added to the journal when changes occur (such as a comment, like, new page). Journal entries are then processed and the entries added to the index (about every 5 seconds). In a cluster, the journal keeps the indexes on each node in sync.
<code>logs/</code>	Confluence's application logs are stored in this directory.

<code>plugin-cache/</code>	All Confluence plugins are stored in the database. To allow for quicker access to classes contained within the plugin JARs, Confluence will cache these plugins in the <code>plugin-cache</code> directory. This directory is updated as plugins are installed and uninstalled from the system and is completely repopulated from the database every time Confluence is restarted. Removing plugins from this directory does not uninstall them.
<code>resources/</code>	The <code>resources</code> directory stores any space logos used in your Confluence instance. Space logos are stored in directories named with the spacekey.
<code>temp/</code>	<p>The <code>temp</code> directory is used for runtime functions such as exporting, importing, file upload and indexing. Files in this directory are temporary and can be safely removed when Confluence is offline. A daily job within Confluence deletes files that are no longer needed.</p> <p>An alternative directory may be specified for temporary data by setting the <code>webwork.multipart.saveDir</code> property in <code>confluence.cfg.xml</code>.</p>
<code>thumbnails/</code>	Stores temporary files for image thumbnails. This directory is essentially a thumbnail cache, and files deleted from this directory will be regenerated the next time the image is accessed.

## Changing the location of the home directory

When Confluence first starts up, it reads the `confluence-init.properties` file to determine where to look for the Home directory.

To change the location of the home directory edit the `confluence.home` property in the `confluence-init.properties` file as follows:

- **Windows**

In Windows, the path `C:\confluence\data` would be written as:

```
confluence.home=C:/confluence/data
```

Note that all backslashes (\) are written as forward slashes (/)

- **Linux / Solaris**

On any Linux-based system, the property is defined using the normal directory syntax:

```
confluence.home=/var/confluence/
```

## Symbolic links

There can be no symbolic links within the Confluence home directory. You must define an absolute path. If disk space is an issue, place the entire `confluence.home` directory on a disk partition where there is enough space. The absolute path of generated files (such as exports) is compared with the absolute path of the `confluence.home` directory when constructing URLs. When a sub-directory has a different path, the URL will be incorrect, and you may receive "Page not found" errors. These measures are in place to prevent "directory traversal" attacks.

## Fixing the Confluence Configuration

The Confluence configuration file: `confluence-cfg.xml` inside the home directory may contain references to the original location of your Confluence home. You will need to edit this file to update these references to also point to the new location. The two properties in this file that need to change are:

- `daily.backup.dir` if you have not configured your backups to be placed elsewhere already
- `hibernate.connection.url` if you are using the embedded HSQL database.

## Database

All other data, including page content, is kept in the database. If you installed Confluence as a trial, or chose to use the [embedded HSQL database](#) during setup, the database will store its files under `database/` in the Confluence Home Directory. Otherwise, the database management system you are connecting to is responsible for where and how your remaining data is stored.

### Temp directory

The temp directory is configured in the Java runtime and some Confluence components write temporary files or lockfiles into this directory.

The temp directory is located in the installation directory as `/temp`.

To change the location of this directory, start the Java Virtual Machine in which confluence is running with the argument:

```
-Djava.io.tmpdir=/path/to/your/own/temp/directory.
```

## Application Server Configuration

The following pages contain information about configuring your application server for Confluence:

- [Configuring URL Encoding on Tomcat Application Server](#)
- [Managing Application Server Memory Settings](#)
- [Switching to Apache Tomcat](#)
- [Java Policy Settings for Enterprise or Webhosting Environments](#)

### Configuring URL Encoding on Tomcat Application Server

Application servers may have different settings for character encodings. We strongly recommend **UTF-8** where possible.

By default, Tomcat uses ISO-8859-1 character encoding when decoding URLs received from a browser. This can cause problems when Confluence's encoding is UTF-8, and you are using international characters in the names of attachments or pages.

#### To configure the URL encoding in Tomcat:

1. Edit `conf/server.xml` and find the line where the `Coyote HTTP Connector` is defined. It will look something like this, possibly with more parameters:

```
<Connector port="8090"/>
```

2. Add a `URIEncoding="UTF-8"` property to the connector:

```
<Connector port="8090" URIEncoding="UTF-8"/>
```

3. Restart Tomcat

#### If you are using mod\_jk

You should apply the same `URIEncoding` parameter as above to the AJP connector if you are using `mod_jk`, and add the following option to your Apache `mod_jk` configuration:

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

```
JkOptions +ForwardURISCompatUnparsed
```

#### More information using Apache with Tomcat

For comprehensive examples of how to use Tomcat and Apache with Confluence, see [Running Confluence behind Apache](#).

### Managing Application Server Memory Settings

The minimum and maximum JVM heap space allocated to the application server affects performance. Confluence administrators may wish to modify this value from the defaults [depending on their server load](#). This document only provides guidelines rather than rules, so administrators optimising for performance should use this document as a starting point only.

For a comprehensive overview of memory management, and memory tuning in Confluence under Sun JRE, please read [Garbage Collector Performance Issues](#)

#### Testing For Optimum Memory Settings

In the general case, both JIRA & Confluence users will benefit from setting the minimum and maximum values identical. In larger installations, there is benefit to memory tuning, if there is a perceived performance issue. If you are experiencing Out of Memory Heap errors, try increasing the -Xmx and -Xms values for your installation to see if this resolves or helps resolve your issue. It's best to increase in small increments (eg 512mb at a time), to avoid having too large a heap, which can cause different problems. If increasing the memory does not help, please lodge a support ticket as there may be other factors contributing.

Memory usage is most likely to be maximised under peak load, and when [creating a site XML backup](#). In many cases, the backup can be the cause of the OOM, so increase -Xmx values and verify if a backup was occurring at the time of OOM. A quick rule of thumb for gauging the success of a memory adjustment is using simple anecdotal evidence from users. Is it snappier? The same? How does it handle while a backup is occurring?

Atlassian recommends in normal use, to disable the XML backup and use a [Production Backup Strategy](#).

- If you normally perform manual XML site backups on your server, test your maximum memory requirements by performing a site XML backup while the server is under maximum load
- If you do not create manual XML site backups, simply monitor the server while under maximum load

#### Applying Memory Settings

See [How to fix out of memory errors by increasing available memory](#).

#### Related Topics

- [Garbage Collector Performance Issues](#)
- [How to fix out of memory errors by increasing available memory](#)
- [Server Hardware Requirements Guide](#)
- [Performance Tuning](#)
- [Troubleshooting Slow Performance Using Page Request Profiling](#)
- [Tomcat JVM options and Modify the Default JVM Settings](#)

### Switching to Apache Tomcat

Apache Tomcat is the only application server supported for Confluence. To move Confluence from an application server (e.g. WebSphere) to Tomcat using the same database, follow the instructions below.

Follow these instructions:

- [1. Before You Start](#)
- [2. Backing Up](#)
- [3. Switching Application Servers](#)
- [4. Applying Customisations](#)
  - [Confluence Server](#)
  - [Plugins](#)
  - [Look and Feel](#)
  - [Performance](#)
  - [Advanced Customisations](#)


- [5. Testing Confluence](#)

## 1. Before You Start


1. The following instructions will only work if you are running the same major version of Confluence on both application servers. If you are running different major versions of Confluence, you will need to [upgrade Confluence](#) before you can switch to Tomcat.
2. Note that you need current [software maintenance](#), as the process for changing application servers involves installing Confluence.
3. If the environment (e.g. the database system, the operating system and so on) that you are running Confluence in has changed, please ensure it still complies with the Confluence [System Requirements](#).
4. If you are using an external database, familiarise yourself with all [known issues](#) for your specific database. Also make sure the Confluence database connector principal (the database user login) has sufficient permissions to modify the database schema.
5. Note any customisations that you have made to Confluence, e.g. enabled/installed plugins, modified layouts, custom themes, etc. You will need to reapply these after you have switched to Tomcat. You can view the list of customisations in the [Reapplying Customisations](#) section below.
6. We recommend that you **do not** run any other applications in your Tomcat application server that is running Confluence, to prevent performance issues.

## 2. Backing Up

Before you switching to Tomcat, you must back up the following:

1. **Back up your Confluence Home directory.** The location of the Home directory is stored in a configuration file called `confluence-init.properties`, which is located inside the `confluence/WEB-INF/classes` directory in your Confluence Installation directory.
2. **Back up your database.** Perform a manual backup of your external database before proceeding with the upgrade and check that the backup was created properly. If you are not a database expert or unfamiliar with the backup-restore facilities of your database, you should try to restore the backup to a different system to ensure that the backup worked before proceeding. This recommendation is not specific to Confluence usage, but it is good practice to ensure that your database backup is not broken.  
 The 'embedded database' is the HSQLDB database supplied with Confluence for evaluation purposes, you don't need to back it up since it is stored in the home directory. But you should not use this database for production systems anyway, so if you happen to accidentally still use HSQLDB in a production system, please [migrate to a proper database](#) before the upgrade.
3. **Back up your Confluence Installation directory.**

## 3. Switching Application Servers

1. Install Confluence on your new application server. We recommend that you [download Confluence](#) as a stand alone archive (rather than using the installer).  
 **Important:** At this stage, just unzip the standalone file to a location of your choice - this will be your new installation directory. You should not start Confluence until directed to in step 3 below.
2. Copy the following files from your old Confluence installation to your new one:
  - `<CONFLUENCE_INSTALL>\confluence\WEB-INF\classes\confluence-init.properties`
  - `<CONFLUENCE_INSTALL>\confluence\WEB-INF\classes\atlassian-user.xml`
  - `<CONFLUENCE_INSTALL>\confluence\WEB-INF\classes\osuser.xml` (*copy this over if you are using JIRA user management*)
  - `<CONFLUENCE_INSTALL>\confluence\WEB-INF\classes\seraph-config.xml` (*copy this over if you using custom SSO*)
  - `<CONFLUENCE_INSTALL>\confluence\WEB-INF\web.xml` (*copy this over if you have previously modified it, e.g. to configure a datasource*)
3. Start Confluence (make sure you shutdown the old server before you start up the new one)
4. If you are running the new application server on a different machine to the old one, carry out the following actions as soon as you start the new server:
  - [Re-index your data.](#)
  - Make sure that the [attachments location](#) is valid for the new server.
5. If you have applied special settings to their Confluence server and/or Confluence look and feel, you will

need to reapply these customisations as described in [below](#).

#### 4. Applying Customisations

After switching to Tomcat, you need to review any customisations and other special configurations you previously used for your Confluence instance, and re-apply if necessary. This section also contains some Tomcat-specific customisations that you may wish to considering applying, if you haven't used Confluence with Tomcat before.

##### Before you apply customisations

Please ensure that your Confluence installation works correctly on Tomcat without any customisations before you apply any of customisations listed below. This will make it easier to identify problems, if you run into trouble during the switch to Tomcat.

##### Confluence Server

- For long-term use, we recommend that you configure Confluence to start automatically when the operating system restarts. For Windows servers, this means [configuring Confluence to run as a Windows service](#).
- If you are using the Confluence edition and you have previously defined a **CATALINA\_HOME environment variable**, please check that it points to the correct path for the new Confluence Tomcat server.
- If you were previously running **Confluence on a non-standard port**, edit your new `<Installation-Direct>\conf\server.xml` file as described in [Change listen port for Confluence](#).

##### Plugins

- If you were previously using any **plugins**, install the latest compatible version and disable any plugins that are incompatible with your new instance of Confluence. The easiest way to do this is to use the [Universal Plugin Manager](#) in the Confluence Administration Console.

##### Look and Feel

- If you are using any **customised themes**, please check that they are displaying as expected. Some further customisation may be required to ensure compatibility with your new version of Confluence.
- If you had previously customised the **default site or space layouts**, you will need to reapply your changes to the new defaults as [described here](#). Please do not just copy your VM (velocity) files across. Ensure that Confluence works without your custom layouts then apply the layout via the Confluence Administration console.

##### Performance

- If the load on your Confluence instance is high, you may need more simultaneous connections to the database. Read more about this in the [Performance Tuning](#) guide.
- If you had previously modified the **memory flags** (`Xms` and `Xmx`) in either the `<Installation-Direct>\bin\setenv.sh` or the `setenv.bat` file, you may want to make the modifications in your new installation. The parameters are specified in the `CATALINA_OPTS` variable. See [How to fix out of memory errors by increasing available memory](#) for more information.

##### Advanced Customisations

- If you were previously running **Confluence over SSL**, you will need to reapply your configuration as described in [Running Confluence Over SSL or HTTPS](#).
- If you were using a custom SSO authenticator, change `seraph-config.xml` to the correct authenticator.
- If you had changed the [Confluence interface text](#), you will need to copy over the `ConfluenceActionSupport.properties` file.
- If you had previously modified the Confluence **source code**, you will need to reapply your changes to the new version.

#### 5. Testing Confluence

Make sure you **test Confluence on the new server** before deploying it in production.

The [Working with Confluence Logs](#) document contains the locations for the application logs, if you need to refer to them.

### Java Policy Settings for Enterprise or Webhosting Environments

Confluence relies on a number of Java libraries. Some of these libraries make use of features of the Java language that may be restricted by Java security policies.

This does not normally cause any problems. The default security configuration of most application servers will happily run Confluence. However, in some shared-hosting or enterprise environments, security settings may be such that Confluence cannot function.

#### Related pages:

- [Application Server Configuration](#)
- [Confluence Administrator's Guide](#)

When you attempt to run Confluence, you may get the following error:

```
java.security.AccessControlException: access denied (java.lang.RuntimePermission
accessDeclaredMembers)
 at
 java.security.AccessControlContext.checkPermission(AccessControlContext.java(Compiled
 Code))
 at
 java.security.AccessController.checkPermission(AccessController.java(Compiled
 Code))
 at java.lang.SecurityManager.checkPermission(SecurityManager.java(Compiled
 Code))
```

The permissions required by Confluence to run are detailed in the sample policy file below. You may need to give this information to your systems administrator so that they can be deployed with the Confluence application.

```
grant codeBase "file:${catalina.home}/webapps/confluence/-" {
 permission java.security.AllPermission;
};

grant {
 permission java.lang.RuntimePermission "accessDeclaredMembers";
 permission java.lang.reflect.ReflectPermission "suppressAccessChecks";
 permission java.lang.RuntimePermission "defineCGLIBClassInJavaPackage";
};
```

## Web Server Configuration

- [Configuring Web Proxy Support for Confluence](#)
- [Running Confluence behind Apache](#)
  - [General Apache Configuration Notes](#)
  - [Using Apache with mod\\_proxy](#)
  - [Using Apache with virtual hosts and mod\\_proxy](#)
  - [Using Apache with mod\\_jk](#)
  - [Using mod\\_rewrite to Modify Confluence URLs](#)
  - [Configuring Apache to Cache Static Content via mod\\_disk\\_cache](#)

### Configuring Web Proxy Support for Confluence

Some of Confluence's macros, such as {rss} and {jiraissues} need to make web requests to remote servers in order to retrieve data. If Confluence is deployed within a data centre or DMZ, it may not be able to access the Internet directly to make these requests. If you find that the {rss} macro does not work, ask your network administrator if Confluence needs to access the Internet through a web proxy.

#### *Configuring an outbound HTTP proxy in Confluence*

Proxy support is configured by passing certain [system properties](#) to the Java Virtual Machine on startup. These properties follow the [conventions defined by Oracle](#):

- `http.proxyHost`
- `http.proxyPort` (default: 80)
- `http.nonProxyHosts` (default: <none>)
- `https.proxyHost`
- `https.proxyPort`

At a minimum, you need to define `http.proxyHost` to configure an HTTP proxy, and `https.proxyHost` to configure an HTTPS proxy. System property configuration is described in the [Configuring System Properties](#).

Properties `http.proxyHost` and `http.proxyPort` indicate the proxy server and port that the http protocol handler will use, and `https.proxyHost` and `https.proxyPort` indicate the same for the https protocol handler.

```
-Dhttp.proxyHost=proxy.example.org -Dhttp.proxyPort=8080
-Dhttps.proxyHost=proxy.example.org -Dhttps.proxyPort=8080
```

Property `http.nonProxyHosts` indicates the hosts which should be connected to directly and not through the proxy server. The value can be a list of hosts, each separated by a pipe character `|`. In addition, a wildcard character (asterisk) `*` can be used for matching. For example:

```
-Dhttp.nonProxyHosts=*.foo.com|localhost
```

Note: You may need to escape the pipe character `|` in some command-line environments.

If the `http.nonProxyHosts` property is not configured, all web requests will be sent to the proxy.

Please note that any command line parameters set are visible from the process list, and thus anyone who has the appropriate access to view the process list will see the proxy information in the clear. To avoid this, you can set these properties in the `catalina.properties` file, located in `confluence-install/conf/`. Add this to the end of the file:

```
http.proxyHost=yourProxyURL
http.proxyPort=yourProxyPort
http.proxyUser=yourUserName
http.proxyPassword=yourPassword
https.proxyHost=yourProxyURL
https.proxyPort=yourProxyPort
https.proxyUser=yourUserName
https.proxyPassword=yourPassword
```

### **Configuring HTTP proxy authentication**

Proxy authentication is also configured by providing [system properties](#) to Java in your application server's configuration file. Specifically, the following two properties:

- `http.proxyUser` – username
- `http.proxyPassword` – secret

### **HTTP proxy (Microsoft ISA) NTLM authentication**

Confluence supports NTLM authentication for outbound HTTP proxies when Confluence is running on a Windows server.

This means that the {rss} and {jiraissues} macro will be able to contact external websites if requests have to go through a proxy that requires Windows authentication. This support is not related to logging in Confluence users automatically with NTLM, for which there is a [user-contributed authenticator](#) available.

To configure NTLM authentication for your HTTP proxy, you need to define a domain [system property](#), `http.auth.ntlm.domain`, in addition to the properties for host, port and username mentioned above:

```
-Dhttp.auth.ntlm.domain=MYDOMAIN
```

### Configuring authentication order

Sometimes multiple authentication mechanisms are provided by an HTTP proxy. If you have proxy authentication failure messages, you should first check your username and password, then you can check for this problem by examining the HTTP headers in the proxy failure with a packet sniffer on the Confluence server. (Describing this is outside the scope of this document.)

To set the order for multiple authentication methods, you can set the [system property](#) `http.proxyAuth` to a comma-separated list of authentication methods. The available methods are: ntlm, digest and basic; this is also the default order for these methods.

For example, to attempt Basic authentication before NTLM authentication, and avoid Digest authentication entirely, you can set the `http.proxyAuth` property to this value:

```
-Dhttp.proxyAuth=basic,ntlm -Dhttps.proxyAuth=basic,ntlm
```

### Troubleshooting

1. There's a diagnostic jsp file in [CONF-9719](#) for assessing the connection parameters.
2. 'Status Code [407]' errors are described in [APR-160](#).
3. Autoproxies are not supported. See [CONF-16941](#).

### Running Confluence behind Apache

This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

### Introduction

For improved performance in high-load environments, you should run Confluence behind a web server. In general, web server caching and thread management is far superior to that provided by your application server's HTTP interface.

To run Confluence behind the [Apache httpd web server](#), there are two main configuration options: `mod_jk` or `mod_proxy`.

Connection type	Unique features	Common features to both <code>mod_proxy</code> and <code>mod_jk</code>
<code>mod_proxy</code> (also known as reverse proxy)	<ul style="list-style-type: none"> <li>• <b>recommended connection method</b></li> <li>• simple HTTP proxy to application server</li> <li>• works with all application servers</li> <li>• if application paths are consistent, there is minimal load on the web server</li> </ul>	<ul style="list-style-type: none"> <li>• application paths must be consistent to avoid complex and slow URL rewriting</li> <li>• works with name-based virtual hosting, both on web server and app server</li> <li>• web server keeps a pool of connections to application server</li> </ul>

mod_jk (also known as AJP)	<ul style="list-style-type: none"> <li>• uses the AJP binary protocol</li> <li>• provides failover (and load balancing, which Confluence supports only with a clustered license)</li> <li>• only works with some application servers (typically Tomcat)</li> <li>• if application paths are consistent, there is some load on the web server to translate requests to AJP</li> </ul>
----------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

### Configuration Guides

Please choose one configuration. Trying to configure for both mod\_proxy and mod\_jk will only lead to confusion and tears.

- [Using Apache with mod\\_proxy](#)
- [Using Apache with mod\\_jk](#)
- [Using Apache with virtual hosts and mod\\_proxy](#)

### Mod\_jk2 not supported

The misleadingly-named mod\_jk2 is an older method of connecting to Tomcat from Apache. Since mod\_jk2 is [no longer supported by the Apache Foundation](#), we do not support this configuration, and are not updating our [mod\\_jk2 documentation](#). Mod\_jk2 also has unresolved problems with Unicode URLs; you need to use either mod\_proxy or mod\_jk for international characters to work correctly in Confluence.

### Caching static content via mod\_disk\_cache

To improve performance of a large Confluence site, we recommend that you move the caching of static content from the JVM into Apache. This will prevent the JVM from having a number of long running threads serving up static content. See [Configuring Apache to Cache Static Content via mod\\_disk\\_cache](#).

### Other related documentation

- [Configuring Tomcat's URI encoding](#)
- [Running Confluence Over SSL or HTTPS](#)

### General Apache Configuration Notes

#### On this page:

- [Prefer Apache mod\\_deflate to Confluence's built-in gzip implementation](#)
- [Ensure keepalive is enabled](#)
- [Enable keepalive for recent MSIE user agents](#)

#### Prefer Apache mod\_deflate to Confluence's built-in gzip implementation

1. Disable gzip in confluence. See [Compressing an HTTP Response within Confluence](#).
2. Enable gzip compression in Apache. For RedHat distributions this can be achieved by adding the following lines:

```
AddOutputFilterByType DEFLATE text/html text/plain text/xml text/css
application/x-javascript

ensure sensible defaults
DeflateBufferSize 8192
DeflateCompressionLevel 4
DeflateMemLevel 9
DeflateWindowSize 15
```

### Ensure keepalive is enabled

```
KeepAlive On
```

### Enable keepalive for recent MSIE user agents

The standard Apache SSL configuration is very conservative when it comes to MSIE and SSL. By default all keepalives are disabled when using HTTPS with MSIE. While MSIE will always be *special*, the issues with SSL and MSIE have been solved since [Service Pack 2 for Windows XP](#), released over 4 years go. For anyone using an XP machine SP2 or above, it is safe to allow keepalive for MSIE 6 and above.

Remove the following lines:

```
SetEnvIf User-Agent ".*MSIE.*" \
 nokeepalive ssl-unclean-shutdown \
 downgrade-1.0 force-response-1.0
```

Add these in their place:

```
BrowserMatch "MSIE [1-5]" nokeepalive ssl-unclean-shutdown downgrade-1.0
force-response-1.0
BrowserMatch "MSIE [6-9]" ssl-unclean-shutdown
```

## RELATED TOPICS

[Running Confluence behind Apache](#)  
[Configuring Tomcat's URI encoding](#)  
[Running Confluence Over SSL or HTTPS](#)

### Using Apache with mod\_proxy

This page describes how to integrate Confluence into an Apache website using mod\_proxy.

There are some common situations where you might use the configuration:

- You have an existing Apache-based website, and want to add Confluence to the mix (for example, <http://www.example.com/confluence>).
- You have two or more Java applications, each running in their own application server on different ports, for example, <http://example:8090/confluence> and <http://example:8080/jira>. By setting up Apache with mod\_proxy, you can have both available on the regular HTTP port (80) – for example, at <http://www.example.com/confluence> and <http://www.example.com/jira>. This allows each application to be restarted, managed and debugged separately.

**Note:** This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

### Base configuration

In these examples, we use the following:

**<http://www.example.com/confluence>** - your intended URL

**<http://example:8090>** - the hostname and port Confluence is currently installed to

**[/confluence](#)** - the intended context path (the part after hostname and port)

Please substitute the examples below with your intended URL's in your own server. Copy/pasting these suggestions will not work on your server.

#### Set the context path

Set your Confluence application path (the part after hostname and port). To do this in Tomcat (bundled with Confluence), edit `conf/server.xml`, locate the "Context" definition:

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

and change it to:

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

Then restart Confluence, and ensure you can access it at **<http://example:8090/confluence>**

#### Set the URL for redirection

Set the URL for redirection. In the same `conf/server.xml` file, locate this code segment:

```
<Connector port="8090" maxHttpHeaderSize="8192"
 maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
 enableLookups="false" redirectPort="8443" acceptCount="100"
 connectionTimeout="20000" disableUploadTimeout="true" />
```

And append the last line:

```
<Connector port="8090" maxHttpHeaderSize="8192"
 maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
 enableLookups="false" redirectPort="8443" acceptCount="100"
 connectionTimeout="20000" disableUploadTimeout="true"
 proxyName="www.example.com" proxyPort="80" />
```

If this isn't working for you and you're using SSL, try adding a scheme attribute to your Connector tag: **scheme="https"**.

Now we have two options:

- If you want a URL like **<http://www.example.com/confluence>**, follow the simple configuration.
- If you want a URL like **<http://confluence.example.com>**, go to the [complex configuration](#).

#### Simple Configuration

##### Configure `mod_proxy`

Now enable `mod_proxy` in Apache, and proxy requests to the application server by adding the example below to your Apache `httpd.conf` (note: the files may be different on your system; See [Integrating JIRA with Apache](#) for th

e process for Ubuntu/Debian layout):

### Apache 2.2

```
Put this after the other LoadModule directives
LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so

Put this in the main section of your configuration (or desired virtual host, if
using Apache virtual hosts)
ProxyRequests Off
ProxyPreserveHost On

<Proxy *>
 Order deny,allow
 Allow from all
</Proxy>

ProxyPass /confluence http://app-server.internal.example.com:8090/confluence
ProxyPassReverse /confluence http://app-server.internal.example.com:8090/confluence
<Location /confluence>
 Order allow,deny
 Allow from all
</Location>
```

### Apache 2.4

```
Put this after the other LoadModule directives
LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so

Put this in the main section of your configuration (or desired virtual host, if
using Apache virtual hosts)
ProxyRequests Off
ProxyPreserveHost On

<Proxy *>
 # Auth changes in 2.4 - see
 http://httpd.apache.org/docs/2.4/upgrading.html#run-time
 Require all granted
</Proxy>

ProxyPass /confluence http://app-server.internal.example.com:8090/confluence
ProxyPassReverse /confluence http://app-server.internal.example.com:8090/confluence
<Location /confluence>
 # Auth changes in 2.4 - see
 http://httpd.apache.org/docs/2.4/upgrading.html#run-time
 Require all granted
</Location>
```

#### ▼ Note to Windows Users

It is recommended that you specify the absolute path to the `mod_proxy.so` and `mod_proxy_http.so` files

#### Complex configuration

Complex configuration involves using the `mod_proxy_html` filter to modify the proxied content en-route. This is required if the Confluence path differs between Apache and the application server. For example:

<b>Externally accessible (Apache) URL</b>	<a href="http://confluence.example.com/">http://confluence.example.com/</a>
<b>Application server URL</b>	<a href="http://app-server.internal.example.com:8090/confluence/">http://app-server.internal.example.com:8090/confluence/</a>

Notice that the application path in the URL is different in each. On Apache, the path is /, and on the application server the path is /confluence.

For this configuration, you need to install the `mod_proxy_html` module, which is not included in the standard Apache distribution.

Alternative solutions are discussed below.

### Apache 2.2

```
Put this after the other LoadModule directives
LoadModule proxy_module modules/mod_proxy.so
LoadModule proxy_http_module modules/mod_proxy_http.so
LoadModule proxy_html_module modules/mod_proxy_html.so

<VirtualHost *>
 ServerName confluence.example.com

 # Put this in the main section of your configuration (or desired virtual host,
 if using Apache virtual hosts)
 ProxyRequests Off
 ProxyPreserveHost On

 <Proxy *>
 Order deny,allow
 Allow from all
 </Proxy>

 ProxyPass / http://app-server.internal.example.com:8090/confluence
 ProxyPassReverse / http://app-server.internal.example.com:8090/confluence

 ProxyHTMLURLMap / /confluence/

 <Location />
 Order allow,deny
 Allow from all
 </Location>
</VirtualHost>
```

### Apache 2.4

```
Put this after the other LoadModule directives
LoadModule proxy_module modules/mod_proxy.so
LoadModule proxy_http_module modules/mod_proxy_http.so
LoadModule proxy_html_module modules/mod_proxy_html.so

<VirtualHost *>
 ServerName confluence.example.com

 # Put this in the main section of your configuration (or desired virtual host,
 if using Apache virtual hosts)
 ProxyRequests Off
 ProxyPreserveHost On

 <Proxy *>
 # Auth changes in 2.4 - see
 http://httpd.apache.org/docs/2.4/upgrading.html#run-time
 Require all granted
 </Proxy>

 ProxyPass / http://app-server.internal.example.com:8090/confluence
 ProxyPassReverse / http://app-server.internal.example.com:8090/confluence

 ProxyHTMLURLMap / /confluence/

 <Location />
 # Auth changes in 2.4 - see
 http://httpd.apache.org/docs/2.4/upgrading.html#run-time
 Require all granted
 </Location>
</VirtualHost>
```

The ProxyHTMLURLMap configuration can become more complex if you have multiple applications running under this configuration. The mapping should also be placed in a Location block if the web server URL is a subdirectory and not on a virtual host. The [Apache Week tutorial](#) has more information how to do this.

### Final Configuration Steps

#### Restart your Apache server

This is needed to pick up on the new configuration. This can be done by running the following on your command line/terminal/shell:

```
sudo apachectl graceful
```

#### Disable HTTP Compression

Having compression run on both the proxy and Tomcat can cause problems integrating with other Atlassian applications, such as JIRA. Please disable HTTP compression as per our [Compressing an HTTP Response within Confluence](#) docs.

#### Set the Confluence Base URL

The last stage is to set the [Base URL](#) to the address you're using within the proxy. In this example, it would be <http://www.example.com/confluence>

### Adding SSL

If you're running Apache in front of Tomcat, it's a good idea to terminate your SSL configuration at Apache, then

forward the requests to Tomcat over HTTP. You can set up Apache to terminate the SSL connection and use the ProxyPass and ProxyPassReverse directives to pass the connection through to Tomcat (or the appropriate application server) which is running Confluence.

1. Create a new SSL host by creating a virtual host on 443
2. The standard http connection on apache could be used to redirect to https if you want or it could just be firewalled.
3. Within the VirtualHost definition:
  - a. define the SSL options (SSLEngine and SSLCertificateFile)
  - b. define the ProxyPass and ProxyPassReverse directives to pass through to Tomcat.

Most of the relevant Apache Config:

```
Listen 443

NameVirtualHost *:443
<VirtualHost *:443>
 SSLEngine On
 SSLCertificateFile /etc/apache2/ssl/apache.pem
 ProxyPass / http://localhost:8090/
 ProxyPassReverse / http://localhost:8090/
</VirtualHost>
```

Apart from the Apache configuration there are a couple of things you will need to do before you get your server working:

1. You will have to change your base URL to point to https addresses. See the documentation on [configuring the server base URL](#).
2. We need to set up the connector to use https. In your installation directory, edit the file server.xml and add this attributes to your connector:

```
proxyName="proxy.example.com" proxyPort="443" scheme="https" secure="true"
```

#### More information

- The [mod\\_proxy\\_html site](#) has documentation and examples on the use of this module in the complex configuration.
- [Apache Week has a tutorial](#) that deals with a complex situation involving two applications and ProxyHTMLURLMap.
- [Using Apache with virtual hosts and mod\\_proxy](#) shows how to configure the special case where you want JIRA and Confluence running on separate application servers on virtual host subdomains.

#### Alternatives

If Tomcat is your application server, you have two options:

- use [mod\\_jk](#) to send the requests to Tomcat
- use Tomcat's virtual hosts to make your Confluence application directory the same on the app server and the web server, removing the need for the URL mapping.

If your application server has an AJP connector, you can:

- use [mod\\_jk](#) to send the requests to your application server.

#### Using Apache with virtual hosts and mod\_proxy

**Note:** This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

#### Introduction

The Apache web server is often used in front of an application server to improve performance in high-load environments. Mod\_proxy simply redirects requests for certain URLs to another web server, so it typically requires no additional configuration on the application server.

This page documents a very common configuration request: configuring a JIRA application and Confluence on two Apache virtual hosts, running on different application servers. This is just a special case of [mod\\_proxy configuration](#).

You can use virtual hosts in your application server if you want to run JIRA applications and Confluence on the same application server.

#### Apache configuration

For this configuration to work properly, the application paths must be the same on both the application servers and the web server. For both JIRA and Confluence below, this is /.

<b>JIRA external URL</b>	<a href="http://jira.example.com/">http://jira.example.com/</a>
<b>JIRA application server URL</b>	<a href="http://jira-app-server.internal.example.com:8080/">http://jira-app-server.internal.example.com:8080/</a>
<b>Confluence external URL</b>	<a href="http://confluence.example.com/">http://confluence.example.com/</a>
<b>Confluence application server URL</b>	<a href="http://confluence-app-server.internal.example.com:8090/">http://confluence-app-server.internal.example.com:8090/</a>

Add the following to your Apache httpd.conf:

```
Put this after the other LoadModule directives
LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so

Put this with your other VirtualHosts, or at the bottom of the file
NameVirtualHost *
<VirtualHost *>
 ServerName confluence.example.com

 ProxyRequests Off
 <Proxy *>
 Order deny,allow
 Allow from all
 </Proxy>

 ProxyPass / http://confluence-app-server.internal.example.com:8090/
 ProxyPassReverse / http://confluence-app-server.internal.example.com:8090/
 <Location />
 Order allow,deny
 Allow from all
 </Location>
</VirtualHost>
<VirtualHost *>
 ServerName jira.example.com

 ProxyRequests Off
 <Proxy *>
 Order deny,allow
 Allow from all
 </Proxy>

 ProxyPass / http://jira-app-server.internal.example.com:8080/
 ProxyPassReverse / http://jira-app-server.internal.example.com:8080/
 <Location />
 Order allow,deny
 Allow from all
 </Location>
</VirtualHost>
```

#### Points to note:

- ProxyPass and ProxyPassReverse directives send traffic from the web server to your application server.
- The application path is the same on the application server and on the web server (both are /).
- Because the above configuration uses [name-based virtual hosting](#), you must configure your DNS server to point both names (jira.example.com, confluence.example.com) to your web server.

#### More information

For different ways to configure mod\_proxy, see [Using Apache with mod\\_proxy](#).

If you use Tomcat, [mod\\_jk](#) provides a different way of connecting Apache via AJP. You can also use the above configuration with just one application server if you use Tomcat's virtual hosts.

#### Using Apache with mod\_jk

- The preferred configuration is [Using Apache with mod\\_proxy](#). This works with any application server, and together with mod\_proxy\_html allows complex URL rewriting to deal with different application paths on the web server and the application server.
- This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you

use it at your own risk.

## Introduction

The Apache web server is often used in front of an application server to improve performance in high-load environments. Mod\_jk allows request forwarding to an application via a protocol called AJP. Configuration of this involves enabling mod\_jk in Apache, configuring a AJP connector in your application server, and directing Apache to forward certain paths to the application server via mod\_jk.

Mod\_jk is sometimes preferred to [mod\\_proxy](#) because AJP is a binary protocol, and because some site administrators are more familiar with it than with mod\_proxy..

The scope of this documentation is limited to configuring the AJP connector in Tomcat 5.x. Other application servers may support AJP connectors; please consult your application server documentation for instructions on how to configure it.

The configuration below assumes your Confluence instance is accessible on the same path on the application server and the web server. For example:

<b>Externally accessible (web server) URL</b>	<a href="http://www.example.com/confluence/">http://www.example.com/confluence/</a>
<b>Application server URL (HTTP)</b>	<a href="http://app-server.internal.example.com:8090/confluence/">http://app-server.internal.example.com:8090/confluence/</a>

The AJP connection of the application server is set to: app-server.internal.example.com:8009.

## Configuring mod\_jk in Apache

The standard distribution of Apache does not include mod\_jk. You need to download it from the [JK homepage](#) and put the mod\_jk.so file in your Apache modules directory.

Next, add the following in **httpd.conf** directly or included from another file:

```
Put this after the other LoadModule directives
LoadModule jk_module modules/mod_jk.so

Put this in the main section of your configuration (or desired virtual host, if
using Apache virtual hosts)
JkWorkersFile conf/workers.properties
JkLogFile logs/mod_jk.log
JkLogLevel info

JkMount /confluence worker1
JkMount /confluence/* worker1
```

## Configuring workers.properties

Create a new file called 'workers.properties', and put it in your Apache conf directory. (The path for workers.properties was one of the configuration settings above.)

```
worker.list=worker1

worker.worker1.host=app-server.internal.example.com
worker.worker1.port=8009
worker.worker1.type=ajp13
```

## Tomcat 5.x configuration

In Tomcat 5, the AJP connector is enabled by default on port 8009. An absolutely minimal Tomcat server.xml is below for comparison. The relevant line is the Connector with port 8009 – make sure this is uncommented in

your server.xml.

```
<Server port="8000" shutdown="SHUTDOWN">
 <Service name="Catalina">

 <!-- Define a HTTP/1.1 Connector on port 8090 -->
 <Connector port="8090" />

 <!-- Define an AJP 1.3 Connector on port 8009 -->
 <Connector port="8009" protocol="AJP/1.3" />

 <Engine name="Catalina" defaultHost="localhost">
 <Host name="localhost" appBase="webapps">
 <Context path="/confluence" docBase="/opt/webapps/confluence-2.2/confluence"/>
 <Logger className="org.apache.catalina.logger.FileLogger"/>
 </Host>
 </Engine>
 </Service>
</Server>
```

Points to note:

- the Connector on port 8009 has protocol of "AJP/1.3". This is critical.
- the Context path of the Confluence application is "/confluence". This must match the path used to access Confluence on the web server.
- we recommend keeping your application Contexts outside the server.xml in Tomcat 5.x. The above example includes them for demonstration only.

#### Improving the performance of the mod\_jk connector

The most important setting in high-load environments is the number of processor threads used by the Tomcat AJP connector. By default, this is 200, but you should increase it to match Apache's maxThreads setting (256 by default):

```
<Connector port="8009" minSpareThreads="5" maxThreads="256" protocol="AJP/1.3" />
```

All the configuration parameters for the AJP connector are covered in [the Tomcat documentation](#).

#### Ensuring UTF-8 compatibility

If you have problems downloading attachments with non-ASCII characters in the filename, add the following to your Apache configuration:

```
JkOptions +ForwardURICompatUnparsed
```

And specify UTF-8 as the URIEncoding in the AJP connector configuration:

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

These settings are discussed further on [Configuring Tomcat's URI encoding](#).

#### More information

The [Tomcat JK website](#) has complete documentation on workers.properties and Apache configuration. You can also find information there on how to use mod\_jk with IIS.

**Note:** In IIS proxy server the maximum file upload is 30mb by default. Contact your server administrator if you need to upload a bigger file.

## Using `mod_rewrite` to Modify Confluence URLs

**Note:** This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

Confluence requires URL rewriting for proper functionality, if Confluence is accessible via different domain names. If Confluence is configured for multiple domains *without* URL rewriting, you will experience an array of problems. See [Various Issues Caused when Server Base URL Does Not Match the URL Used to Access Confluence](#).

An example of why you may want to access Confluence from different domains:

- From an internal network:  
`http://wiki`
- The externally visible domain:  
`http://wiki.domain.com`

### Using URL rewriting to access Confluence over multiple domains

To configure Confluence over multiple domains:

1. Add a DNS entry mapping `http://wiki` to the externally visible IP address of the Confluence server.
2. Set Confluence's [server base URL](#) to `http://wiki.domain.com`.
3. Add Apache HTTP proxy, using the instructions from [Running Confluence behind Apache](#).
4. Add the `mod_rewrite` module to change the URL.

### Further information

You may be interested in the [UrlRewriteFilter](#) that is Java web filter that works in a similar way of the **Apache's** `mod_rewrite`.

## Configuring Apache to Cache Static Content via `mod_disk_cache`

To improve performance of a large Confluence site, we recommend that you move the caching of static content from the JVM into Apache. This will prevent the JVM from having a number of long running threads serving up static content.

Static content in Confluence includes most JavaScript, CSS and image files which are included with the application or an installed plugin. This content will be cached by Apache in this configuration. User-provided content like space logos, attachments or embedded images are not considered static content and will not be cached.

**Note:** This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

### Configuring Apache `mod_disk_cache`

To configure Apache to cache static Confluence content:

1. Add a `mod_disk_cache` stanza to the virtual host configuration:

```
<IfModule mod_disk_cache.c>
"/s" is where Confluence serves "static" stuff. Instruct Apache
to cache it:
CacheEnable disk /s
CacheIgnoreHeaders Set-Cookie
CacheRoot "/var/cache/mod_proxy"
</IfModule>
```

2. Configure Apache to load `mod_disk_cache`. For example, in our server configuration this is done in `/etc/httpd/conf/httpd.conf`:

```
LoadModule disk_cache_module modules/mod_disk_cache.so
```

3. Restart Apache after both modifications are complete.

#### Notes

- Please refer to the Apache documentation for [mod\\_disk\\_cache](#).
- If you encounter problems where users are served stale content, you may need to purge the Apache cache directory (`/var/cache/mod_proxy` in the above configuration) after a Confluence or plugin upgrade. This is a simple 3 step process:
  - Shut down Apache.
  - Clear the cache directory. For example: `sudo rm -r /var/cache/mod_proxy/*`
  - Restart Apache.
- Ensure that you are running the `htcacheclean` daemon in order to prevent excessive use of disk space. In our situation we ran it like this:

```
sudo htcacheclean -d30 -n -t -p /var/cache/mod_proxy -l 512M
```

This will purge content once the cache reaches 512M every 30 minutes. See the Apache documentation for [htcacheclean](#) for details of the options.

## Starting Confluence Automatically on System Startup

You can configure Confluence to start automatically on system startup, allowing it to recover automatically after a reboot.

### Content by label

There is no content with the specified labels



## Start Confluence Automatically on Linux

On Linux/Solaris, the best practice is to install, configure and run each service (including Confluence) as a dedicated user with only the permissions they require.

To install, configure and run Confluence automatically on Linux/Solaris:

1. Create a `confluence` user for instance, using the following command:

```
sudo useradd --create-home -c "Confluence role account" confluence
```

2. Create a directory to install Confluence into:

```
sudo mkdir /usr/local/confluence
sudo chown confluence: /usr/local/confluence
```

3. Log in as the `confluence` user to install Confluence:

```
sudo su - confluence
cd /usr/local/confluence/
tar zxvf /tmp/confluence-5.6.4.tar.gz
ln -s confluence-5.6.4/ current
```

#### 4. Edit

<<CONFLUENCE\_INSTALL\_DIRECTORY>>/confluence/WEB-INF/classes/confluence-init.properties file, and set `confluence.home=/usr/local/confluence/<Confluence_Data_Home>` (ensure you have removed the comment '#')

#### 5. Then back as root, create the file `/etc/init.d/confluence` (code shown below), which will be responsible for starting up Confluence after a reboot (or when manually invoked).

⚠ If you are running Ubuntu Jaunty (or later) do not perform this step. Please use the [instructions](#) further down this page.

```
#!/bin/sh -e
Confluence startup script
#chkconfig: 2345 80 05
#description: Confluence

Define some variables
Name of app (JIRA, Confluence, etc)
APP=confluence
Name of the user to run as
USER=confluence
Location of application's bin directory
CATALINA_HOME=/usr/local/confluence/current
Location of Java JDK
export JAVA_HOME=/usr/lib/jvm/java-7-oracle

case "$1" in
 # Start command
 start)
 echo "Starting $APP"
 /bin/su -m $USER -c "$CATALINA_HOME/bin/start-confluence.sh &> /dev/null"
 ;;
 # Stop command
 stop)
 echo "Stopping $APP"
 /bin/su -m $USER -c "$CATALINA_HOME/bin/stop-confluence.sh &> /dev/null"
 echo "$APP stopped successfully"
 ;;
 # Restart command
 restart)
 $0 stop
 sleep 5
 $0 start
 ;;
 *)
 echo "Usage: /etc/init.d/$APP {start|restart|stop}"
 exit 1
 ;;
esac

exit 0
```


#### 6. Make this file executable:

```
sudo chmod +x /etc/init.d/confluence
```

7. Set this file to run at the appropriate runlevel. For example, use `sudo chkconfig --add confluence` on Redhat-based systems, `sudo update-rc.d confluence defaults` or `rcconf` on Debian-based systems.
8. You should now be able to start Confluence with the init script. A successful startup output typically looks like this:

```
$ sudo /etc/init.d/confluence start
Starting Confluence:
If you encounter issues starting up Confluence, please see the
Installation guide at
http://confluence.atlassian.com/display/DOC/Confluence+Installation
+Guide
Using CATALINA_BASE: /usr/local/confluence/current
Using CATALINA_HOME: /usr/local/confluence/current
Using CATALINA_TMPDIR: /usr/local/confluence/current/temp
Using JRE_HOME: /usr/lib/jvm/java-1.7.0-oracle
done.
```

You should then see this running at `http://<server>:8090/`

 The port for this will be whatever is defined in your Confluence `server.xml` file.

#### Adding Confluence as a service for Ubuntu Jaunty (or later)

To continue configuring Confluence to start automatically as a service on Ubuntu Jaunty (or later):

1. After **logging in** as the `confluence` user to install Confluence, create `start` and `stop` scripts in `/usr/local/confluence`:

Example `startscript`:

```
#!/bin/bash
export JAVA_HOME=/usr/lib/jvm/java-7-oracle-1.7.0.71/
export JDK_HOME=/usr/lib/jvm/java-7-oracle-1.7.0.71/
cd /usr/local/confluence/current/bin
./startup.sh
```

Example `stopscript`:

```
#!/bin/bash
export JAVA_HOME=/usr/lib/jvm/java-7-oracle-1.7.0.71/
export JDK_HOME=/usr/lib/jvm/java-7-oracle-1.6.0.71/
cd /usr/local/confluence/current/bin
./shutdown.sh
```

2. Make both of these scripts executable. For example, by issuing the command: `sudo chmod a+x /usr/local/confluence/start /usr/local/confluence/stop`.
3. Karmic and later: Create two text files in `/etc/init/` called `confluence-up.conf` and `confluence-down.conf`:

confluence-up:

```
start on runlevel [2345]

script

 date >> /tmp/confluence-startup.out
 exec sudo -u confluence /usr/local/confluence/start >>
/tmp/confluence-startup.out 2>&1

end script
```

confluence-down:

```
start on runlevel [16]

expect fork
respawn

exec sudo -u confluence /usr/local/confluence/stop >>
/tmp/confluence-shutdown.out 2>&1
```

... and make them readable to all users:

```
sudo chmod a+r /etc/init/confluence-up.conf /etc/init/confluence-down.conf
```

1. Jaunty, Intrepid: Create two text files in `/etc/event.d/` called `confluence-up` and `confluence-down`:

confluence-up:

```
start on runlevel 2
start on runlevel 3
start on runlevel 4
start on runlevel 5

exec sudo -u confluence /usr/local/confluence/start >>
/tmp/confluence-startup.out 2>&1
```

confluence-down:

```
start on runlevel 1
start on runlevel 6

exec sudo -u confluence /usr/local/confluence/stop >>
/tmp/confluence-shutdown.out 2>&1
```

... and make them readable to all users:

```
sudo chmod a+r /etc/event.d/confluence-up /etc/event.d/confluence-down
```

#### RELATED TOPICS

### Starting Confluence Automatically on System Startup

#### Start Confluence Automatically on Windows as a Service

For long-term use, we recommend that you configure Confluence to start automatically when the operating system restarts. For Windows servers, this means configuring Confluence to run as a Windows service.

There are two ways to install the Confluence distribution as a service: using the [Confluence installer](#) or manually as described below.

#### On this page:

- [Reasons for Starting Confluence as a Service](#)
- [Changing the User Running the Service](#)
- [Manually Installing the Confluence Distribution as a Service](#)
- [Managing Confluence as a Service](#)
- [Upgrading Confluence](#)
- [Troubleshooting Confluence while Running as a Windows Service](#)
- [Requesting Support](#)

#### Problem with 64-bit Windows

If you are running 64-bit Windows, please note that you may encounter problems with Apache Tomcat running as a Windows service if you are using a 64-bit JDK. Refer to our [knowledge base article](#) for more information.

#### Reasons for Starting Confluence as a Service

Installation as a Windows service offers these advantages:

- Reduced risk of shutting down Confluence by accident (If you start Confluence manually, a console window opens and there is a risk of someone accidentally shutting down Confluence by closing the window).
- Automated Confluence recovery after server restart.
- Improved troubleshooting through logging server output to file.

You can read more about Windows services in the [Microsoft Developer Network](#).

#### Changing the User Running the Service

If you wish to run the service as a non-administrator user for security, or if you are using network drives for backups, attachments or indexes, you can run the service as another user. To change users, open the Apache Tomcat Confluence properties, go to the 'Log On' tab and enter the required username and password. Go to your Windows Control Panel -> User Accounts and confirm that the user has write permissions for the <CONFLUENCE-INSTALL> and <CONFLUENCE-HOME> directories, and all subfolders. Note that any network drives must be specified by UNC and not letter mappings (eg. \\backupserver\confluence not z:\confluence).

For more detail, see [Creating a Dedicated User Account on the Operating System to Run Confluence](#).

#### Manually Installing the Confluence Distribution as a Service

From your Windows-based server:

1. Open a command prompt in the <CONFLUENCE-INSTALL>/bin directory.
2. Confirm that the JAVA\_HOME variable is set to the JDK base directory with the command:

```
echo %JAVA_HOME%
```

Note that any directory in the path with spaces (eg. C:\Program Files must be converted to its eight-character equivalent (e.g. C:\Progra~1).

3. If you are installing Confluence on a Windows 2008 server, be sure to run the command prompt using 'run as administrator'. (Otherwise running 'service.bat', as described in the next step, will fail.)
4. Use the following command to install the service with default settings:

```
service.bat install Confluence
```

NB: This will create a service called **Apache Tomcat Confluence**.

5. Now, to have the service start automatically when the server starts, run:

```
tomcat7 //US//Confluence --Startup auto
```

(if you are using a different version of tomcat, specify your tomcat version, e.g. tomcat6)


6. If you have a less than a 1024 megabytes of memory, skip this step. For users with large Confluence installations, you can increase the maximum memory Confluence can use. (The default is 1024MB). For example, you can set the maximum memory to 2048 megs using:

```
tomcat7 //US//Confluence --JvMx 2048
```

7. If you do not have any JVM parameters that you pass to your distribution of Confluence, you can skip this step. If you do, add them to the service using:

```
tomcat7 //US//Confluence ++JvmOptions="-Djust.an.example=True"
```

8. For further configuration options, please refer to the [Tomcat Windows Service How-To](#) guide
9. Go to your Windows Control Panel -> Administrative Tools -> Services -> Apache Tomcat Confluence and right-click on Properties to verify the settings are correct.

 Confluence is now installed as a service, but will not automatically start up until the next server reboot

10. Start the Confluence service with the command:

```
net start Confluence
```

### Managing Confluence as a Service

You can manage the Confluence service from the command prompt.

- Stop Confluence with:

```
net stop Confluence
```

- Uninstall the Confluence service with:

```
service.bat remove Confluence
```

### Upgrading Confluence

After upgrading Confluence, you can either uninstall and reinstall the Windows service or change the StartPath

parameter to your new folder. Refer to the [Tomcat documentation](#) for help.

### Troubleshooting Confluence while Running as a Windows Service

- Check the Knowledge Base articles:
  - [Getting 'The image file tomcat6.exe is valid, but is for a machine type other than the current machine'](#)
  - [Confluence Does Not Start Due to Windows Firewall](#)
  - [Unable to start Confluence Windows service after allocating JVM memory](#)
  - [Unable to Configure Confluence to Run as a Service on Tomcat 5](#)
  - [Unable to Install Service on Windows Vista](#)
- If none of the above solves your problem, please refer to the complete [list of known issues](#) in our Knowledge Base.
- When investigating memory issues or bugs, it may be useful to view information from Confluence's garbage collection. To turn on the verbose garbage collection, use the command:

```
tomcat6 //US//Confluence
++JvmOptions="-Xloggc:<CONFLUENCE-INSTALL>\logs\atlassian-gc.log"
```

- The Confluence 2.9 installer does not work when installed as service, due to a missing semi-colon in `service.bat`. Please refer to reported issue [CONF-12785](#).
- You can use a [Sysinternals](#) tool called **Procmon.exe** from the The Microsoft Windows Sysinternals Team, to check that the error occurred at the specific time when the Confluence service started. You need to match the time when Tomcat failed, as captured by this tool, against the time in the Windows Event Viewer.

#### Note

We do not recommend that you run this tool for too long as it may disrupt other Atlassian applications. Once you have captured the required information you will need to press **Ctrl + E** to stop capturing.

### Requesting Support

If, after following the troubleshooting guide above, you still cannot make Confluence run as a Windows Service or if there is an error when setting the JVM configuration for the service, you can [create a support request](#).

Please provide the following information when creating your support request, because we will need it to assist you:

- Are you running a 32 bit or 64 bit Windows?
- Give us the result of running **java -version** from Windows command line console.
- A screen shot of your Windows Registry setting for Tomcat.
- If you have modified **service.bat**, please give us a copy of this file for review.
- What application server are you using? eg. Are you using the Confluence distribution?
- Your [atlassian-confluence.log](#) file.

### RELATED TOPICS

[Starting Confluence Automatically on System Startup](#)

[How to fix out of memory errors by increasing available memory](#)

## Configuring Confluence

This section focuses on settings and configurations within the Confluence application.

For guidelines on external configuration, see [Configuring a Confluence Environment](#).

#### Related pages:

- [Customising your Confluence Site](#)
- [Confluence Administrator's Guide](#)

- [Viewing System Information](#)
- [Configuring the Server Base URL](#)
- [Configuring the Confluence Search and Index](#)
- [Configuring Mail](#)
- [Configuring Character Encoding](#)
- [Other Settings](#)
- [Configuring System Properties](#)
- [Working with Confluence Logs](#)
- [Configuring Confluence Security](#)
- [Scheduled Jobs](#)
- [Configuring the Whitelist](#)
- [Configuring the Time Interval at which Drafts are Saved](#)

## Viewing System Information

The System Information screen provides information about Confluence's configuration, and the environment in which Confluence has been deployed.

To view your system information:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **System Information** in the left-hand panel.

Notes:

- The handy **memory graph** helps you keep track of Confluence's memory usage.
- Your system configuration information is helpful to Atlassian Support when diagnosing errors you may face using Confluence. When logging a **support request** or **bug report**, please provide as much detail as possible about your installation and environment.

### Related pages:

- [Cache Statistics](#)
- [Live Monitoring Using the JMX Interface](#)
- [Tracking Customisations Made to your Confluence Installation](#)

## Live Monitoring Using the JMX Interface

With the JMX interface (introduced in Confluence 2.8), you can monitor the status of your Confluence instance in real time. This will provide you with useful data such as the resource usage of your instance and its database latency, allowing you to diagnose problems or performance issues. To read the JMX data, you will need to use a JMX client.

### Disable JMX

If you experience any problems during Confluence startup that are related to JMX, it is possible to disable the JMX registration process. Please place [jmxContext.xml](#) in your `<confluence-install>/confluence/WEB-INF/classes` folder to do so.

### What is JMX?

JMX ([Java Management eXtensions](#)) is a technology for monitoring and managing Java applications. JMX uses objects called MBeans (Managed Beans) to expose data and resources from your application.

#### 1. Enabling JMX Remote with Tomcat

By default, Confluence uses the Apache Tomcat web server. To use JMX, you must enable it on your Tomcat server, by carrying out the steps under the [Apache Tomcat documentation](#), entitled [Enabling JMX Remote](#). With those steps completed, restart your Tomcat server.

For the stand-alone, add the startup parameter `-Dcom.sun.management.jmxremote` to `setenv.sh` or `setenv.bat`.

See instructions for [Increasing JIRA application memory](#) - enter it in the same place as PermGen Memory.

## 2. Selecting your JMX Client

You need to use a JMX client in order to view the JMX output from Confluence. [JConsole](#) is a readily available JMX client that is included with the [supported](#) Java Developer Kit (version 5 onwards). The full name is the 'Java Monitoring and Management Console', but we will refer to it as JConsole for the purposes of this document.

## 3. Adding the JMX Client to your Path

You must add the location of the JConsole binary file to your path environment variable. As JConsole resides in the 'bin' (binaries) folder under your Java directory, the path should resemble something like this:

```
JDK_HOME/bin/
```

In this example, replace 'JDK\_HOME' with the full system path to your Java directory.

## 4. Configuring JConsole

To configure JConsole:

1. Run the JConsole application.
2. You will be prompted to create a new connection. Choose **remote process** and enter the hostname of your Confluence instance and a port of your choosing.

To connect easily, add the startup parameters to setenv.bat or setenv.sh:  
-Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=8086  
-Dcom.sun.management.jmxremote.authenticate=false  
Port 8086 is unlikely to be used. Then, connect remotely using port 8086.

JConsole, or any JMX client, will not see applications which are not owned by the same user. For example under Windows, if an application is started as a service, it is the System User which owns the process, and not the Current User.

3. Click **Connect**.

Note: Other JMX clients besides JConsole can read JMX information from Confluence.

### What can I monitor with JMX?

The JMX interface allows you to see live internal information from your Confluence instance, via the following MBeans:

#### IndexingStatistics

This MBean shows information related to search indexing.

Property name	Function	Values
Flushing	Shows state of cache (i.e. flushing, or not).	True/False
LastElapsedMilliseconds	Time taken during last indexing.	Milliseconds
LastElapsedReindexing	Time taken during last re-indexing.	Milliseconds
TaskQueueLength	Shows number of tasks in the queue.	Integer

#### SystemInformation

This MBean shows information related to database latency. It also contains most of the information presented on the [System Information](#) page.

Property name	Function	Values
DatabaseExampleLatency	Shows the latency of an example query performed against the database.	Milliseconds

**RequestMetrics**

This MBean shows information related to system load and error pages served.

Property name	Function	Values
AverageExecutionTimeForLastTenRequests	Average execution time for the last ten requests.	Milliseconds
CurrentNumberOfRequestsBeingServed	Number of requests being served at this instant.	Integer
ErrorCount	Number of times the Confluence error page was served.	Integer
NumberOfRequestsInLastTenSeconds	Obviously, the Number Of Requests In the Last Ten Seconds.	Integer

**MailServer-SMTPServer**

This MBean shows information related to email dispatch attempts and failures. There will be an MBean for every SMTP Mailserver that has been configured in the Confluence instance.

Property name	Function	Values
EmailsAttempted	The number of email messages Confluence has tried to send.	Integer
EmailsSent	The number of email messages sent successfully.	Integer

**MailTaskQueue**

This MBean shows information related to the email workload.

Property name	Function	Values
ErrorQueueSize	Number of errors in the queue.	Integer
Flushing	Shows state (i.e. flushing, or not)	True/False
FlushStarted	Time that operation began.	Time
RetryCount	The number of retries that were performed.	Integer
TaskSize	Number of email messages queued for dispatch.	Integer

**SchedulingStatistics**

This MBean shows information related to current jobs, scheduled tasks and the time that they were last run.

**High CPU consuming threads**

For Java 1.6, add the [Top Threads Plugin](#) to monitor whether CPU is spiking. Download it to a directory and run JConsole like this:

```
JConsole -pluginpath /path/to/topthreads.jar
```

This works only with JDK 1.6, but that can be on the remote machine if the server is running a lower version.

Please note, adding live monitoring to a production instance may itself have an impact on performance.

**Tracking Customisations Made to your Confluence Installation**

The **'Modification'** section of the Confluence **'System Information'** screen lists the files that have been changed since your Confluence application was installed. You will find this information particularly useful when upgrading Confluence to a new version, because you will need to re-apply all customisations after the upgrade.

To see the modifications made to files in your Confluence installation:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Select **'System Information'** in the 'Administration' section of the left-hand panel.
3. Scroll down to the section titled **'Modification'**.

*Screenshot: Modifications tracker on the Confluence System Information screen*

Modification	
Modified	decorators/main.vmd, pages/page-breadcrumbs.vmd, template/includes/macros.vmd, decorators/mail.vmd, decorators/space.vmd, template/includes/personal-sidebar.vmd
Removed	No files removed

#### Notes

- The modification tracker does not detect changes to class files from the `confluence.jar` or other JAR files. If you modify classes, the Confluence modification detection does not report the modification. See issue [CONF-20993](#).

#### Viewing Site Statistics

Note that the site activity information is **disabled by default**. See notes [below](#).

If enabled, the global activity screen displays statistics on the activity in your Confluence site.

These include:

- How many pages and blog posts have been viewed, added or updated over a given period.
- Which spaces are the most popular (most frequently viewed).
- Which spaces are the most active (most frequently edited).
- Which people are the most active contributors/editors of content.

#### Related pages:

- [How do I get more statistics from Confluence?](#)
- [Cache Statistics](#)
- [View Space Activity](#)
- [Live Monitoring Using the JMX Interface](#)

#### To view the activity on your site:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **'Global Activity'** in the 'Administration' section of the left-hand panel (only appears if enabled - see [below](#)).

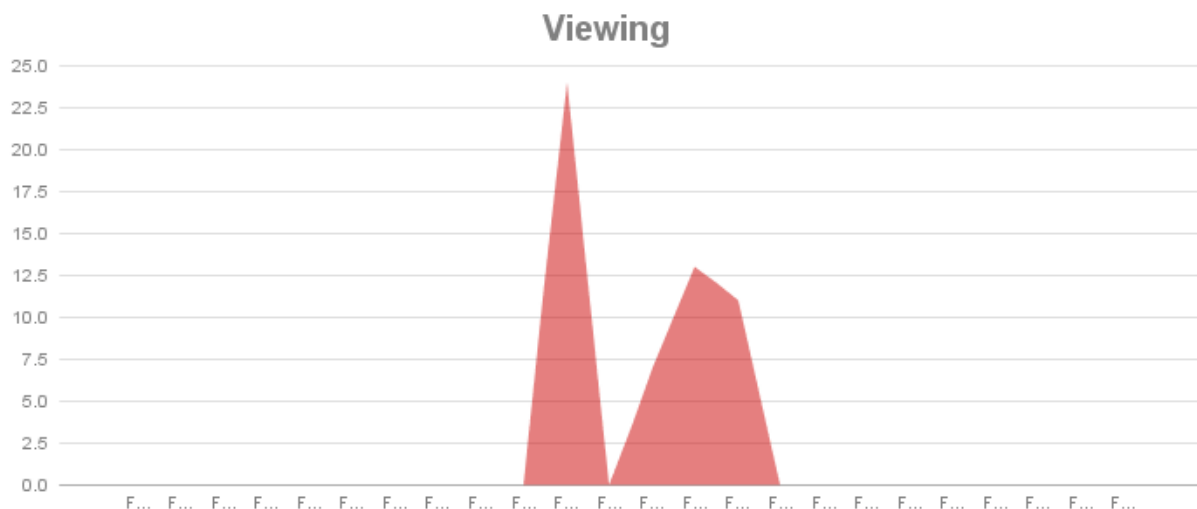
*Screenshot: Global Activity*

## Global Activity

Activity for 15 February 2013

Period: [month](#) | [week](#) | [day](#)

Day: [<< Previous](#) | [Now](#) | [Next >>](#)



This graph shows how many times pages and blog posts have been viewed over the current time period.



This graph shows how many pages, blog posts and comments have been created or updated over the current time period.

#### Most popular spaces (Views)

1. Angry Nerds (39)
2. Documentation (8)
3. Sample Space (6)
4. Stars (1)
5. Josh User (1)

#### Most active spaces (Edits)

1. Angry Nerds (16)
2. Documentation (4)
3. Sample Space (2)

#### Most active contributors (Edits)

1. Rach Admin (8)
2. Josh User (7)
3. Sophie Staunton (4)
4. Ewan User (3)

The top ten most popular and most active pages and/or blog posts will be listed, with a link to each.

#### Notes

- The [Confluence Usage Stats plugin](#), which provides the 'Global Activity' screen, is known to cause performance problems on large installations. This plugin is **disabled by default**. A status report on the progress of the performance issues with this plugin is available in this issue: [USGTRK-15](#).
- Your Confluence system administrator can enable the plugin, but please be aware of the possible impact upon your site's performance.
- The plugin is sometimes called 'Confluence Usage Tracking'.
- If your Confluence site is clustered, the global activity information will not be available.

## Viewing System Properties

After adding memory, setting a proxy, or changing other Java options, it can be difficult to diagnose whether the system has picked them up. This page tells you how to view the system properties that your Confluence site is using.

You can see the expanded system properties on the 'System Information' screen of the Confluence Administration Console. You do not need to restart Confluence before viewing the information.

To see the system properties recognised by your Confluence installation:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **System Information** in the left-hand panel.
3. Scroll down to the section titled **System Properties**.

## Configuring the Server Base URL

The **Server Base URL** is the URL via which users access Confluence. The base URL **must** be set to the same URL by which browsers will be viewing your Confluence site.

Confluence will automatically detect the base URL during setup, but you may need to set it manually if your site's URL changes or if you set up Confluence from a different URL to the one that will be used to access it publicly.

You need to have [System Administrator](#) permissions in order to perform this function.

### To configure the Server Base URL:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **General Configuration** in the left-hand panel
3. Choose **Edit**
4. Enter the new URL in the **Server Base URL** text box
5. Choose **Save**

## Example

If Confluence is installed to run in a non-root context path (that is, it has a context path), then the server base URL should include this context path. For example, if Confluence is running at:

```
http://www.foobar.com/confluence
```

then the server base URL should be:

```
http://www.foobar.com/confluence
```

## Notes

- **Using different URLs.** If you configure a different base URL or if visitors use some other URL to access Confluence, it is possible that you may encounter errors while viewing some pages.
- **Changing the context path.** If you change the context path of your base URL, you may also need to edit the web server's `server.xml` file to reflect the new path:
  1. Stop the Confluence server.
  2. Go to your Confluence 'destination directory'. This is the directory where the Confluence installation files are stored. For example, `C:\Program Files\Atlassian\Confluence`. Let's call this directory '{CONFLUENCE\_INSTALLATION}'.
  3. Edit the configuration file at `{CONFLUENCE_INSTALLATION}\conf\server.xml`.
  4. Change the value of the `path` attribute in the `Context` element to reflect the context path. For example, if Confluence is running at `http://www.foobar.com/confluence`, then your `path` attribute should look like this:

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

5. Save the file.
- **Proxies.** If you are running behind a proxy, ensure that the proxy name matches the base URL. For example: `proxyName="foobar.com" proxyPort="443" scheme="https"`. This will make sure we are passing the information correctly.
  - This information needs to be added in the `Connector` element at `{CONFLUENCE_INSTALLATION}\conf\server.xml`.

## Configuring the Confluence Search and Index

Confluence administrators can adjust the behaviour of the Confluence search, and manage the index used by the search.

- [Configuring Indexing Language](#)
- [Configuring Quick Navigation](#)
- [Content Index Administration](#)
- [Enabling OpenSearch](#)
- [Rebuilding the Ancestor Table](#)
- [Setting Up Confluence to Index External Sites](#)
- [Setting Up an External Search Tool to Index Confluence](#)

### Related pages:

- [Search](#)
- [Confluence Administrator's Guide](#)


## Configuring Indexing Language

Changing the indexing language defined in Confluence may improve the accuracy of Confluence search results, if the majority of the content of your site is in some language other than English. Confluence supports content indexing in English (default), German, Russian, Chinese, CJK, Custom Japanese, French, Brazilian, Czech and Greek.

### Related pages:

- [Choosing a Default Language](#)
- [Installing a Language Pack](#)
- [Content Index Administration](#)
- [How to rebuild the content indexes from scratch](#)

### To configure the indexing language:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration
2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Select the **Indexing Language** from the dropdown list in the **Formatting and**

**International Settings** section.

5. Choose **Save**.

## Configuring Quick Navigation

When a user is searching Confluence (see [Searching Confluence](#)) the quick navigation aid automatically offers a dropdown list of pages and other items, matched by title to the search query. By default, this feature is enabled, with the maximum number of simultaneous quick navigation requests set to 40. These options can be modified as described below.

The maximum number of simultaneous quick navigation requests defines the maximum number of individuals who can use this feature simultaneously on the same Confluence server. If your Confluence server serves a large number of individuals who use this feature regularly, some of whom are being denied access to it, you may wish to increase this value.

To configure quick navigation:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Further Configuration** in the left-hand panel.

3. Choose **Edit**.

4. To disable Quick Navigation, deselect the **Quick Navigation** checkbox.

5. To modify the maximum number of simultaneous quick navigation requests, enter the appropriate number in the field beside **Max Simultaneous Requests**.

6. Choose **Save**.

### Related pages:

- [Search](#)

## Content Index Administration

The content index, also called the search index, supports Confluence's search functionality. It is also used for a number of related functions such as building email threads in the mail archive, the [space activity](#) feature, and lists of recently-updated content. The [Gliffy plugin](#) also uses the index for some of its functionality.

For reasons of efficiency, Confluence does not immediately add content to the index. New and modified Confluence content is first placed in a queue and the queue is processed once every five seconds (by default).

### Viewing the content index summary

**To see information about your Confluence site's content indexing:**

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **'Content Indexing'** under the heading **'Administration'** in the left-hand panel.

*Screenshot: Index summary*

### On this page:

- [Viewing the content index summary](#)
- [Rebuilding the search index](#)
- [The 'Did You Mean' index is no longer relevant](#)
- [Slow reindexing](#)
- [Viewing the index browser](#)
- [More hints and tips](#)

### Related pages:

- [Scheduled Jobs](#)
- [Search](#)
- [Configuring the Confluence Search and Index](#)

## Search Index

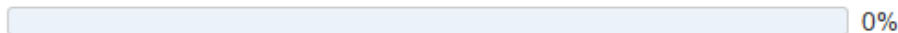
The search index allows searching of Confluence content. If you are having troubles with search, you may need to rebuild the search index. Please note, rebuilding the search index can severely affect the performance of your instance - it can take hours for some large instances.

**BUILT**

100%

## Did You Mean Index

You will need to build this index to make "Did You Mean" work. After this has finished, "Did You Mean" will be automatically turned on. **Please note, this feature only provides suggestions for the English language.**

**NOT BUILT**

0%

### Rebuilding the search index

The search index is maintained automatically, but you may need to rebuild it manually under circumstances such as these:

- Your searching and mail threading are malfunctioning.
- After an upgrade. If a re-index is required after an upgrade, it will be noted in an upgrade subsection of the relevant [Confluence Release Notes](#).

### To rebuild the search index:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **'Content Indexing'** under the heading **'Administration'** in the left-hand panel.
3. Choose the **'Rebuild'** button in either the 'Search Index' section.  
(If the indexes has never been built, its button will indicate 'Build' instead of 'Rebuild.')

*Screenshot: Content indexing*

## Search Index

The search index allows searching of Confluence content. If you are having troubles with search, you may need to rebuild the search index. Please note, rebuilding the search index can severely affect the performance of your instance - it can take hours for some large instances.

IN PROGRESS...

0% (Elapsed Time: 00:00:00)

Rebuild

## Did You Mean Index

You will need to build this index to make "Did You Mean" work. After this has finished, "Did You Mean" will be automatically turned on. Please note, this feature only provides suggestions for the English language

**'Did You Mean' index build temporarily disabled while Search Index build is in progress.**

NOT BUILT

0%

### The 'Did You Mean' index is no longer relevant

The 'Did You Mean' feature is no longer available in Confluence. This index is therefore redundant, and will be removed at some time in the future.

### Slow reindexing

Does the reindexing take a long time to complete? The length of time depends on the following factors:

- Number of pages in your Confluence instance.
- Number, type and size of attachments.
- Amount of memory allocated to Confluence.
- Disk throughput.

It may help to increase the heap memory allocation of Confluence by following the instructions here [Increasing JIRA application memory](#). The process is basically the same for Confluence or JIRA applications.

If you are running an older version of Confluence and find that the index rebuild is not progressing, you may need to shut down Confluence, and restart it with the following *Java system property* set: `bucket.indexing.threads.fixed=1`. This will cause the re-indexing to happen in a single thread and be much more stable (but slower).


### Viewing the index browser

Confluence uses a search engine called Lucene. If you need to see more details of the indexed pages in your Confluence site, you can download and run [Luke](#). Luke is a development and diagnostic tool that accesses existing Lucene indexes and allows you to display and modify their content in several ways.

Start Luke and use it to open the `index` directory, located in your [Confluence Home directory](#). For example: `c:\confluence\data\confluence-home\index`.

**Note:** Confluence 5.2 (and later) use Lucene 4.3 (or later). If the [Luke](#) library has not been updated to support the latest version of Lucene, you can compile Luke yourself, from the [fork on Github](#) – please read the warnings and notes in the README file of that repository.

### More hints and tips

- If you are still experiencing problems after performing the above rebuild, the next step might be to [remove the index and rebuild it from scratch](#).
-  The [space activity](#) feature uses the index to store data. If you remove the index file, the existing activity data will disappear.
- A tip for the development community: If you have the Confluence source, you can look for references to the SmartListManager to find the screens and lists that rely on the content index.

## Enabling OpenSearch

With OpenSearch autodiscovery, you can add Confluence search to your Firefox or IE7 search box (see [Searching Confluence from your Browser's Search Box](#)). By default, OpenSearch autodiscovery is enabled. This feature can be enabled or disabled as described below.

### To enable or disable OpenSearch autodiscovery:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Further Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Select the **Open Search** checkbox to enable this feature (deselect to disable).
5. Choose **Save**.

### Related pages:

- [Search](#)
- [Confluence Administrator's Guide](#)

## Rebuilding the Ancestor Table

This process could take hours in a production instance and new content may not be save-able during this process. **It is highly recommended that you take a full backup of your database and then execute this process during a maintenance window.**

In Confluence, the ancestor table defines what pages are ancestors or descendants of other pages (which can be used by search restrictions with the ancestorids restriction). Occasionally, the ancestor table will become out of sync. When this happens, you can rebuild the table to restore everything to normal.

Access this URL:

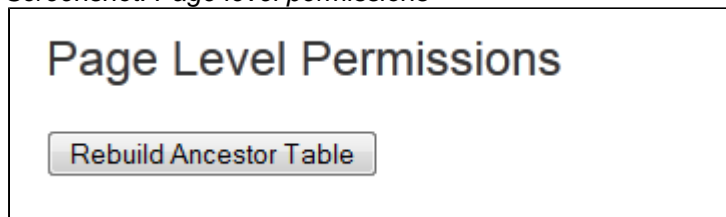
```
http://yoursite/admin/permissions/pagepermsadmin.action
```

### Related pages:

- [Search](#)
- [Confluence Administrator's Guide](#)

After rebuilding the ancestor table, you'll need to flush the "Inherited Content Permissions" cache in [Cache Statistics](#), otherwise the inherited permissions may not be applied immediately to all pages. You may also need to [rebuild the content index](#) so that the permissions take effect in search results.

*Screenshot: Page level permissions*



**Note:** There was a known issue with accessing this option from Confluence 5.4.1 to Confluence 5.5.6. See

[CONF-32174](#) - Rebuild Ancestor Table is Inaccessible

**RESOLVED**

## Setting Up Confluence to Index External Sites

Confluence cannot easily index external sites due to [technical reasons](#), but there are two alternatives:

1. [Embed External Pages Into Confluence](#)
2. [Replace Confluence Search](#)

### Related pages:

- [Setting Up an External Search Tool to Index Confluence](#)
- [Configuring the Confluence Search and Index](#)

### Technical reasons

Confluence indexes pages using a customised [Lucene](#) search engine that returns matching pages, mail and blog posts for which the searcher has view permission. It would require significant source code modifications to enable Confluence to process search results from external pages, as the indexing process has been customised to utilise internal Confluence metadata. Note that users can still [index content from new attachment filetypes](#).

### Embedding external pages into Confluence

If you only have a small number of external sites to index, you may prefer to enable the [HTML-include Macro](#) and use it embed the external content inside normal Confluence pages.

### Replacing the Confluence search

Use your own programmer resources to replace Confluence's internal search with a crawler that indexes both Confluence and external sites. This advanced option is easier than modifying the internal search engine. It requires removing Confluence internal search from all pages and replacing the internal results page with your own crawler front-end.

1. Setup a replacement federated search engine to index the Confluence site, as well as your other sites, and provide the results that way. You would need to host a web crawler, such as these [open-source crawlers](#). Note that you can perform a search in Confluence via the [remote API](#)
2. Replace references to the internal search by [modifying the site layout](#) so that it links to your search front-end
3. Host another site containing the search front-end. You may wish to insert it into a suitable context path in your application server so that it appears to be from a path under Confluence. Tomcat sets Confluence's paths from the Confluence install\confluence\WEBINF\web.xml file.

## Setting Up an External Search Tool to Index Confluence

Any web crawler can be configured to index Confluence content, for example the Google Search Appliance or similar. If a login is required to view content that will be indexed, you should create a Confluence user specifically for the search crawler to use. Grant this user view rights to all content you wish to index, but deny that user all delete and administration rights. This ensures that an aggressive crawler will not be able to perform actions that could modify the site.

External applications can also use the search function in the [Confluence remote API](#).

### Related pages:

- [Setting Up Confluence to Index External Sites](#)
- [Configuring the Confluence Search and Index](#)

## Configuring Mail

- [Configuring a Server for Outgoing Mail](#)
- [Setting Up a Mail Session for the Confluence Distribution](#)
- [Configuring the Recommended Updates Email Notification](#)
- [The Mail Queue](#)
- [Customising the eMail Templates](#)

### Configuring a Server for Outgoing Mail

Configuring your Confluence server to send email messages allows your Confluence users to:

- [Receive emailed notifications and daily reports of updates.](#)
- [Send a page via email.](#)

You can personalise email notifications by configuring the 'From' field to include the name and email address of the Confluence user who made the change.

You need [System Administrator](#) permissions in order to configure Confluence's email server settings.

**On this page:**

- [Configuring Confluence to send email messages](#)
- [Testing the email settings](#)

**Related pages:**

- [The Mail Queue](#)
- [Setting Up a Mail Session for the Confluence Distribution](#)

## Configuring Confluence to send email messages

### To configure Confluence to send outgoing mail:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Select **Mail Servers** under **Configuration** in the left-hand panel. This will list all currently configured SMTP servers.
3. Click **Add New SMTP Server** (or **edit** an existing server).
4. Edit the following fields as required:
  - **Name:** By default, this is simply 'SMTP Server'.
  - **From Address:** Enter the email address that will be displayed in the 'from' field for email messages originating from this server. This field is mandatory. You will not be able to complete the Confluence mail server configuration until this field has been specified.
  - **From Name:** Enter the name that will be displayed in the 'from' field for email messages originating from this server. This is the text which appears before the user's registered email address (in square brackets). This field accepts the following variables, which reference specific details defined in the relevant Confluence user's profile:

Variable	Description
<code>\${fullname}</code>	The user's full name.
<code>\${email}</code>	The user's email address.
<code>\${email.hostname}</code>	The domain/host name component of the user's email address.

The default is '`${fullname}` (Confluence)'.

Hence, if Joe Bloggs made a change to a page he was watching and the Confluence site's 'From Address' was set to `confluence-administrator@example-company.com`, then the 'From' field in his email notification would be: Joe Bloggs (Confluence)  
<confluence-administrator@example-company.com>.

- **Subject Prefix:** Enter some text to appear at the beginning of the subject line.
5. Enter your **Hostname**, **Port**, **User name** and **Password** details.

If your SMTP host uses the Transport Layer Security (TLS) protocol select **Use TLS**.

**OR**

Specify the **JNDI location** of a mail session configured in your application server. For more information on how to set up a JNDI mail session, see [Setting Up a Mail Session for the Confluence Distribution](#).

### Testing the email settings

A Confluence administrator can test the email server as follows:

1. Set up a mail server as described above.
2. Click **Send Test Email** to check that the server is working. Check that you get the test email in your inbox.
3. You can flush the email queue to send the email message immediately. Go to **Mail Queue**, and click **Flush Mail Queue**. See [The Mail Queue](#).

A user can test that notifications are working as follows:

1. Go to your user profile (using the **Settings** link) and edit your email preferences. See [Email Notifications](#).
2. Enable **Notify On My Actions**. (By default, Confluence does not send you notifications for your own changes.)
3. Go to a page you wish to get notifications about.
4. Choose **Watch** at the top-right of the page. See [Watch Pages, Spaces and Blogs](#).
5. Edit the page, make a change, and save the page.
6. Check your email inbox. You may need to wait a while for the email message to arrive.

### Setting Up a Mail Session for the Confluence Distribution

The simplest way to set up a mail server through the Confluence Administration console. See [Configuring a Server for Outgoing Mail](#).


If you want to add different options or parameters you can also set up a mail session for the Confluence distribution. In the example below we'll set up Gmail.

To set up a mail session for the Confluence distribution:

1. Stop Confluence.
2. Move (don't copy) mail-x.x.x.jar from <confluence-install>/confluence/WEB-INF/lib to <confluence-install>/lib.  
**Note:** where x.x.x. represents the version numbers on the jar files in your installation. Do not leave a renamed backup of the jar files in /confluence/WEB-INF/lib. Even with a different file name, the files will still be loaded as long as it remains in the directory
3. Add the following to your server.xml file found in <confluence-install>/conf/ (add it just before the </Context> tag):

#### For Confluence 3.5.x

```
<Resource name="mail/GmailSMTPServer"
 auth="Container"
 type="javax.mail.Session"
 mail.smtp.host="smtp.gmail.com"
 mail.smtp.port="465"
 mail.smtp.auth="true"
 mail.smtp.user="yourEmailAddress@gmail.com"
 password="yourPassword"
 mail.smtp.starttls.enable="true"
 mail.transport.protocol="smtps"
 mail.smtp.socketFactory.class="javax.net.ssl.SSLSocketFactory"
/>
```

4. Restart Confluence.
5. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration
6. Choose **Mail Servers**.
7. Choose either **Edit an existing configuration**, or **Add a new SMTP mail server**.
8. Edit the server settings as necessary, and set the **JNDI Location** as:

```
java:comp/env/mail/GmailSMTPServer
```

Note that the JNDI Location is case sensitive and must match the resource name specified in server.xml.

9. Save your changes and send a test email.

### Configuring the Recommended Updates Email Notification

Confluence sends a regular email report to subscribers, containing the top content that is relevant to the person receiving the message, from spaces they have permission to view. This is called the 'Recommended Updates' notification.

If you have Confluence Administrator or System Administrator [permissions](#), you can configure the default settings that determine how often the Recommended Updates notification is sent. When new users are added to Confluence, the default settings will be applied to their user profiles.

Confluence users can choose their personal settings, which will override the defaults. See [Email Notifications](#).

#### Initial settings of the defaults

When you install Confluence, the initial values of the default settings are as follows:

- The default frequency is weekly.
- If your Confluence site has public signup enabled, the Recommended Updates notification is disabled by default. If public signup is not enabled, the notification is enabled by default.

You can change the above settings, specifying a different default value for the site.

#### Notes:

- The Recommended Updates notification is sent only to people who have a user profile in Confluence. If your Confluence site uses external user management, such as LDAP, then people will receive the report only after they have logged in for the first time. (The first login creates their user profile.)
- The daily email message is sent at 1 p.m. in the user's configured time zone.
- The weekly email message is sent at 1 p.m. on Thursdays in the user's configured time zone.

#### On this page:

- [Initial settings of the defaults](#)
- [Configuring the Recommended Updates notification](#)
- [Disabling the Recommended Updates notification for the entire site](#)

#### Related pages:

- [Email Notifications](#)
- [Confluence Administrator's Guide](#)

### Configuring the Recommended Updates notification

You can set the the default send option (send / do not send) and the default schedule (daily or weekly).

#### To configure the Recommended Updates email notification:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Click **Recommended Updates Email** in the left-hand panel.

#### Disabling the Recommended Updates notification for the entire site

You can also turn off the **recommended updates** notification for the entire site, by disabling the 'Confluence daily summary email' plugin. See [Disabling and enabling add-ons](#).

## The Mail Queue

Email messages waiting to be sent are queued in a mail queue and periodically flushed from Confluence once a minute. A Confluence administrator can also manually flush messages from the mail queue.

If there is an error sending messages, the failed email messages are sent to an error queue from which you can either try to resend them or delete them.

### To view the mail queue:

1. Choose the **cog icon**




, then choose **General Configuration** under Confluence Administration

2. Choose **Mail Queue** in the left-hand panel. This will display the email messages currently in the queue.
3. Choose **Flush Mail Queue** to send all email messages immediately.
4. Choose **Error Queue** to view failed email messages. You can try to **Resend** the messages, which will flush the mails back to the mail queue, or you can **Delete** them from here.

#### Related pages:

- [Configuring a Server for Outgoing Mail](#)
- [Setting Up a Mail Session for the Confluence Distribution](#)

 *The information on this page does not apply to Confluence Cloud.*

## Configuring Character Encoding

This page explains the encoding settings that are applicable in Confluence and how they relate to application behaviour.

To avoid problems with character encoding, make sure the encoding used across the different components of your system are the same. In general, **always set all character encodings to UTF-8**:

- Confluence character encoding – described below.
- Database – see [Configuring Database Character Encoding](#).
- Application server – see [Configuring URL Encoding on Tomcat Application Server](#)

### Configuring the Confluence character encoding

By default, Confluence uses UTF-8 character encoding to deliver its pages.

**Note:** While it is possible to change the character encoding, we recommend that you leave this as it is unless you are certain of what you are doing.

In summary: Changing the Confluence character encoding will change your HTTP request and response encoding and your filesystem encoding as used by exports and Velocity templates.

### To change the Confluence character encoding via the UI:

#### On this page:

- [Configuring the Confluence character encoding](#)
- [More details about character encoding](#)
  - [Java character encoding](#)
  - [Confluence character encoding](#)
  - [Database encoding](#)
  - [Filesystem encoding](#)
- [Problems with character encodings](#)
- [Notes](#)

#### Related pages:

- [Configuring Confluence](#)
- [Confluence Administrator's Guide](#)
- [Application Server Configuration](#)
- [Database Configuration](#)

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Enter the new character encoding of your choice in the text box next to **Encoding**.
5. Choose **Save**.

**Note:** At runtime, the character encoding is available in `Settings.defaultEncoding`.

### More details about character encoding

There are three places where character encoding matters to Confluence:

1. **Database encoding** - usually the most important; it is where almost all user data is stored.
2. **Filesystem encoding** - important for attachment storage (pre-2.2), reading Velocity templates and writing exported files.
3. **HTTP request and response encoding** - important for form parsing, correct rendering by the browser and browser interpretation of encoded URLs.

Problems generally arise when Confluence thinks one of the above encoding is different to what it actually is. For example, Confluence might believe the database is using ISO-8859-1 encoding, when in fact it is UTF-8 encoded.

In certain cases (for example, Microsoft Windows), it might not be possible to use a fully Unicode filesystem (that is, a default Windows installation does not support Unicode filenames properly). If so, keep UTF-8 for the other two and be aware that your operating system might have limitations around international attachments (pre-2.2), backup and restore of international data, etc.

### Java character encoding

Java *always* uses the multibyte UTF-16 character encoding for all `String` data\*. This means that each of the encodings above defines how, at that particular point, characters are converted to and from Java's native UTF-16 format into some other format that the browser, filesystem or database might understand.

So when a request comes in to Confluence, we convert it from the request encoding to UTF-16. Then we store that data into the database, converting from UTF-16 to the database's encoding. Retrieving information from the database and sending it back to the browser is the same process in the opposite direction.

\*A `char` represents single Unicode code point from the Base Multilingual Plane (BMP), encoded as UTF-16. Multiple `chars` are used as surrogate pairs for characters beyond U+FFFF.

### Confluence character encoding

The Confluence character encoding is used in the following parts of the system:

- `ConfluenceWebWorkConfiguration` sets `webwork.i18n.encoding` to the this encoding, which `WebWork` uses in the response `Content-Type` header.
- `AbstractEncodingFilter` sets the HTTP request encoding to this encoding. This seems unnecessary, since the `Content-Type` header from the client should include the encoding used. This affects form submissions and file uploads.
- `VelocityUtils` reads in Velocity templates using this encoding when reading templates from disk.
- `AbstractXmlExporter` creates its output using this encoding.
- `GeneralUtil` uses this encoding when doing `URLEncode` and `URLDecode`. Different browsers have different support for character sets in URLs, so it's uncertain how much benefit this provides.

See [Configuring Confluence Character Encoding](#) (described above.)

### Database encoding

The database encoding is the responsibility of your JDBC drivers. The drivers are responsible for reading and writing from the database in its native encoding and translating this data to and from Java Strings (which

are UTF-16). For some drivers, such as [MySQL](#), you must set Unicode encoding explicitly in the JDBC URL. For others, the driver is smart enough to determine the database encoding automatically.

Ideally, your database itself should be in a Unicode encoding (and we recommend [doing this](#) for the simplest configuration), but that is not necessary as long as:

- the database encoding supports all the characters you want to store in Confluence
- your JDBC drivers can properly convert from the database encoding to UTF-16 and vice-versa.

See [Configuring Database Character Encoding](#).

### Filesystem encoding

The filesystem encoding is mostly ignored by Confluence, except for the cases where the above configuration setting above plays a part (exports, velocity). When attachments are uploaded, they are written as a stream of bytes directly to the filesystem. It is the same when they are downloaded: the bytes from the file `InputStream` are written directly to the HTTP response.

In some places in Confluence, we use the *default filesystem encoding* as determined by the JVM and stored in the `file.encoding` system property (it can be overridden by setting this property at startup). This encoding is used by the Java `InputStreamReader` and `OutputStreamWriter` classes by default. This encoding should probably never be used; for consistent results across all filesystem access we should be using the encoding set in the General Configuration.

In certain cases we explicitly hard-code the encoding used to read or write data to the filesystem. Two important examples are:

- importing Mbox mailboxes which are known to be ISO-8859-1
- Confluence Bandana config files are always stored as UTF-8.

Some application servers, Tomcat for example, have an encoding setting that modifies Confluence URLs before they reach the application. This can prevent access to international pages and attachments (really anything with international characters in the URL). See [configuring your Application Server URL encoding](#).

### Problems with character encodings

If Confluence has the wrong idea about encoding for one of the above, it manifests itself in different ways:

1. Incorrect database encoding - user data is corrupted between saving and restoring from the database. This often happens after a delay, as we cache data as it is written to the database and only later retrieve the corrupted copy from the database.
2. Incorrect/non-Unicode filesystem encoding - international filenames break attachment download/upload/removal (pre-2.2); exports break with international content or attachments.
3. Incorrect HTTP encoding - incorrect encoding selected by browser, resulting in incorrect rendering of characters. Changing browser encoding causes page to render properly. Broken URLs when linking to pages or attachments with non-ASCII characters.

See [Troubleshooting Character Encodings](#).

### Notes

- Mac users please note that **MacRoman** encoding is compatible with UTF-8. You do not need to change your encoding settings if you are already using MacRoman.
- This is a good article by Joel Spolsky: [The Absolute Minimum Every Software Developer Absolutely, Positively Must Know About Unicode and Character Sets \(No Excuses!\)](#)

### Troubleshooting Character Encodings

Often users may have problems with certain characters in a Confluence instance. Symptoms may include:

- Non-ASCII characters appearing as question marks (?)
- Page links with non-ASCII characters not working
- Single characters being displayed as two characters
- Garbled text appearing

In most cases, it is due to a mis-configuration in one of the components that Confluence uses.

Follow these steps to diagnose the problem.

### 1. Run the encoding test

Confluence includes an encoding test that can reveal problems with your configuration.

To perform the test, access the Encoding Test page via the `<confluence base-url>/admin/encodingtest.action` page on your Confluence instance. You will be required to copy and paste a line of text and submit a form. The test will take the text and pass it through Confluence, the application server and the database, and return the results.

You should also test pasting some sample text (Japanese for example) if you are experiencing problems with a specific language.

Example:

```
http://confluence.atlassian.com/admin/encodingtest.action
```

or

```
http://<host address>:<port>/admin/encodingtest.action
```

If the text displayed in the encoding test is different to what was entered, then there are problems with your character encoding settings.

A successful test looks like the following:

*Screenshot: Successful encoding test*

## Character Encoding Test Results

The encoding test has now been run. Below, you can compare the raw text delivered from Confluence against the text returned by your browser in web forms, and the text as it appears after a round-trip through the database. All the test results should appear identical.

**İñtërnâtiônàlizætiøn** This image is how the sample text below *should* appear. If it does not, please file a support request at <http://support.atlassian.com>, including a screenshot of this page, and all of your [System Information](#).

### Test 1: Raw text

This is the test string generated in Confluence

İñtërnâtiônàlizætiøn

### Test 2: Form submission

This is the test string pasted by you into the web form and submitted back to Confluence

İñtërnâtiônàlizætiøn

### Test 3: Database round-trip (select as lower-case)

This is the string from Test 2 after being stored in the database and then retrieved as lower-case

ïñtërnâtiônàlizætiøn

**Expected result (converting Java string to lowercase)**

ïñtërnâtiônàlizætiøn

### Test 4: Database round-trip (select as upper-case)

This is the string from Test 2 after being stored in the database and then retrieved as upper-case

İÑTËRNÂTIÔNÀLIZÆTIØN

**Expected result (converting Java string to uppercase)**

İÑTËRNÂTIÔNÀLIZÆTIØN

### Test 5: International file name support

**Try to write a file to the confluence home directory with the test string as the file name**

✔ **File was written successfully**

### Test 6: Detect international file name mangling

**Detect whether the file system is mangling the file name when it is saved**

✔ **The file name has been preserved**

### MySQL 3.x

MySQL 3.x is known to have some problems with the upper- and lower-casing of some characters, and may fail the last two tests. For more information, see [MySQL 3.x Character Encoding Problems](#).

## 2. Ensure the same encoding is used across all components

As mentioned in the [Configuring Encoding](#) document, the same character encoding should be used across the database, application server and web application (Confluence).

- To change the character encoding used in **Confluence**, see [Configuring Character Encoding](#).
- To change the character encoding used in the **application server**, please ensure you set the [Application Server URL encoding](#) and view your application server's documentation on any other settings required to enable your encoding.
- To change the character encoding used in the **database**, see [Configuring Database Character Encoding](#).

## 3. Requesting support

If there are still problems with character encoding after following the above steps, [create a support request](#), and our support staff will aid in solving your problem.

Entering in the following details will help us to identify your problem:

- Attach screenshots of the problem
- Attach the results of the encoding test (above)

- Select which application server (and version) you are using
- Select which database (and version) you are using
- Copy the contents of the [System Information](#) page into the 'Description' field

### "€" Euro character not displaying properly

The € (euro) symbol is a three byte character, with byte values in file (UTF-8) of 0xE2, 0x82, 0xAC.

Sometimes, if the character encoding is not set consistently among all participating entities of the system, Confluence, server and the database, one may experience strange behaviour.

...

I write a page with a Euro sign in it (€). All is well, the Euro sign shows up in the wiki markup text-box, and the preview, and the display of the saved page.

One day later, the Euro sign has changed into a question mark upside down!

...

What is going on? Why does the Euro sign mysteriously change? How do I prevent it?

Interestingly enough the character encoding test passes with no problems, demonstrating that Confluence and the connected Database both recognise the € symbol.

There are two potential reasons for this behaviour:

*Database and Confluence is using utf-8 encoding. The connection is not.*

When data transferred to it via the connection which does not use utf-8 encoding gets encoded incorrectly. Hence, updating the connection encoding may resolve this problem from now on, yet it probably would not affect already existing data.

*Database is not using utf-8. Confluence and your connection are.*

If your Database encoding is not set to UTF-8, yet is using some other encoding such as *latin1*, it could be one of the potential reasons why you lose the "€" characters at some stage. It could be occurring due to **caching**. When Confluence saves data to the database, it may also keep a local cached copy. If the database encoding is set incorrectly, the Euro character may not be correctly recorded in the database, but Confluence will continue to use its cached copy of that data (which is encoded correctly). The encoding error will only be noticed when the cache expires, and the incorrectly encoded data is fetched from the database.

For instance the *latin1* encoding would store and display all 2-byte UTF8 characters correctly except for the euro character which is replaced by '?' before being stored. As Confluence's encoding was set to UTF-8, the 2-byte UTF-8 characters were stored in *latin1* database assuming that they were two *latin1* different characters, instead of one utf8 character. Nevertheless, this is not the case for 3-byte utf8 characters, such as the Euro symbol.

Please ensure that you set the character encoding to UTF-8 for all the entities of your system as advised in this [guide](#).

### MySQL 3.x Character Encoding Problems

MySQL 3.x is known to have some problems upper- and lower-casing certain (non-ASCII) characters.

#### Diagnosing the problem

1. Follow the instructions for [Troubleshooting Character Encodings](#).
2. If the upper- and lower-cased strings displayed on the Encoding Test are different, then your database is probably affected.

An example (faulty) output of the Encoding Test is shown below:

*Screenshot: Encoding Test Output (excerpt)*

**Test 4: Database round-trip (select as upper-case)**

**This is the string from Test 2 after being stored in the database and then retrieved as upper-case**

**INtERNAtIOnAlIZæTIøn** ←

**Expected result (converting Java string to uppercase)**

**INtERNAtIOnAlIZæTIøn**

## Solution

Upgrade to a newer version of MySQL. (4.1 is confirmed to work.)

## Other Settings

- [Configuring a WebDAV client for Confluence](#)
- [Configuring HTTP Timeout Settings](#)
- [Configuring Number Formats](#)
- [Configuring Shortcut Links](#)
- [Configuring Time and Date Formats](#)
- [Enabling the Remote API](#)
- [Enabling Threaded Comments](#)
- [Installing a Language Pack](#)
- [Installing Patched Class Files](#)

## Configuring a WebDAV client for Confluence

WebDAV allows users to access Confluence content via a WebDAV client, such as 'My Network Places' in Microsoft Windows. Provided that the user has permission, they will be able to read and write to spaces, pages and attachments in Confluence. Users will be asked to log in and the standard Confluence content access permissions will apply to the equivalent content available through the WebDAV client.

Mapping a Confluence WebDAV network drive requires a set of specific criteria to be met. For specific information, please see [Windows Network Drive Requirements](#).

### Introduction to Confluence's WebDAV Client Integration

By default, all WebDAV clients have permission to write to Confluence. Write permissions include the ability for a WebDAV client to create, edit, move or delete content associated with spaces, pages and attachments in a Confluence installation.

On the '**WebDAV Configuration**' screen in the Confluence Administration Console, you can:

- [Deny a WebDAV client write permissions to a Confluence installation](#) using a regular expression (regex).
- [Disable or enable strict path checking](#).
- [Enable or disable access to specific virtual files/folders](#).

### Note:

- The 'WebDav Configuration' page is only available if the WebDAV plugin has been enabled. Note that this plugin is bundled with Confluence, and can be enabled or disabled by the System Administrator.
- The settings on the 'WebDav Configuration' page do **not** apply to external attachment storage configuration.

### Using a WebDAV Client to Work with Pages

The following sections tell you how to set up a WebDAV client natively for a range of different operating

### On this page:

- [Introduction to Confluence's WebDAV Client Integration](#)
- [Using a WebDAV Client to Work with Pages](#)
- [Restricting WebDAV Client Write Access to Confluence](#)
- [Disabling Strict Path Checking](#)
- [Virtual Files and Folders](#)
- [Known Issues](#)

### Related pages:

- [Global Permissions Overview](#)
- [Disabling and enabling add-ons](#)
- [Disabling or Enabling Confluence Add-ons](#)
- [Attachment Storage Configuration](#)

systems. WebDAV clients typically appear as drives in your operating system's file browser application, such as Windows Explorer in Microsoft Windows, or Konqueror in Linux.

### Accessing Confluence in Finder on Mac OSX

To use Finder to view and manage Confluence content:

1. In Finder choose **Go > Connect to Server**.
2. Enter your Confluence URL in this format:

```
http://<confluenceURL>/plugins/servlet/confluence/default
```

For example if your Confluence URL is <http://ourconfluence.atlassian.net/wiki> you would enter:

```
http://ourconfluence.atlassian.net/wiki/plugins/servlet/confluence/default
```

You must use HTTP, you'll be able to connect with HTTPS but you won't see any content.

3. Enter your **Confluence** username and password and click **Connect**.  
Note: use your username (jsmith), not your email address, unless your email address is your username.
4. Confluence will appear as a shared drive in Finder.

You can use the same URL to connect using a third party WebDAV client, like CyberDuck.

### Accessing Confluence in Explorer in Microsoft Windows

This section covers the two methods for configuring a WebDAV client natively in Microsoft Windows:

- As a network drive
- As a web folder

If possible, use the network drive method as this will enable more comprehensive WebDAV client interaction with Confluence than that provided by a web folder. However, your Confluence instance must meet several environmental constraints if you use this method. If you cannot configure your instance to meet these requirements, then use the web folder method or third-party WebDAV client software.

If you run into any problems with the procedures in this section, please refer to the [WebDAV Troubleshooting](#) page.

#### Windows Network Drive

To map a Confluence WebDAV client network drive, your Confluence instance must be configured so that *all* of the following criteria is met:

- Uses HTTP (not HTTPS)
- Listens on port 80 (not 8090, the default for Confluence distributions)
- Has no context root
- There is an issue ([WBDV-208](#)) that can prevent Network Drives from being mapped. Please use the Network Folders steps below as a workaround.

The reason for these restrictions results from limitations in Microsoft's Mini-Redirector component. For more information, please refer to Microsoft's [server discovery issue](#).

#### To map a Confluence WebDAV client network drive in Microsoft Windows:

1. In Windows XP, go to **My Computer -> Tools menu -> Map Network Drive**.  
In Windows Vista, go to **Computer -> Map Network Drive**.  
The 'Map Network Drive' dialog box opens.
2. Specify the following input to map the WebDAV client as a network drive:
  - **Drive:** <Any drive letter> (for example, Z:)
  - **Folder:** \\<hostname>\webdav (for example, \\localhost\webdav)

### 3. Click **'Finish'**.

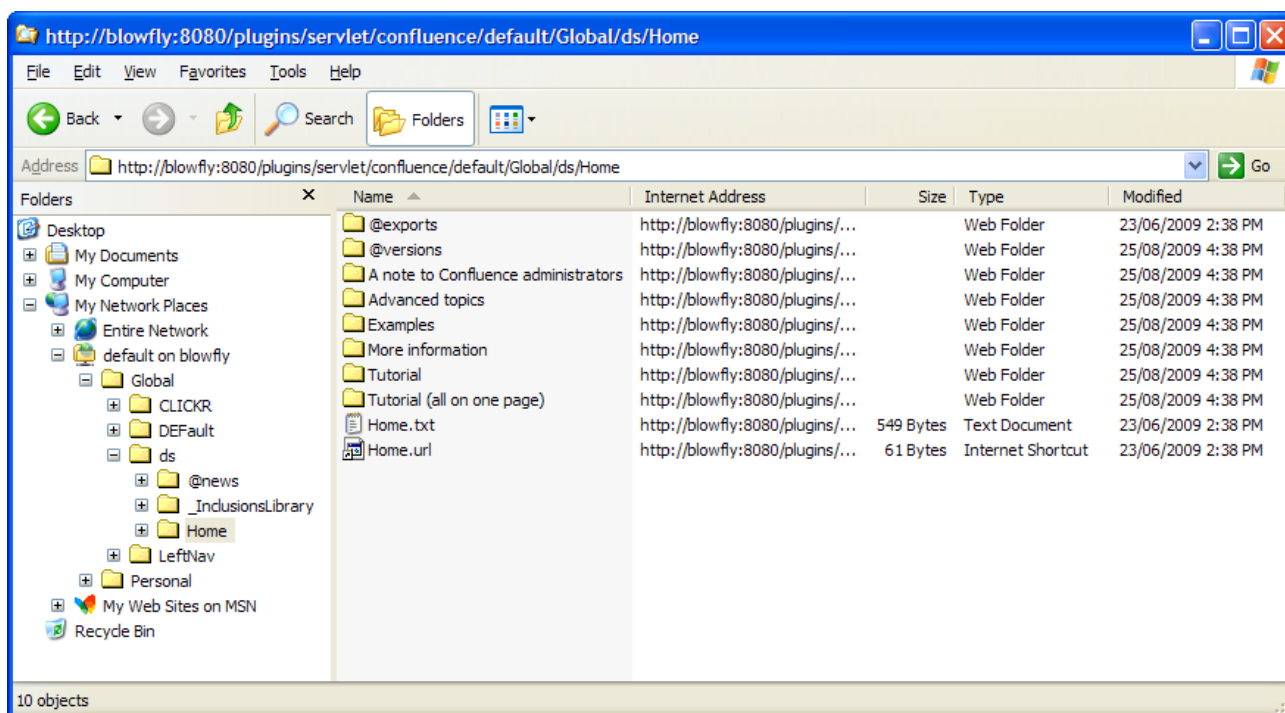
When prompted for login credentials, specify your Confluence username and password.

#### Windows Web Folder

#### To map a Confluence WebDAV client web folder in Windows XP:

1. Go to **My Network Places** and choose **'Add a network place'**. The 'Add Network Place Wizard' opens.
2. Click **'Next'**, ensure that **'Choose another network location'** is selected and then click **'Next'** again.
3. In the 'Internet or network address' field, enter the URL for the Confluence WebDAV location (for example, `http://<confluence server url>/confluence/plugins/servlet/confluence/default` or `http://<confluence server url>/plugins/servlet/confluence/default`) and then click **'Next'**.  
When prompted for login credentials, specify your Confluence username and password.
4. Provide a meaningful name for your web folder and proceed with the remainder of the wizard.
5. Click **'Finish'**.

*Screenshot: A Confluence WebDAV Client Web Folder in Windows XP*



#### To map a Confluence WebDAV client web folder in Windows Vista:

This procedure is very similar to the one for Windows XP. However, the following procedure includes the slight interface differences that are specific to Windows Vista.

1. Open the 'Map Network Drive' dialog box (refer to first step of the procedure above for mapping a network drive) and choose **'Connect to a Web site that you can use to store your documents and pictures'**. The 'Add Network Location' wizard opens.
2. Click **'Next'**, ensure that **'Choose a custom network location'** is selected and then click **'Next'** again.
3. In the 'Internet or network address' field, enter the URL for the Confluence WebDAV location (for example, `http://<confluence server url>/confluence/plugins/servlet/confluence/default` or `http://<confluence server url>/plugins/servlet/confluence/default`) and then click **'Next'**.  
When prompted for login credentials, specify your Confluence username and password.
4. Provide a meaningful name for your network location/web folder and proceed with the remainder of the wizard.
5. Click **'Finish'**.

#### Setting up a WebDAV client in Linux or Solaris

There are many tools and mechanisms available for configuring WebDAV clients in these operating systems. Therefore, we have chosen to demonstrate this using the file manager [Konqueror](#), which is part of the Linux

## K Desktop Environment.

### To set up a Confluence WebDAV client in Konqueror:

1. Open Konqueror.
2. In the 'Location' field, enter the URL for the Confluence WebDAV location using the 'protocol' webdav s (for example, webdavs://<confluence server url>/confluence/plugins/servlet/confluence/default or webdavs://<confluence server url>/plugins/servlet/confluence/default) and press Enter.  
If prompted for login credentials, specify your Confluence username and password.  
You should be able to click to load many, but not all files. In practice, you would normally save a modified file locally, then drag it to the Konqueror window to upload it to Confluence.

### Restricting WebDAV Client Write Access to Confluence

In earlier versions of the WebDAV plugin, separate options for restricting a WebDAV client's write permissions (that is, create/move, edit and delete actions), were available. However, in the current version of this plugin, they have been simplified and combined into a general write permission restriction that covers all of these actions.


WebDAV clients are now denied write permission to your Confluence installation by setting a regex that matches specific content within the WebDAV client's user agent header. Upon setting a regex, it will be added to a list of restricted WebDAV clients. Any WebDAV clients whose user agent header matches a regex in this list will be denied write permission to your Confluence installation.

*Example: A PROPFIND method header generated by a Microsoft Web Folder WebDAV client, showing the user agent header field:*

```
PROPFIND /plugins/servlet/confluence/default HTTP/1.1
Content-Language: en-us
Accept-Language: en-us
Content-Type: text/xml
Translate: f
Depth: 1
Content-Length: 489
User-Agent: Microsoft Data Access Internet Publishing Provider DAV
Host: 127.0.0.1:8082
Connection: Keep-Alive
```

**Note:** Unlike earlier versions of the WebDAV plugin which could only restrict write permissions for **all** WebDAV clients, the current version of this plugin allows you to restrict write permissions to specific WebDAV clients selectively.

### To restrict a WebDAV client's write access permissions to your Confluence installation:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **'WebDav Configuration'** in the left panel. The 'WebDAV Configuration' page is displayed.
3. Enter a regex that matches a specific component of the user agent header sent by the WebDAV client you want to restrict.
4. Click the **'Add new regex'** button. The regex is added to the list of restricted WebDAV clients.  
You can repeat steps 3 and 4 to add a regex for each additional WebDAV client you want to restrict.
5. Click the **'Save'** button to save the configuration changes.

### To restore one or more restricted WebDAV client's write access permissions to your Confluence installation:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration

2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Select the regex(es) from the list that match(es) the user agent header sent by the restricted WebDAV client(s) you want to restore.
4. Click the '**Remove selected regexes**' button. The regexes you had selected are removed from the list of restricted WebDAV clients.
5. Click the '**Save**' button to save the configuration changes.

Screenshot: WebDAV configuration

## WebDAV Configuration

### Denying Modifications From WebDAV Clients

You can deny certain WebDAV clients from writing to Confluence. To do that, please add a regular expression matching the clients [user agent](#) headers. For instance, "Microsoft.\*" will deny write operations by Microsoft WebDAV clients.

### Disable Strict Path Checking

Strict path checking is enabled by default to respect the file system heirarchy that the plugin exposes. Unfortunately, that can cause problems like [WBDV-175](#). The problem is related to corrupt page ancestor data and we recommend that users affected by the problem fix it by [rebuilding the ancestor table](#). If that does not fix the problem or if the approach can't be taken for any reason, users can disable the strict path checking here.



**Please note that disabling strict path checking may cause other problems with certain WebDAV clients.** This is a highly experimental workaround.

Disable strict path check : ☐

### Disabling Strict Path Checking

If you observe any idiosyncrasies with your WebDAV client, such as a folder that does exist on your Confluence site but is missing from the client, you can disable the WebDAV plugin's strict path checking option, which may minimise these problems.

#### To disable the WebDAV plugin's strict path checking option:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Clear the 'Disable strict path check' check box.  
 You can re-enable this option at a later point in time by simply selecting this check box.
4. Click the '**Save**' button to save this configuration change.

### Virtual Files and Folders

In the unlikely event that you observe any problems with the WebDAV client's performance or stability, you can enable access to automatically generated (that is, virtual) files and folders.

#### Note:

By default, these options are hidden on the 'WebDAV Configuration' page. To make them visible, you must

append the parameter `?hiddenOptionsEnabled=true` to the end of your URL and reload the page. For example:

```
<Confluence base
URL>/admin/plugins/webdav/config.action?hiddenOptionsEnabled=true
```

Screenshot: The Hidden Virtual Files and Folders Option

.url	@exports	@versions
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Save

#### To enable or disable access to virtual files and folders:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Amend your URL as described in the note above and reload the 'WebDav Configuration' page.
4. Select or clear the check box options in the 'Virtual Files and Folders' section as required.
5. Click the '**Save**' button to save the configuration changes.

#### Known Issues

Please refer to the [WebDAV plugin documentation](#) for a description of the known issues and suggested workarounds.

### Configuring HTTP Timeout Settings

When macros such as the [RSS Macro](#) make HTTP requests to servers which are down, a long timeout value is used. You can set this timeout value through a system parameter to avoid this.

#### To configure the HTTP Timeout Settings:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Select '**General Configuration**' under the 'Configuration' heading in the left-hand panel.
3. Find the '**Connection Timeouts**' section in the lower portion of the screen.
4. Click '**Edit**' to adjust the settings:
  - **Adjust External connections enabled:** This setting allows system administrators to disable external connections so macros like the [RSS Macro](#) won't be allowed to make connections to an external server. It provides protection against external servers providing insecure HTML, timing out or causing performance problems. The default setting is 'true'.
  - **Connection Timeout (milliseconds):** Sets the maximum time for a connection to be established. A value of zero means the timeout is not used. The default setting is ten seconds (10000).
  - **Socket Timeout (milliseconds):** Sets the default socket timeout ([SO\\_TIMEOUT](#)) in milliseconds, which is the maximum time Confluence will wait for data. A timeout value of zero is interpreted as an infinite timeout. The default setting is ten seconds (10000).


### Configuring Number Formats

There are two number format settings in Confluence:

- Long number format. For example: #####
- Decimal number format. For example: #####.#####

Confluence uses the guidelines in this Java document from Oracle: [Class NumberFormat](#).

### To change the number formats in Confluence:

1. Choose 
  - > **General Configuration**
2. Choose **Edit**
3. Update the **Long Number Format** and **Decimal Number Format** to suit your requirements
4. Choose **Save**

### Configuring Shortcut Links

Shortcut links provide a quick way of linking to resources that are frequently referenced from Confluence. When you create a shortcut link, you assign a key to an URL so that, when editing, a user can type just the key instead of the complete URL.

#### Example: Creating a shortcut to Google

Most Google searches look like this: <http://www.google.com/search?q=>. If you create a shortcut for this search with the key 'google', every time a user needs to use <http://www.google.com/search?q=searchterms>, they can just type [ **search terms**@google] instead.


Here is a screenshot showing the shortcuts currently defined on <http://confluence.atlassian.com>:

Key	Expanded Value	Default Alias	Operations
cache	<a href="http://www.google.com/search?q=cache:">http://www.google.com/search?q=cache:</a>		<a href="#">Remove</a>
imdb	<a href="http://us.imdb.com/Title?">http://us.imdb.com/Title?</a>		<a href="#">Remove</a>
jira	<a href="http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=">http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=</a>	JIRA Issue %s	<a href="#">Remove</a>
googlegroups	<a href="http://groups.google.com/groups?q=">http://groups.google.com/groups?q=</a>		<a href="#">Remove</a>
google	<a href="http://www.google.com/search?q=">http://www.google.com/search?q=</a>		<a href="#">Remove</a>
dictionary	<a href="http://www.dict.org/bin/Dict?Database=*&amp;Form=Dict1&amp;Strategy=*&amp;Query=">http://www.dict.org/bin/Dict?Database=*&amp;Form=Dict1&amp;Strategy=*&amp;Query=</a>		<a href="#">Remove</a>

Shortcut links are added and maintained by Confluence administrators from the **Administration Console**.

### Creating shortcut links

#### To create a shortcut link:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration
2. Choose **Shortcut Links** in the left-hand panel.
3. Enter a **Key** for your shortcut. This is the shortcut name a user will use to reference the URL.
4. Enter the **Expanded Value**. This is the URL for the link. You can use '%s' in the URL to specify where the user's input is inserted. If there is no '%s' in the URL, the user's input will be put at the end.
5. Enter a **Default Alias**. This is the text of the link which will be displayed on the page where the shortcut is used, with the user's text being substituted for '%s'.
6. Choose **Submit**.

### Using shortcut links

Enter a shortcut link on the **Advanced** tab of the Insert Link dialog. See [Links](#) for details.

Specify in the link what should be appended to the end of the shortcut URL, followed by an at-sign (@) and the key of the shortcut. Shortcut names are case-insensitive. So, for example, using the keys shown in the above screenshot:

To link to...	Type this	Resulting URL
a JIRA issue	CONF-1000@JIRA	<a href="http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=CONF-1000">http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=CONF-1000</a>
a Google search	Atlassian Confluence@Google	<a href="http://www.google.com/search?q=Atlassian+Confluence">http://www.google.com/search?q=Atlassian+Confluence</a>

### Deleting shortcut links

Shortcut links are listed on the **Shortcut Links** tab of the Administration Console. Click **Remove** to delete the shortcut.

## Configuring Time and Date Formats

You can localise the formats that Confluence uses to display dates and times within the web interface. The settings use the syntax of Java's `SimpleDateFormat` class, as described in this document: [Java SimpleDateFormat](#).

There are three time and date format settings:

- **Time format:** Used when displaying only the time of day. For example, when a blog post is published. Example of configuration: `h:mm a`
- **Date time format:** Used when displaying both the date and the time of day. For example, in historical versions of pages. Example of configuration: `MMM dd, yyyy HH:mm`
- **Date format:** Used when displaying only the date. For example, the creation and most recent modification dates of pages. Example of configuration: `MMM dd, yyyy`

### Related pages:

- [Choosing a Default Language](#)
- [Installing a Language Pack](#)
- [Confluence Administrator's Guide](#)

### To change the time and date formats:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Enter the values for **Time Format**, **Date Time Format** and **Date Format**, to suit your requirements.
5. Choose **Save**.

## Enabling the Remote API

Confluence provides XML-RPC and SOAP remote APIs (application programming interfaces). You need to enable the APIs from the **Administration Console** before you can access Confluence remotely.

You need [System Administrator](#) permissions in order to perform this function.

### Related pages:

- [Confluence Remote API Reference](#)

### To enable the remote API:

1. Choose the **cog icon**



- , then choose **General Configuration** under Confluence Administration
2. Click **Further Configuration** in the left-hand panel.
3. Click **Edit**.
4. Click the check box next to **Remote API (XML-RPC & SOAP)**.
5. Click **Save**.

## Enabling Threaded Comments

Comments on pages or blog posts are displayed in one of two views:

- **Threaded:** Shows the comments in a hierarchy of responses. Each reply to a comment is indented to indicate the relationships between the comments.
- **Flat:** Displays all the comments in one single list and does not indicate the relationships between comments.

By default, comments are displayed in **threaded** mode. A Confluence Administrator (see [Global Permissions Overview](#)) can enable or disable the threaded view for the entire Confluence site.

### Related pages:

- [Comment on pages and blog posts](#)
- [Confluence Administrator's Guide](#)

### To enable or disable the threaded view:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Select **Further Configuration** in the left-hand panel
3. Choose **Edit**
4. Select or unselect the **Threaded Comments** checkbox to enable or disable threaded mode
5. Choose **Save**

## Installing a Language Pack

Confluence ships with a number of bundled language packs. These languages appear as options on the 'Language Configuration' screen in the Administration Console when [choosing a default language](#) and as 'Language' options for users in their [user settings](#). You can make additional languages available for selection by installing language packs. Please note, you must be a Confluence administrator to install a language pack.

Language packs are plugins. The process of installing a language pack is the same as [installing a new plugin](#).

### Related pages:

- [Choosing a Default Language](#)
- [Configuring Indexing Language](#)
- [Installing add-ons](#)

### Installing a Language Pack using the Universal Plugin Manager

To install a language pack using the Universal Plugin Manager:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Find New Add-ons** in the left-hand panel.
3. Find the language pack on the Atlassian Marketplace.
4. Choose **Install** to install the language pack.

### Installing a Language Pack Manually

To install a language pack manually, you will need to upload the language pack plugin as described below. The language pack plugin will be enabled by default once you have installed it.

Plugins are distributed as JAR or OBR (OSGi Bundle Repository) files. To install a plugin:

1. Choose the **cog icon**



- , then choose **General Configuration** under Confluence Administration
2. Choose **Manage Add-ons**.
3. Choose **Upload Plugin**.
4. Choose **Browse** to find the plugin file you wish to install from your hard drive and select it, or enter a network location by URL.
5. Choose **Upload**.  
The plugin will be uploaded to Confluence and will be automatically installed.
6. Check the list of user-installed plugins to ensure that the add-on is available.
7. Enable the plugin if necessary. (Some plugins will be enabled by default when they are installed. Others will have to be manually enabled from the 'Manage Add-ons' page.)

### Finding more Language Packs

- You can download official language packs from the [Atlassian Marketplace](#). You can also download language packs developed by the Confluence user community from the [Language Pack Translations page](#).

### Showing User Interface Key Names for Translation

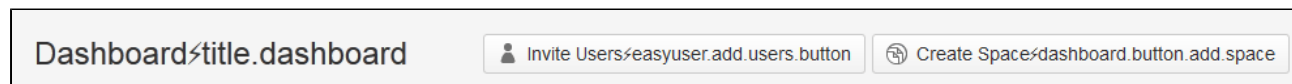
This feature is useful if you are working on creating translations of the Confluence user interface. After opening the Confluence dashboard, you can add this text to the end of your Confluence URL:

```
?i18ntranslate=on
```

Then press Enter.

This will cause each element of the user interface to display its special **key name**. This makes it easier to find the context for each key within the user interface. You can then search for the key on <http://translations.atlassian.com> where you can enter an appropriate translation for your custom language pack.

The key names are displayed with a 'lightning bolt' graphic. For example:



To turn off the translation view, add this code to the end of the Confluence URL:

```
?i18ntranslate=off
```

### Installing Patched Class Files

Atlassian support or the Atlassian bug-fixing team may occasionally provide patches for critical issues that have been resolved but have not yet made it into a release. Those patches will be class files which are attached to the relevant issue in our [JIRA bug-tracking system](#).

#### Installation Instructions for the Confluence Distribution

Follow these steps to install a patched class file:

1. Shut down your confluence instance.
2. Copy the supplied class files to <installation-directory>/confluence/WEB-INF/classes/<subdirectories>, where:
  - <installation-directory> must be replaced with your Confluence Installation directory. (If you need more information, read about the [Confluence Installation Directory](#).)
  - <subdirectories> must be replaced by the value specified in the relevant JIRA issue. This value will be different for different issues. In some cases, the subdirectories will not exist and you will need to create them before copying the class files. Some issues will contain the patch in the form of a ZIP file which will contain the desired directory structure.
3. Restart your Confluence instance for the changes to become effective.

**i** Class files in the `/WEB-INF/classes` directory of a web application will be loaded before classes located in JAR files in the `/WEB-INF/lib` directory. Therefore, classes in the first directory will effectively replace classes of the same name and package which would otherwise be loaded from the JAR files.

### Reverting the patch

To revert the patch, simply remove the class files from the `<installation-directory>/confluence/WEB-INF/classes/` folder (taking care to only remove those that apply to the patch you wish to revert), then restart the instance.

Once the issue that the patch relates to is resolved, you should upgrade to the version of Confluence that contains the fix, and revert the patch. Patches are often naive and untested and may not solve the problem in the most efficient way. As such, an official fix should be preferred in all cases.

## Configuring System Properties

This page describes how to set Java properties and options on startup for Confluence.

See [How to fix out of memory errors by increasing available memory](#) for specific instructions for OutOfMemory Errors.

### Linux

#### To configure System Properties in Linux installations:

1. From `<confluence-install>/bin` (Stand-alone) or `<Tomcat-home>/bin` (EAR-WAR installation), open **setenv.sh**.
2. Find the section **CATALINA\_OPTS=** (this is **JAVA\_OPTS=** in Confluence 5.5 and earlier)
3. Refer to the list of parameters [below](#).

**i** Add all parameters in a space-separated list, inside the quotations.

#### On this page:

- [Linux](#)
- [Windows \(starting from .bat file\)](#)
- [Windows Service](#)

#### Related pages:

- [Recognise d System Properties](#)
- [How to fix out of memory errors by increasing available memory](#)

### Windows (starting from .bat file)

#### To Configure System Properties in Windows Installations When Starting from the .bat File:

1. From `<confluence-install>/bin` (Stand-alone) or `<Tomcat-home>/bin` (EAR-WAR installation), open **setenv.bat**.
2. Find the section **set CATALINA\_OPTS=%CATALINA\_OPTS%** (this is **set JAVA\_OPTS=%JAVA\_OPTS%** in Confluence 5.5 and earlier)
3. Refer to the list of parameters [below](#).

**i** Add all parameters in a space-separated list. Make sure to keep the string `%CATALINA_OPTS%` in place.

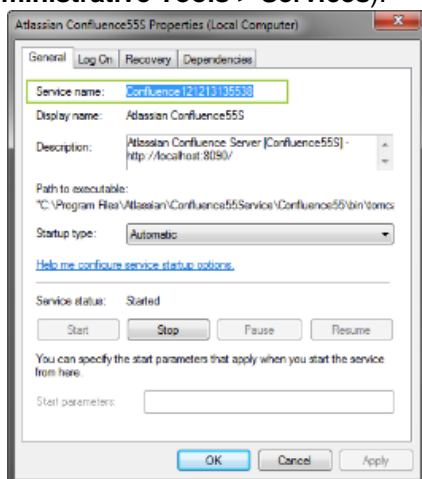
### Windows Service

There are two ways to configure system properties when you [Start Confluence Automatically on Windows as a Service](#), either via [command line](#) or [in the Windows Registry](#)

#### Setting Properties for Windows Services via Command Line

### To set properties for Windows Services via a command line:

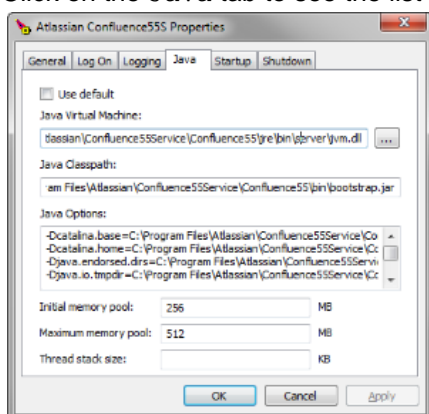
1. Identify the name of the service that Confluence is installed as in Windows (Go to **Control Panel > Administrative Tools > Services**):



- i** In the above example, the service name is *Confluence121213135538*.
2. Open the command window (Choose **Start > cmd.exe**)
  3. `cd` to the `bin` directory of your Confluence instance (or the `bin` directory of your Tomcat installation if you are running Confluence EAR/WAR).
  4. Run the following command:

```
tomcat7w //ES//<SERVICENAME>
```

- i** In the above example, it would be `tomcat7w //ES//Confluence121213135538`  
The Tomcat version number may be different if you are using an earlier version of Confluence.
5. Click on the **Java** tab to see the list of current start-up options:

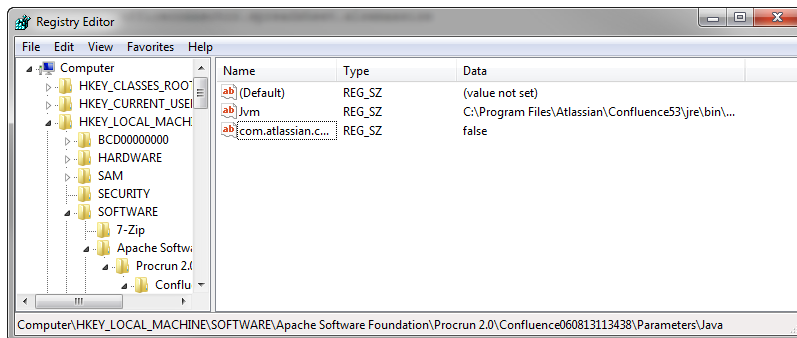


6. Append any new option on its own new line by adding to the end of the existing Java Options. Refer to the list of parameters [below](#).

### Setting Properties for Windows Services via the Windows Registry

In some versions of Windows, there is no option to add Java variables to the service. In these cases, you must add the properties by viewing the option list in the registry.

1. Go to the Registry Editor (**Start > regedit.exe**).
2. Find the Services entry:  
**32-bit:** HKEY\_LOCAL\_MACHINE >> SOFTWARE >> Apache Software Foundation >> Procrun 2.0 >> Confluence  
**64-bit:** HKEY\_LOCAL\_MACHINE >> SOFTWARE >> Wow6432Node >> Apache Software Foundation >> Procrun 2.0 >> Confluence service name



3. To change existing properties double-click the appropriate value.
4. To change additional properties, double-click options.
5. Refer to the list of parameters [below](#). Enter each on a separate line.

## Verifying Your Settings

To see what Confluence is using, check [Viewing System Properties](#).

## Recognised System Properties

Since	Default Value	Module...	Effect
<b>atlassian.forceSchemaUpdate</b>			
1.0	false	atlassian-config	By default, Confluence will only run its database schema update when it detects that it has been upgraded. This flag will force Confluence to perform the schema update on system startup.
<b>confluence.home</b>			
1.0	Any filesystem path	Confluence and atlassian-config	If this system property is set, Confluence will ignore the contents of the <code>confluence-init.properties</code> file, and use this property as the setting for the Confluence Home directory.
<b>confluence.dev.mode</b>			
1.0	false	Confluence	Enables additional debugging options that may be of use to Confluence developers (additionally it changes spring bean creation to use lazy initialization by default to decrease startup time). Do not enable this flag on a production system.
<b>confluence.disable.mailpolling</b>			
2.4	false	Confluence	If set to "true", will prevent Confluence from retrieving mail for archiving within spaces. Manually triggering "check for new mail" via the web UI will still work. This property has no effect on outgoing mail
<b>confluence.i18n.reloadbundles</b>			

1.0	true	Confluence	Setting this property will cause Confluence to reload its i18n resource bundles every time an internationalised string is looked up. This can be useful when testing translations, but will make Confluence run <i>insanely slowly</i> .
<b>confluence.ignore.debug.logging</b>			
1.0	true	Confluence	Confluence will normally log a severe error message if it detects that DEBUG level logging is enabled (as DEBUG logging generally causes a significant degradation in system performance). Setting this property will suppress the error message.
<b>confluence.jmx.disabled</b>			
3.0	false	Confluence	If set to "true", will disable Confluence's JMX monitoring. This has the same effect as setting the "enabled" property to false in WEB-INF/classes/jmxContext.xml
<b>confluence.optimize.index.module</b>			
2.2	20	Confluence	Number of index queue flushes before the index is optimised.
<b>confluence.plugins.bundled.disable</b>			
2.9	false	Confluence	Starts confluence without bundled plugins. May be useful in a development environment to make Confluence start quicker, but since bundled plugins are necessary for some of Confluence's core functionality, this property should not be set on a production system.
<b>atlassian.indexing.contentbody.maxsize</b>			
3.0		atlassian	Confluence stores the body of attachments in the index so the content of attachments are searchable. Set this property to set the maximum indexable size (in bytes). Limiting the size of attachment indexing decreases the size of the index when large attachments are present and prevents excerpts of large attachments being displayed in search results. It does not limit the size of attachments that can be uploaded.
<b>atlassian.mail.fetchdisabled</b>			
3.5	false	Confluence	Disables mail fetching services for IMAP and POP
<b>atlassian.mail.senddisabled</b>			

3.5	false	Confluence and atlassian-mail	Disables sending of mail
<b>atlassian.disable.caches</b>			
2.4	true	atlassian-plugins, atlassian-cache-servlet	Setting this property will disable conditional get and expires: headers on some web resources. This will significantly slow down the user experience, but is useful in development if you are frequently changing static resources and don't want to continually flush your browser cache.
<b>confluence.html.encode.automatic</b>			
2.9		Confluence	Setting this property forces the antixss encoding on or off, overriding the behaviour dictated by settings. The default behaviour differs between Confluence versions.
<b>org.osgi.framework.bootdelegation</b>			
2.10	empty	atlassian-plugins	Comma-separated list of package names to provide from application for OSGi plugins. Typically required when profiling Confluence. For example: "com.jprofiler., <b>com.yourkit.</b> ".
<b>confluence.diff.pool.size</b>			
3.1	20	Confluence	Maximum number of concurrent diffs. When that number is exceeded, additional attempts by RSS feeds to create diffs are ignored and logged. (The RSS requests succeed, they are just missing diffs).
<b>confluence.diff.timeout</b>			
3.1	1000	Confluence	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.
<b>confluence.html.diff.timeout</b>			
4.0	10000	Confluence	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.
<b>atlassian.user.experimentalMapping</b>			

2.10	false	Confluence	Setting this property changes the relationship between local users and local groups to reduce performance degradation when adding a local user to a local group with a large number of users. Please note, setting this property can slow down other user management functions. We recommend that you set it only if you are experiencing performance problems when adding local users to large local groups. Please refer to <a href="#">CONF-12319</a> , fixed in Confluence 3.1.1.
<b>confluence.import.use-experimental-importer</b>			
3.2	false	Confluence	Setting this property changes Confluence to use the Experimental XML Importer. It is designed to be a more stable implementation but, at the time of the release of 3.2, the importer is largely untested and thus not supported.
<b>atlassian.webresource.disable.minification</b>			
3.3	false	atlassian-plugins	Disables automatic minification of JavaScript and CSS resources served by Confluence.
<b>index.queue.thread.count</b>			
3.3	See "Effect"	Confluence	<p>Sets the number of threads to be used for the reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. There is no default value, i.e.</p> <ul style="list-style-type: none"> <li>• If you don't set <code>index.queue.thread.count</code>, the number of threads to be used are calculated based on the number of objects that need to be reindexed and the number of processors available (a maximum of 50 threads will be used).</li> <li>• If you set <code>index.queue.thread.count=2</code>, then two threads will be used to reindex the content (regardless of the number of objects to be reindexed or the number of processors available)</li> <li>• If you set <code>index.queue.thread.count=200</code>, then ten threads (the maximum allowed) will be used to reindex the content.</li> </ul> <p>Note: For Confluence versions from 3.3 to 5.6 the maximum threads is 10.</p>
<b>index.queue.batch.size</b>			
3.3	1500	Confluence	Size of batches used by the indexer. Reducing this value will reduce the load that the indexer puts on the system, but indexing takes longer. Increasing this value will cause indexing to be completed faster, but puts a higher load on the system. Normally this setting does not need tuning.

<b>password.confirmation.disabled</b>			
3.4	false	Confluence	This property disables the password confirmation functionality that Confluence uses as an additional security measure. With this property set, Confluence will <i>not</i> require password confirmation for the following actions: <a href="#">administrative actions</a> , change of email address and <a href="#">Captcha for failed logins</a> . Disabling password confirmations is useful if you are using a custom authenticator.
<b>confluence.browser.language.enabled</b>			
3.5	true	Confluence	Setting this property to "false" disables the detection of browser language headers, effectively restoring Confluence behaviour to that of earlier releases. Setting this property to "true" enables the detection of the language headers sent by the browser. Confluence will change the UI language based on the browser headers. See documentation on how users can <a href="#">choose a language preference</a> .
<b>upm.pac.disable</b>			
Universal Plugin Manager 1.5	false	Universal Plugin Manager (UPM)	When this property is set to true, then UPM will not try to access the <a href="#">The Atlassian Marketplace</a> . This is useful for application servers that do not have access to the Internet. See the <a href="#">UPM documentation</a> .
<b>confluence.reindex.documents.to.pop</b>			
3.5.9	20	Confluence	Indicates how many objects each indexing thread should process at a time during a full re-index. Please note that this number does not include attachments
<b>confluence.reindex.attachments.to.pop</b>			
3.5.9	10	Confluence	Indicates how many attachments each indexing thread should process at a time during a full re-index.
<b>confluence.upgrade.active.directory</b>			
3.5.11	false	Confluence	Forces Confluence to treat any LDAP directories it migrates as Active Directory, rather than relying on looking for sAMAccountName in the username attribute. This is necessary if you are upgrading from before Confluence 3.5, and need to use an attribute other than sAMAccountName to identify your users and are seeing LDAP: error code 4 - Sizelimit Exceeded exceptions in your logs. For more details, see <a href="#">Unable to Log In with Confluence 3.5 or Later Due to 'LDAP error code 4 - Sizelimit Exceeded'</a>

<b>confluence.context.batching.disable</b>			
4.0	false	Confluence	Disables batching for web resources in contexts (e.g. editor, main, admin). Useful for diagnosing the source of javascript or CSS errors.
<b>com.atlassian.logout.disable.session.invalidation</b>			
4.0	false	Confluence	Disables the session invalidation on log out. As of 4.0 the default behaviour is to invalidate the JSession assigned to a client when they log out. If this is set to true the session is kept active (but the user logged out). This may be valuable when using external authentication systems, but should generally not be needed.
<b>officeconnector.spreadsheet.xlsxmaxsize</b>			
4.0.5	2097152	Office Connector	Indicates the maximum size in bytes of an Excel file that can be viewed using the <code>viewxls</code> macro. If empty, the maximum size defaults to 2Mb. See <a href="#">CO NF-21043</a> for more details.
<b>com.atlassian.confluence.extra.calendar3.display.events.calendar.maxpercalendar</b>			
	200	Team Calendars	Specifies the maximum number of events per calendar. This property is effective only if the <a href="#">Team Calendars</a> plugin is installed on your Confluence site.
<b>com.atlassian.confluence.allow.downgrade</b>			
4.3.2, 5.0-OD-10	false	Confluence	Allows Confluence to start up against the home directory of a newer version of Confluence. Note that running Confluence like that is unsupported. You should only turn this on if you know what you are doing. See <a href="#">After Downgrading, Confluence Will No Longer Run</a> for details.
<b>reindex.thread.count</b>			
5.2		Confluence	Sets the number of threads to be used for a one-off reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. There is no default value. This system property does not affect the incremental indexing that Confluence does.
<b>reindex.attachments.thread.count</b>			
5.2	4	Confluence	Sets the number of threads to be used for reindexing attachments specifically, and allows you to reduce the concurrency for these more memory intensive index items.

<b>confluence.mbox.directory</b>			
5.4.1		Confluence	Setting this property defines the directory on your Confluence Server where mailboxes can be imported from (for use with the Confluence Mail Archiving add-on). Mailboxes are not able to be imported from any other location. We recommend administrators create a directory in the Confluence Home directory specifically for this purpose. Mail cannot be imported from the server until this system property is set.
<b>confluence.search.max.results</b>			
5.5	1000	Confluence	Setting this property changes the maximum number of items Confluence Search will return.
<b>confluence.upgrade.recovery.file.enabled</b>			
5.5	true	Confluence	By default, Confluence creates an upgrade recovery file before and after an upgrade. The operation can take long time on large databases and can be safely turned off if there is a process to back up database and verify the backup before performing an upgrade. Setting this property to false will disable upgrade recovery file creation.
<b>confluence.junit.report.directory</b>			
5.5		Confluence	Setting this property defines the directory on your Confluence Server where JUnit Reports can be imported from (for use in the JUnit Report Macro). No other locations are permitted. We recommend administrators create a directory in the Confluence Home directory specifically for this purpose. JUnit Test result files cannot be imported from the server until this system property is set.
<b>officeconnector.textextract.word.docxmaxsize</b>			
5.5.3	16777216	Confluence	This property defines the maximum amount of memory that can be allocated to extracting text from a DOCX file (for the purposes of search etc). The value refers to the amount of text inside the file (in bytes), ignoring any images.
<b>cluster.login.rememberme.enabled</b>			
5.6	False		<p>In a cluster, setting this property to <code>True</code> will enable the 'Remember Me' checkbox on the login page. This is not recommended in a cluster, and is disabled by default (i.e. 'Remember me' is always on and users can move seamlessly between nodes).</p> <p>This system property has no effect in standalone Confluence.</p>

<b>confluence.cluster.hazelcast.listenPort</b>			
5.6	5801		In a cluster, this property can be used to override the default port that Hazelcast will bind to, for example, if the port is unavailable, or you need to run more than one node on the same host (not recommended). It defaults to 5801.
<b>confluence.document.conversion.threads</b>			
5.7		Confluence	The number of threads allocated to the file conversion service is calculated dynamically based on the amount of memory assigned to the instance and the number of CPU cores (usually 4 to 6 threads). This property can be used to change the number of threads. Decrease threads to resolve OOME issues, increase threads to resolve problems with documents spending too long in the queue.
<b>confluence.document.conversion.threads.wait</b>			
5.7	1000	Confluence	Set this property to change the maximum number of items that can be queued for conversion. Any file conversion requests that are made when this maximum limit has been reached are aborted.
<b>confluence.cluster.node.name</b>			
5.7		Confluence	Set this property to give each node in your Data Center cluster a human readable name (displayed in email notifications and in the footer). If left unset, Confluence will assign a node identifier to each node.
<b>confluence.document.conversion.fontpath</b>			
5.8.7		Confluence	<p>Set this property to define a directory where you can add additional fonts to be used when rendering files (in previews and thumbnails).</p> <p>This is useful if you need to support previewing files with specific fonts, or fonts with multibyte characters (such as Japanese).</p>
<b>confluence.document.conversion.words.defaultfontname</b>			
5.8.7		Confluence	<p>Set this property to change the default font for rendering Word (.doc and .docx) files in Confluence.</p> <p>Specify just the name of the font (not the path).</p>
<b>confluence.document.conversion.slides.defaultfontname.regular</b>			

5.8.7		Confluence	<p>Set this property to change the default font for rendering regular fonts in Powerpoint (.ppt and .pptx) files in Confluence.</p> <p>Specify just the name of the font (not the path).</p>
<b>confluence.document.conversion.slides.defaultfontname.asian</b>			
5.8.7	TakaoPGothic	Confluence	<p>Set this property to change the default font for rendering asian fonts in Powerpoint (.ppt and .pptx) files in Confluence.</p> <p>Specify just the name of the font (not the path).</p>
<b>confluence.document.conversion.slides.defaultfontname.symbol</b>			
5.8.7		Confluence	<p>Set this property to change the default font for rendering symbols in Powerpoint (.ppt and .pptx) files in Confluence.</p> <p>This is the font that will be used for bullets and other symbols when the font <code>Symbol</code> is not found.</p> <p>Specify just the name of the font (not the path).</p>

## Recognised System Properties

Confluence supports some configuration and debugging settings that can be enabled through Java system properties. System properties are usually set by passing the `-D` flag to the Java virtual machine in which Confluence is running. See the full instructions: [Configuring System Properties](#).

Since	Default Value	Module...	Effect
<b>atlassian.forceSchemaUpdate</b>			
1.0	false	atlassian-config	By default, Confluence will only run its database schema update when it detects that it has been upgraded. This flag will force Confluence to perform the schema update on system startup.
<b>confluence.home</b>			
1.0	Any filesystem path	Confluence and atlassian-config	If this system property is set, Confluence will ignore the contents of the <code>confluence-init.properties</code> file, and use this property as the setting for the Confluence Home directory.
<b>confluence.dev.mode</b>			
1.0	false	Confluence	Enables additional debugging options that may be of use to Confluence developers (additionally it changes spring bean creation to use lazy initialization by default to decrease startup time). Do not enable this flag on a production system.
<b>confluence.disable.mailpolling</b>			

2.4	false	Confluence	If set to "true", will prevent Confluence from retrieving mail for archiving within spaces. Manually triggering "check for new mail" via the web UI will still work. This property has no effect on outgoing mail
<b>confluence.i18n.reloadbundles</b>			
1.0	true	Confluence	Setting this property will cause Confluence to reload its i18n resource bundles every time an internationalised string is looked up. This can be useful when testing translations, but will make Confluence run <i>insanely slowly</i> .
<b>confluence.ignore.debug.logging</b>			
1.0	true	Confluence	Confluence will normally log a severe error message if it detects that DEBUG level logging is enabled (as DEBUG logging generally causes a significant degradation in system performance). Setting this property will suppress the error message.
<b>confluence.jmx.disabled</b>			
3.0	false	Confluence	If set to "true", will disable Confluence's JMX monitoring. This has the same effect as setting the "enabled" property to false in <code>WEB-INF/classes/jmxContext.xml</code>
<b>confluence.optimize.index.module</b>			
2.2	20	Confluence	Number of index queue flushes before the index is optimised.
<b>confluence.plugins.bundled.disable</b>			
2.9	false	Confluence	Starts confluence without bundled plugins. May be useful in a development environment to make Confluence start quicker, but since bundled plugins are necessary for some of Confluence's core functionality, this property should not be set on a production system.
<b>atlassian.indexing.contentbody.maxsize</b>			
3.0		atlassian	Confluence stores the body of attachments in the index so the content of attachments are searchable. Set this property to set the maximum indexable size (in bytes). Limiting the size of attachment indexing decreases the size of the index when large attachments are present and prevents excerpts of large attachments being displayed in search results. It does not limit the size of attachments that can be uploaded.
<b>atlassian.mail.fetchdisabled</b>			

3.5	false	Confluence	Disables mail fetching services for IMAP and POP
<b><i>atlassian.mail.senddisabled</i></b>			
3.5	false	Confluence and atlassian-mail	Disables sending of mail
<b><i>atlassian.disable.caches</i></b>			
2.4	true	atlassian-plugins, atlassian-cache-servlet	Setting this property will disable conditional get and expires: headers on some web resources. This will significantly slow down the user experience, but is useful in development if you are frequently changing static resources and don't want to continually flush your browser cache.
<b><i>confluence.html.encode.automatic</i></b>			
2.9		Confluence	Setting this property forces the antixss encoding on or off, overriding the behaviour dictated by settings. The default behaviour differs between Confluence versions.
<b><i>org.osgi.framework.bootdelegation</i></b>			
2.10	empty	atlassian-plugins	Comma-separated list of package names to provide from application for OSGi plugins. Typically required when profiling Confluence. For example: "com.jprofiler., <b>com.yourkit.</b> ".
<b><i>confluence.diff.pool.size</i></b>			
3.1	20	Confluence	Maximum number of concurrent diffs. When that number is exceeded, additional attempts by RSS feeds to create diffs are ignored and logged. (The RSS requests succeed, they are just missing diffs).
<b><i>confluence.diff.timeout</i></b>			
3.1	1000	Confluence	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.
<b><i>confluence.html.diff.timeout</i></b>			
4.0	10000	Confluence	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.
<b><i>atlassian.user.experimentalMapping</i></b>			

2.10	false	Confluence	Setting this property changes the relationship between local users and local groups to reduce performance degradation when adding a local user to a local group with a large number of users. Please note, setting this property can slow down other user management functions. We recommend that you set it only if you are experiencing performance problems when adding local users to large local groups. Please refer to <a href="#">CONF-12319</a> , fixed in Confluence 3.1.1.
<b><i>confluence.import.use-experimental-importer</i></b>			
3.2	false	Confluence	Setting this property changes Confluence to use the Experimental XML Importer. It is designed to be a more stable implementation but, at the time of the release of 3.2, the importer is largely untested and thus not supported.
<b><i>atlassian.webresource.disable.minification</i></b>			
3.3	false	atlassian-plugins	Disables automatic minification of JavaScript and CSS resources served by Confluence.
<b><i>index.queue.thread.count</i></b>			
3.3	See "Effect"	Confluence	<p>Sets the number of threads to be used for the reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. There is no default value, i.e.</p> <ul style="list-style-type: none"> <li>• If you don't set <code>index.queue.thread.count</code>, the number of threads to be used are calculated based on the number of objects that need to be reindexed and the number of processors available (a maximum of 50 threads will be used).</li> <li>• If you set <code>index.queue.thread.count=2</code>, then two threads will be used to reindex the content (regardless of the number of objects to be reindexed or the number of processors available)</li> <li>• If you set <code>index.queue.thread.count=200</code>, then ten threads (the maximum allowed) will be used to reindex the content.</li> </ul> <p>Note: For Confluence versions from 3.3 to 5.6 the maximum threads is 10.</p>
<b><i>index.queue.batch.size</i></b>			
3.3	1500	Confluence	Size of batches used by the indexer. Reducing this value will reduce the load that the indexer puts on the system, but indexing takes longer. Increasing this value will cause indexing to be completed faster, but puts a higher load on the system. Normally this setting does not need tuning.
<b><i>password.confirmation.disabled</i></b>			

3.4	false	Confluence	This property disables the password confirmation functionality that Confluence uses as an additional security measure. With this property set, Confluence will <i>not</i> require password confirmation for the following actions: <a href="#">administrative actions</a> , change of email address and <a href="#">Captcha for failed logins</a> . Disabling password confirmations is useful if you are using a custom authenticator.
<b>confluence.browser.language.enabled</b>			
3.5	true	Confluence	Setting this property to "false" disables the detection of browser language headers, effectively restoring Confluence behaviour to that of earlier releases. Setting this property to "true" enables the detection of the language headers sent by the browser. Confluence will change the UI language based on the browser headers. See documentation on how users can <a href="#">choose a language preference</a> .
<b>upm.pac.disable</b>			
Universal Plugin Manager 1.5	false	Universal Plugin Manager (UPM)	When this property is set to true, then UPM will not try to access the <a href="#">The Atlassian Marketplace</a> . This is useful for application servers that do not have access to the Internet. See the <a href="#">UPM documentation</a> .
<b>confluence.reindex.documents.to.pop</b>			
3.5.9	20	Confluence	Indicates how many objects each indexing thread should process at a time during a full re-index. Please note that this number does not include attachments
<b>confluence.reindex.attachments.to.pop</b>			
3.5.9	10	Confluence	Indicates how many attachments each indexing thread should process at a time during a full re-index.
<b>confluence.upgrade.active.directory</b>			
3.5.11	false	Confluence	Forces Confluence to treat any LDAP directories it migrates as Active Directory, rather than relying on looking for sAMAccountName in the username attribute. This is necessary if you are upgrading from before Confluence 3.5, and need to use an attribute other than sAMAccountName to identify your users and are seeing LDAP: error code 4 - Sizelimit Exceeded exceptions in your logs. For more details, see <a href="#">Unable to Log In with Confluence 3.5 or Later Due to 'LDAP error code 4 - Sizelimit Exceeded'</a>
<b>confluence.context.batching.disable</b>			
4.0	false	Confluence	Disables batching for web resources in contexts (e.g. editor, main, admin). Useful for diagnosing the source of javascript or CSS errors.

<b><i>com.atlassian.logout.disable.session.invalidation</i></b>			
4.0	false	Confluence	Disables the session invalidation on log out. As of 4.0 the default behaviour is to invalidate the JSession assigned to a client when they log out. If this is set to true the session is kept active (but the user logged out). This may be valuable when using external authentication systems, but should generally not be needed.
<b><i>officeconnector.spreadsheet.xlsxmaxsize</i></b>			
4.0.5	2097152	Office Connector	Indicates the maximum size in bytes of an Excel file that can be viewed using the <code>viewxls</code> macro. If empty, the maximum size defaults to 2Mb. See <a href="#">CONF-21043</a> for more details.
<b><i>com.atlassian.confluence.extra.calendar3.display.events.calendar.maxpercalendar</i></b>			
	200	Team Calendars	Specifies the maximum number of events per calendar. This property is effective only if the <a href="#">Team Calendars</a> plugin is installed on your Confluence site.
<b><i>com.atlassian.confluence.allow.downgrade</i></b>			
4.3.2, 5.0-OD-10	false	Confluence	Allows Confluence to start up against the home directory of a newer version of Confluence. Note that running Confluence like that is unsupported. You should only turn this on if you know what you are doing. See <a href="#">After Downgrading, Confluence Will No Longer Run</a> for details.
<b><i>reindex.thread.count</i></b>			
5.2		Confluence	Sets the number of threads to be used for a one-off reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. There is no default value. This system property does not affect the incremental indexing that Confluence does.
<b><i>reindex.attachments.thread.count</i></b>			
5.2	4	Confluence	Sets the number of threads to be used for reindexing attachments specifically, and allows you to reduce the concurrency for these more memory intensive index items.
<b><i>confluence.mbox.directory</i></b>			

5.4.1		Confluence	Setting this property defines the directory on your Confluence Server where mailboxes can be imported from (for use with the Confluence Mail Archiving add-on). Mailboxes are not able to be imported from any other location. We recommend administrators create a directory in the Confluence Home directory specifically for this purpose. Mail cannot be imported from the server until this system property is set.
<b><i>confluence.search.max.results</i></b>			
5.5	1000	Confluence	Setting this property changes the maximum number of items Confluence Search will return.
<b><i>confluence.upgrade.recovery.file.enabled</i></b>			
5.5	true	Confluence	By default, Confluence creates an upgrade recovery file before and after an upgrade. The operation can take long time on large databases and can be safely turned off if there is a process to back up database and verify the backup before performing an upgrade. Setting this property to false will disable upgrade recovery file creation.
<b><i>confluence.junit.report.directory</i></b>			
5.5		Confluence	Setting this property defines the directory on your Confluence Server where JUnit Reports can be imported from (for use in the JUnit Report Macro). No other locations are permitted. We recommend administrators create a directory in the Confluence Home directory specifically for this purpose. JUnit Test result files cannot be imported from the server until this system property is set.
<b><i>officeconnector.textextract.word.docxmaxsize</i></b>			
5.5.3	16777216	Confluence	This property defines the maximum amount of memory that can be allocated to extracting text from a DOCX file (for the purposes of search etc). The value refers to the amount of text inside the file (in bytes), ignoring any images.
<b><i>cluster.login.rememberme.enabled</i></b>			
5.6	False		<p>In a cluster, setting this property to <code>True</code> will enable the 'Remember Me' checkbox on the login page. This is not recommended in a cluster, and is disabled by default (i.e. 'Remember me' is always on and users can move seamlessly between nodes).</p> <p>This system property has no effect in standalone Confluence.</p>
<b><i>confluence.cluster.hazelcast.listenPort</i></b>			

5.6	5801		In a cluster, this property can be used to override the default port that Hazelcast will bind to, for example, if the port is unavailable, or you need to run more than one node on the same host (not recommended). It defaults to 5801.
<b>confluence.document.conversion.threads</b>			
5.7		Confluence	The number of threads allocated to the file conversion service is calculated dynamically based on the amount of memory assigned to the instance and the number of CPU cores (usually 4 to 6 threads). This property can be used to change the number of threads. Decrease threads to resolve OOME issues, increase threads to resolve problems with documents spending too long in the queue.
<b>confluence.document.conversion.threads.wait</b>			
5.7	1000	Confluence	Set this property to change the maximum number of items that can be queued for conversion. Any file conversion requests that are made when this maximum limit has been reached are aborted.
<b>confluence.cluster.node.name</b>			
5.7		Confluence	Set this property to give each node in your Data Center cluster a human readable name (displayed in email notifications and in the footer). If left unset, Confluence will assign a node identifier to each node.
<b>confluence.document.conversion.fontpath</b>			
5.8.7		Confluence	<p>Set this property to define a directory where you can add additional fonts to be used when rendering files (in previews and thumbnails).</p> <p>This is useful if you need to support previewing files with specific fonts, or fonts with multibyte characters (such as Japanese).</p>
<b>confluence.document.conversion.words.defaultfontname</b>			
5.8.7		Confluence	<p>Set this property to change the default font for rendering Word (.doc and .docx) files in Confluence.</p> <p>Specify just the name of the font (not the path).</p>
<b>confluence.document.conversion.slides.defaultfontname.regular</b>			
5.8.7		Confluence	<p>Set this property to change the default font for rendering regular fonts in Powerpoint (.ppt and .pptx) files in Confluence.</p> <p>Specify just the name of the font (not the path).</p>

<b>confluence.document.conversion.slides.defaultfontname.asian</b>			
5.8.7	TakaoPGothic	Confluence	<p>Set this property to change the default font for rendering asian fonts in Powerpoint (.ppt and .pptx) files in Confluence.</p> <p>Specify just the name of the font (not the path).</p>
<b>confluence.document.conversion.slides.defaultfontname.symbol</b>			
5.8.7		Confluence	<p>Set this property to change the default font for rendering symbols in Powerpoint (.ppt and .pptx) files in Confluence.</p> <p>This is the font that will be used for bullets and other symbols when the font Symbol is not found.</p> <p>Specify just the name of the font (not the path).</p>

## Working with Confluence Logs

Confluence uses Apache's [log4j](#) logging service. This allows a developer or administrator to control the logging behavior and the log output file by editing a configuration file, without touching the application binary. There are six known [log4j logging levels](#).

If you request help from Atlassian Support, we will almost always ask for the `atlassian-confluence.log` from the `confluence-home/logs` directory. To access the logs, go to go to



> **General Configuration > Atlassian Support Tools.**

If you cannot access logs via Support Tools, check the properties file at `<confluence-installation>/confluence/WEB-INF/classes/confluence-init.properties`, look for the `confluence.home` setting in that file, then find the logs in the Confluence home directory.

### Finding the Confluence Log Files

This section describes Confluence's default logging behaviour, assuming that you have not changed the destination of the logs. In order to unify logging across different application servers, Confluence uses the `atlassian-confluence.log` as its primary log, not the application server log.

Both the Confluence and Confluence EAR/WAR distributions follow the same default behaviour:

- When you start Confluence, log entries will be sent to the application server logs until Confluence has completed its initial bootstrap. Any log entries written to the console will be repeated into the log in the Confluence home directory as described below.
- Once the initial startup sequence is complete, all logging will be to `<confluence-home>/logs/atlassian-confluence.log`. For example: `c:/confluence/data/logs/atlassian-confluence.log`.

Note that the default location is the Confluence **home directory**, not the application server's log file. The home directory is specified in `<confluence-installation>/confluence/WEB-INF/classes/confluence-init.properties`.

#### On this page:

- [Finding the Confluence Log Files](#)
- [Finding the Log Configuration File](#)
- [Changing the Destination of the Log Files](#)
- [Changing the Logging Levels](#)
- [Using Some Specific Confluence Logging Options](#)
- [Scanning Log Files for Known Problems](#)
- [Notes](#)

#### Related pages:

- [Enabling Detailed SQL Logging](#)
- [Enabling user access logging](#)
- [Generating a Thread Dump](#)
- [Enabling Page Request Profiling](#)

## Finding the Log Configuration File

Confluence's logging behaviour is defined in the following properties file:

`<CONFLUENCE-INSTALL>/confluence/WEB-INF/classes/log4j.properties`

This file is a standard log4j configuration file, as described in the [Apache log4j documentation](#).

## Changing the Destination of the Log Files

**Terminology:** In log4j, an output destination is called an 'appender'.

To change the destination of the log files, you need to stop Confluence and then change the settings in the '**Logging Location and Appender**' section of the `log4j.properties` file. The location of this file is described [above](#).

In the standard properties file, you will find entries for two appenders:

- `com.atlassian.confluence.logging.ConfluenceHomeLogAppender` – This is a custom appender which controls the default logging destination described [above](#). This appender allows the following settings:
  - `MaxFileSize`
  - `MaxBackupIndex`
- `org.apache.log4j.RollingFileAppender` – If you want to log to a different location, uncomment the `RollingFileAppender` line and change the destination file in the line below it. Comment out the previous lines referring to the `ConfluenceHomeLogAppender`.

Confluence ships with the full suite of appenders offered by log4j. Read more about appenders in the [log4j documentation](#).

## Changing the Logging Levels

See [Configuring Logging](#) for instructions on how to change the logging configuration of Confluence.

## Using Some Specific Confluence Logging Options

This section contains some pointers to specific log configurations you may need.

### ***Log the Details of SQL Requests made to the Database***

You may want to increase Confluence's logging so that it records individual SQL requests sent to the database. This is useful for troubleshooting specific problems.

You can enable detailed SQL logging in two ways:

- At runtime – see [instructions above](#).
- Via the logging properties file – see the [detailed instructions](#).

### ***Log the Details of Users Viewing/Accessing each Confluence Page***

You can configure the log to show which users are accessing which pages in Confluence. This can only be done via the logging properties file – see the [detailed instructions](#).

## Scanning Log Files for Known Problems

Confluence provides an inbuilt log scanner that will check your Confluence logs for errors and attempt to match them against known issues in our knowledge base and bug tracker. See [Troubleshooting Problems and Requesting Technical Support](#).

## Notes

- **Finding the thread dumps.** Thread dumps are logged to the application server log file.

## Configuring Logging

We recommend that you configure Confluence's logging to your own requirements. You can change the log settings in two ways:

- Configure logging in Confluence Administration – Your changes will be in effect only until you next restart Confluence.
- Edit the properties file – Your changes will take effect next time you start Confluence, and for all subsequent sessions.

Both methods are described below. In some rare circumstances you may *also* need to configure the [logging.properties](#) file.

**Terminology:** In log4j, a 'logger' is a named entity. Logger names are case-sensitive and they follow a hierarchical naming standard. For example, the logger named `com.foo` is a parent of the logger named `com.foo.Bar`.

### Configure logging in Confluence Administration

You can change some of Confluence's logging behaviour via the **Administration Console** while Confluence is running. Any changes made in this way will apply only to the currently-running Confluence lifetime. The changes are not written to the `log4j.properties` file and are therefore discarded when you next stop Confluence.

Not all logging behaviour can be changed via the Administration Console. For logging configuration not mentioned below, you will need to stop Confluence and then [edit the logging properties file](#) instead.

The '**Logging and Profiling**' screen shows a list of all currently defined loggers. On this screen you can:

- Turn [page profiling](#) on or off.
- Turn detailed SQL logging on or off.
- Add a new logger for a class/package name.
- Remove a logger for a class/package name.
- Set the logging level (INFO, WARN, FATAL, ERROR or DEBUG) for each class or package name.
- Reset all logging levels to a predefined profile.

Changing the logging configuration

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Select '**Logging and Profiling**' in the '**Administration**' section of the left-hand panel.
  - i** You need to have [System Administrator](#) permissions in order to perform this function.
3. The '**Logging and Profiling**' screen appears, as shown below. Use the following guidelines to change the logging behaviour while Confluence is running:
  - '**Performance Profiling**' — See [Page Request Profiling](#).
  - '**SQL Logging**' — Click the '**Enable SQL Logging**' button to log the details of SQL requests made to the database.
    - i** If you need to enable logging of SQL parameter values, you will need to change the setting in the [properties](#) file. This option is not available via the Administration Console.
  - '**Log4j Logging**' — Click one of the profile buttons to reset all your loggers to the predefined profiles:
    - The '**Production**' profile is a fairly standard profile, recommended for normal production conditions.
    - The '**Diagnostic**' profile gives more information, useful for troubleshooting and debugging. It results in slower performance and fills the log files more quickly.
  - '**Add New Entry**' — Type a class or package name into the text box and click the '**Add Entry**' button. The new logger will appear in the list of '**Existing Levels**' in the lower part of the screen.
  - '**Existing Levels**' - These are the loggers currently in action for your Confluence instance.
    - You can change the logging level by selecting a value from the '**New Level**' dropdown list. Read the [Apache documentation](#) for a definition of each level.
    - Click the '**Remove**' link to stop logging for the selected class/package name.
4. Click the '**Save**' button to save any changes you have made in the '**Existing Levels**' section.

## Screenshot: Changing Log Levels and Profiling

## Logging and Profiling

Performance Profiling  
Profiling is currently OFF.  
[Enable Profiling](#)

---

SQL Logging  
[Enable SQL Logging](#)

---

Log4j Logging  
loglevel.profileDesc  
[Production](#) [Diagnostic](#)

**OR:**  
Customise specific logging settings

**Add New Entry**

Class/Package Name	New Level	
<input type="text"/>	INFO ▼	<a href="#">Add entry</a>

**Existing Levels**

Class/Package Name	Current Level	New Level	
atlassian.plugin	INFO	INFO ▼	<a href="#">Remove</a>
com.atlassian.confluence.admin.actions	INFO	INFO ▼	<a href="#">Remove</a>

**Editing the Properties File**

To configure the logging levels and other settings on a permanent basis, you need to stop Confluence and then change the settings in the `log4j.properties` file, described [above](#).

The properties file contains a number of entries for different loggers that can be uncommented if you are interested in logging from particular components. Read more in the [Apache log4j documentation](#).

See [Working with Confluence Logs](#) for some guidelines on specific configuration options you may find useful.

**Configuring Levels for java.util.logging in logging.properties**

A few libraries used by Confluence use `java.util.logging` rather than `log4j` or `slf4j`. These libraries include:

- `com.sun.jersey`
- `org.apache.shindig`
- `net.sf.ehcache`

Confluence's `logging.properties` file is set to redirect `java.util.logging` at specific levels to `log4j` via `slf4j`.

To increase logging levels for these libraries you must first configure the `logging.properties` file in `<CONFLUENCE-INSTALL>/confluence/WEB-INF/classes/`. The logging levels are different from `log4j` and are listed [here](#).

For example, to increase logging for `shindig` change the following line in the `logging.properties` file:

```
org.apache.shindig.level = INFO
```

to

```
org.apache.shindig.level = FINE
```

And then use one of the methods above **as well** to configure the log4j level.

## log4j Logging Levels

### Logging Levels

- **DEBUG** - designates fine-grained informational events that are most useful to debug an application (*what is going on*)
- **INFO** - announcements about the normal operation of the system - scheduled jobs running, services starting and stopping, user-triggered processes and actions
- **WARN** - any condition that, while not an error in itself, may indicate that the system is running sub-optimally
- **ERROR** - a condition that indicates something has gone wrong with the system
- **FATAL** - a condition that indicates something has gone wrong so badly that the system can not recover
- **TRACE** - n/a within confluence

There are two ways to modify the logging levels, as described in [Working with Confluence Logs](#).

1. Modifying the runtime log levels via the **Administration Console**.
2. Manually modifying the <Confluence-Install>\confluence\WEB-INF\classes\log4j.properties file.

### Default Log Level

The standard Confluence log level **WARN** is a way for Confluence to communicate with the server administrator. Logging at WARN level and higher should be reserved for situations that require some kind of attention from the server administrator, and for which corrective action is possible.

See [log4j manual](#) for more information.

## Troubleshooting SQL Exceptions

If you get an exception similar to those shown below, it is a good idea to increase the logging levels of your Confluence instance. If you request [Atlassian support](#), this additional logging will help us work out the cause of the error.

Increased logging levels will enable us to diagnose errors like these:

```
org.springframework.dao.DataIntegrityViolationException: (HibernateTemplate): data
integrity violated by SQL ''; nested exception is java.sql.BatchUpdateException:
Duplicate entry '1234' for key 1
 at
 org.springframework.jdbc.support.SQLStateSQLExceptionTranslator.translate(SQLStateS
QLExceptionTranslator.java:88)
 caused by: java.sql.BatchUpdateException: Duplicate entry '1234' for key 1
 at
 com.mysql.jdbc.ServerPreparedStatement.executeBatch(ServerPreparedStatement.java:64
7)
```

or

```
(HibernateTemplate): data integrity violated by SQL ''; nested exception is
java.sql.BatchUpdateException: ORA-00001: unique constraint
(CONFLUENCE.SYS_C0012345) violated
```

This document outlines the steps to take to increasing logging on your system.


### Changing the logging levels via the Administration Console

With Confluence 2.7 and later, you can adjust logging levels at runtime via the Administration Console — read the [instructions](#). Below we tell you how to edit the log4j files directly.

1. Open `confluence/WEB-INF/classes/log4j.properties` and uncomment the following lines. The double `##` lines are comments, leave them intact.

```
log hibernate prepared statements/SQL queries (equivalent to setting
'hibernate.show_sql' to 'true')
#log4j.logger.net.sf.hibernate.SQL=DEBUG

log hibernate prepared statement parameter values
#log4j.logger.net.sf.hibernate.type=DEBUG
```

-  If you can not locate these lines in your `log4j.properties` file, please add them to the end of it.
2. Restart Confluence.
3. Redo the steps that led to the error.
4. Zip up your logs directory and attach it your support ticket.
5. If you are using Oracle and received a **constraint error**, please ask your database administrator which *table* and *column* the constraint (that is, `CONFLUENCE.SYS_C0012345`) refers to and add that information to your support ticket.
6. Open `confluence/WEB-INF/classes/log4j.properties` again and remove the 4 lines you added in step 1. (The additional logging will impact performance and should be disabled once you have completed this procedure.)

#### RELATED TOPICS

[Enabling Detailed SQL Logging](#)

[Working with Confluence Logs](#)

[Troubleshooting failed XML site backups](#)

### Configuring Confluence Security

This section gives guidelines on configuring the security of your Confluence site:

- [Confluence Security Overview and Advisories](#)
- [Confluence Cookies](#)
- [Configuring Secure Administrator Sessions](#)
- [Using Fail2Ban to limit login attempts](#)
- [Securing Confluence with Apache](#)
- [Trackback and External Referrers](#)
- [Best Practices for Configuring Confluence Security](#)
- [Hiding the People Directory](#)
- [Configuring Captcha for Spam Prevention](#)
- [Hiding External Links From Search Engines](#)
- [Configuring Captcha for Failed Logins](#)
- [Configuring XSRF Protection](#)
- [User Email Visibility](#)
- [Anonymous Access to Remote API](#)

#### Related pages:

- [Permissions and Restrictions](#)
- [Configuring a Confluence Environment](#)
- [Confluence Administrator's Guide](#)

- [Running Confluence Over SSL or HTTPS](#)
- [Connecting to LDAP or JIRA applications or Other Services via SSL](#)
- [Configuring RSS Feeds](#)
- [Preventing and Cleaning Up Spam](#)

## Confluence Security Overview and Advisories

This document is for system administrators who want to evaluate the security of the Confluence web application. The page addresses overall application security and lists the security advisories issued for Confluence. As a public-facing web application, Confluence's application-level security is important. This document answers a number of questions that commonly arise when customers ask us about the security of our product.

Other topics that you may be looking for:

- For information about user management, groups and permissions, please refer to the [internal security overview](#).
- For guidelines on configuring the security of your Confluence site, see the [administrator's guide to configuring Confluence security](#).

## Application Security Overview

### Password Storage

When Confluence's internal user management is used, since version 3.5 of Confluence passwords are hashed through the salted [PKCS5S2 implementation provided by Embedded Crowd](#) before being stored in the database. There is no mechanism within Confluence to retrieve a user's password – when password recovery is performed, a reset password link is generated and mailed to the user's registered address.

When external user management is enabled, password storage is delegated to the external system.

#### On this page:

- [Application Security Overview](#)
- [Finding and Reporting a Security Vulnerability](#)
- [Publication of Confluence Security Advisories](#)
- [Severity Levels](#)
- [Our Security Bugfix Policy](#)
- [Published Security Advisories](#)

### Buffer Overflows

Confluence is a 100% pure Java application with no native components. As such it is highly resistant to buffer overflow vulnerabilities – possible buffer overruns are limited to those that are bugs in the Java Runtime Environment itself.

### SQL Injection

Confluence interacts with the database through the Hibernate Object-Relational mapper. Database queries are generated using standard APIs for parameter replacement rather than string concatenation. As such, Confluence is highly resistant to SQL injection attacks.

### Script Injection

Confluence is a self-contained Java application and does not launch external processes. As such, it is highly resistant to script injection attacks.

### Cross-Site Scripting

As a content-management system that allows user-generated content to be posted on the web, precautions have been taken within the application to prevent cross-site scripting attacks:

- The wiki markup language in Confluence does not support dangerous HTML markup
- Macros allowing the insertion of raw HTML are disabled by default
- HTML uploaded as a file attachment is served with a content-type requesting the file be downloaded, rather than being displayed inline
- Only system administrators can make HTML-level customisations of the application

When cross-site scripting vulnerabilities are found in the Confluence web application, we endeavour to fix them as quickly as possible.

#### Transport Layer Security

Confluence does not directly support SSL/TLS. Administrators who are concerned about transport-layer security should set up SSL/TLS at the level of the Java web application server, or the HTTP proxy in front of the Confluence application.

For more information on configuring Confluence for SSL, see: [Running Confluence Over SSL or HTTPS](#)

#### Session Management

Confluence delegates session management to the Java application server in which it is deployed. We are not aware of any viable session-hijacking attacks against the Tomcat application server shipped with Confluence. If you are deploying Confluence in some other application server, you should ensure that it is not vulnerable to session hijacking.

#### Plugin Security

Administrators install third party plugins **at their own risk**. Plugins run in the same virtual machine as the Confluence server, and have access to the Java runtime environment, and the Confluence server API.

Administrators should always be aware of the source of the plugins they are installing, and whether they trust those plugins.

#### Administrator Trust Model

Confluence is written under the assumption that anyone given [System Administrator privileges](#) is trusted. System administrators are able, either directly or by installing plugins, to perform any operation that the Confluence application is capable of.

As with any application, you should not run Confluence as the root/Administrator user. If you want Confluence to listen on a privileged network port, you should set up port forwarding or proxying rather than run Confluence with additional privileges. The extra-careful may consider running Confluence inside a `chroot` jail.

#### Stack Traces

To help debug support cases and provide legendary support, Confluence provides stack traces through the web interface when an error occurs. These stack traces include information about what Confluence was doing at the time, and some information about your deployment server.

Only non-personal information is supplied such as operating system and version and Java version. With proper network security, this is not enough information to be considered dangerous. No usernames or passwords are included.

#### Finding and Reporting a Security Vulnerability

Atlassian's approach to reporting security vulnerabilities is detailed in [How to Report a Security Issue](#).

#### Publication of Confluence Security Advisories

Atlassian's approach to releasing security advisories is detailed in [Security Advisory Publishing Policy](#).

#### Severity Levels

Atlassian's approach to ranking security issues is detailed in [Severity Levels for Security Issues](#).

#### Our Security Bugfix Policy

Our approach to releasing patches for security issues is detailed in our [Security Bugfix Policy](#).

## Published Security Advisories

### Confluence Cookies

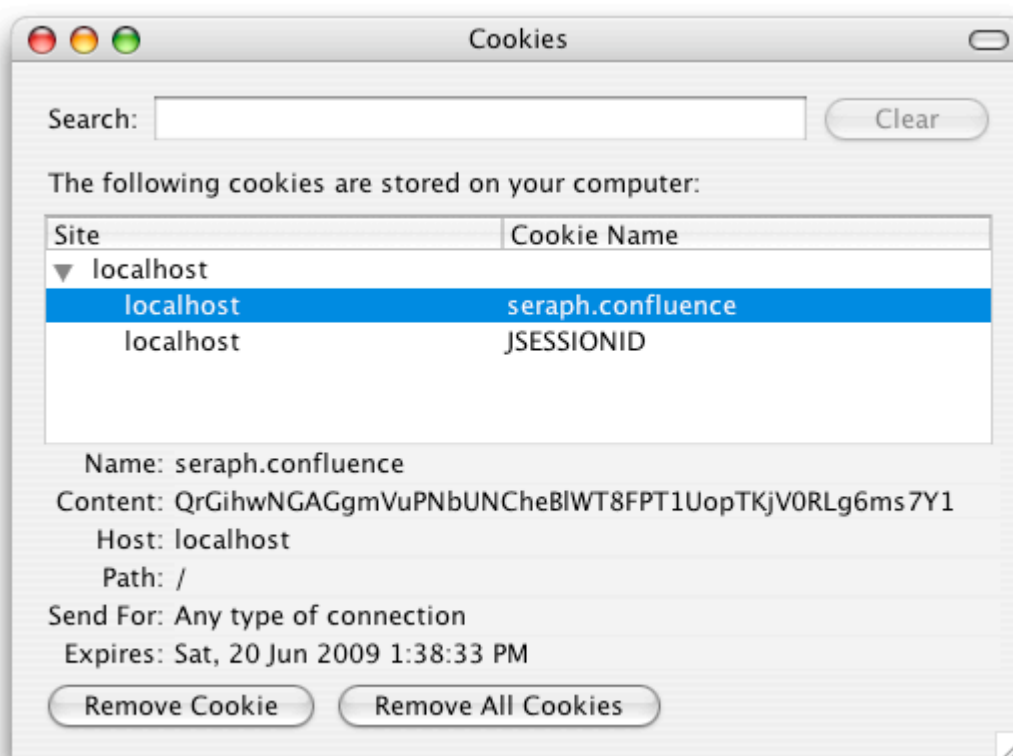
This page lists cookies stored in Confluence users' browsers which are generated by Confluence itself. This page does not list cookies that may originate from 3rd-party Confluence plugins.

#### Authentication cookies

Confluence uses [Seraph](#), an open source framework, for HTTP cookie authentication. Confluence uses two types of cookies for user authentication:

- The JSESSIONID cookie is created by the application server and used for session tracking purposes. This cookie contains a random string and the cookie expires at the end of every session or when the browser is closed.
- The 'remember me' cookie, `seraph.confluence`, is generated by Confluence when the user selects the **Remember me** check box on the login page.

 You can read about cookies on the [Wikipedia page about HTTP cookies](#).



#### On this page:

- [Authentication cookies](#)
- [Other Confluence cookies](#)
- [Notes](#)

#### The 'remember me' cookie

The 'remember me' cookie, `seraph.confluence`, is a long-lived HTTP cookie. This cookie can be used to authenticate an unauthenticated session. Confluence generates this cookie when the user selects the **Remember me** check box on the login page.

#### Cookie key and contents

By default, the cookie key is `seraph.confluence`, which is defined by the `login.cookie.key` parameter in the `CONFLUENCE-INSTALLATION/confluence/WEB-INF/classes/seraph-config.xml` file.

The cookie contains a unique identifier plus a securely-generated random string (i.e. token). This token is generated by Confluence and is also stored for the user in the Confluence database.

#### *Use of cookie for authentication*

When a user requests a web page, if the request is not already authenticated via session-based authentication or otherwise, Confluence will match the 'remember me' cookie (if present) against the token (also if present), which is stored for the user in the Confluence database.

If the token in the cookie matches the token stored in the database and the cookie has not expired, the user is authenticated.

#### *Life of 'remember me' cookies*

You can configure the maximum age of the cookie. To do that you will need to modify the `CONFLUENCE-INSTALLATION/confluence/WEB-INF/classes/seraph-config.xml` file and insert the following lines below the other `init-param` elements:

```
<init-param>
 <param-name>autologin.cookie.age</param-name>
 <param-value>2592000</param-value><!-- 30 days in seconds -->
</init-param>
```

#### *Automatic cleanup of 'remember me' tokens*

Every cookie issued by Confluence has a corresponding record in the database. A scheduled job runs on the 20th of every month to clean up expired tokens. The name of the trigger is `clearExpiredRememberMeToken` `sTrigger`.

*Note:* The only purpose of this job is to prevent the database table from growing too big. For authentication purposes, Confluence will ignore expired tokens even if they still exist in the database.

#### *Is it possible to disable the 'remember me' feature?*

Confluence does not offer an option for disabling the 'Remember Me' feature. See the [workaround](#).

#### **Other Confluence cookies**

There are several cookies that Confluence uses to store basic 'product presentation' states. Confluence users' authentication details are not stored by these cookies.

Cookie Key	Purpose	Cookie Contents	Expiry
<b>doc-sidebar</b>	Remembers the user's preference for the width of the navigation sidebar in the Confluence documentation theme.	The width of the sidebar in pixels. For example, 300px	One year from the date it was set or was last updated.
<b>confluence.list.pages.cookie</b>	Remembers the user's last chosen tab in the "list pages" section.	The name of the last selected tab. For example, list-content-tree	One year from the date it was set or was last updated.

<b>confluence.browse.space.cookie</b>	Remembers the user's last chosen tab in the "browse space" section	The name of the last selected tab. For example, space-pages	One year from the date it was set or was last updated.
<b>confluence-language</b>	Remembers the user's language chosen on the login page. This cookie relates to a feature that allows a user to change Confluence's language from (and including) the login page, when the language presented to the user prior to logging in is not appropriate.	A locale relating to the chosen language. For example, de_DE	360 days from the date it was set or was last updated.
<b>AJS.conglomerate.cookie</b>	Tracks which general tabs were last used or expansion elements were last opened or closed.	One or more key-value strings which indicate the states of your last general tab views or expansion elements.	One year from the date it is set or was last updated.

#### Notes

- The *autocomplete* feature in browser text fields (which are typically noticeable when a user logs in to Confluence) is a browser-specific feature, not a Confluence one. Confluence cannot enable or disable this autocompletion, which is typically set through a browser's settings.

### Configuring Secure Administrator Sessions

Confluence protects access to its administrative functions by requiring a secure administration session to use the [Confluence administration console](#) or [administer a space](#). When a Confluence administrator (who is logged into Confluence) attempts to access an administration function, they are prompted to log in again. This logs the administrator into a temporary secure session that grants access to the Confluence/space administration console.

The temporary secure session has a rolling timeout (defaulted to 10 minutes). If there is no activity by the administrator in the Confluence/space administration console for a period of time that exceeds the timeout, then the administrator will be logged out of the secure administrator session (note, they will remain logged into Confluence). If the administrator does click an administration function, the timeout will reset.

To configure secure administrator sessions:

1. Choose the **cog icon**




, then choose **General Configuration** under Confluence Administration

2. Choose **Security Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Configure the setting as follows:
  - To disable secure administrator sessions, uncheck the **Enable** check box next to **Secure administrator sessions**. When this setting is disabled, administrators will no longer be required to log into a secure session to access the administration console.
  - To change the timeout for secure administrator sessions, update the value next to **minutes before invalidation**. The default timeout for a secure administration session is 10 minutes.
5. Choose **Save**.

#### Notes

- **Disabling password confirmation.** Confluence installations that use a custom authentication mechanism may run into problems with the Confluence security measure that requires password confirmation. If necessary, you can set the `password.confirmation.disabled` system property to disable the password confirmation functionality. See [Recognised System Properties](#). See issue [CONF-20958](#) "Confluence features that require password confirmation (websudo, captcha) do not work with custom authentication".
- **WebSudo.** The feature that provides secure administrator sessions is also called 'WebSudo'.
- **Manually ending a secure session.** An administrator can choose to manually end their secure session by clicking the '**drop access**' link in the banner displayed at the top of their screen. For example:


 You have temporary access to administrative functions. [Drop access](#) if you no longer require it. For more information, refer to [the documentation](#).

- **Note for developers.** Secure administrator sessions can cause exceptions when developing against Confluence or deploying a plugin. Please read this FAQ: [How do I develop against Confluence with Secure Administrator Sessions?](#) Note: The Confluence XML-RPC and REST APIs are not affected by secure administration sessions.

## Using Fail2Ban to limit login attempts

### What is Fail2Ban?

We need a means of defending sites against brute-force login attempts. [Fail2Ban](#) is a Python application which trails logfiles, looks for [regular expressions](#) and works with Shorewall (or directly with iptables) to apply temporary blacklists against addresses that match a pattern too often. This can be used to limit the rate at which a given machine hits login URLs for Confluence.

 The information on this page *does not apply to* Confluence Cloud.

### Prerequisites

- Requires [Python](#) 2.4 or higher to be installed
- Needs a specific file to follow, which means your Apache instance needs to log your Confluence access to a known logfile. You **should adjust the configuration below** appropriately.

### How to set it up

*This list is a skeletal version of the instructions*

- There's an RPM available for RHEL on the [download page](#), but you can also download the source and set it up manually
- Its configuration files go into `/etc/fail2ban`
- The generic, default configuration goes into `.conf` files (`fail2ban.conf` and `jail.conf`). Don't change these, as it makes upgrading difficult.
- Overrides to the generic configuration go into `.local` files corresponding to the `.conf` files. These only need to contain the specific settings you want overridden, which helps maintainability.
- Filters go into `filter.d` — this is where you define regexps, each going into its own file
- Actions go into `action.d` — you probably won't need to add one, but it's handy to know what's available
- "jails" are a configuration unit that specify one regexp to check, and one or more actions to trigger when the threshold is reached, plus the threshold settings (e.g. more than 3 matches in 60 seconds causes that address to be blocked for 600 seconds)
- Jails are defined in `jail.conf` and `jail.local`. Don't forget the `enabled` setting for each one — it can be as bad to have the wrong ones enabled as to have the right ones disabled.

### Running Fail2Ban

- Use `/etc/init.d/fail2ban {start|stop|status}` for the obvious operations
- Use `fail2ban-client -d` to get it to dump its current configuration to STDOUT. Very useful for troubleshooting.
- Mind the CPU usage; it can soak up resources pretty quickly on a busy site, even with simple regexp
- It can log either to syslog or a file, whichever suits your needs better

## Common Configuration

### jail.local

```
The DEFAULT allows a global definition of the options. They can be
override
in each jail afterwards.

[DEFAULT]

"ignoreip" can be an IP address, a CIDR mask or a DNS host. Fail2ban
will not
ban a host which matches an address in this list. Several addresses
can be
defined using space separator.
ignoreip = <space-separated list of IPs>

"bantime" is the number of seconds that a host is banned.
bantime = 600

A host is banned if it has generated "maxretry" during the last
"findtime"
seconds.
findtime = 60

"maxretry" is the number of failures before a host get banned.
maxretry = 3

[ssh-iptables]

enabled = false

[apache-shorewall]

enabled = true
filter = cac-login
action = shorewall
logpath = /var/log/httpd/confluence-access.log
bantime = 600
maxretry = 3
findtime = 60
backend = polling
```

## Configuring for Confluence

The following is an example only, and you should adjust it for your site.

### filter.d/confluence-login.conf

```
[Definition]
```

```
failregex = <HOST>.*"GET /login.action
```

```
ignoreregex =
```

## Securing Confluence with Apache

The following outlines some basic techniques to secure a Confluence instance using Apache. These instructions are basic to-do lists and should not be considered comprehensive. For more advanced security topics see the "Further Information" section below.

- [Using Apache to limit access to the Confluence administration interface](#)
- [Using Fail2Ban to limit login attempts](#)

### Further Information

#### [Running Confluence behind Apache](#)

#### Using Apache to limit access to the Confluence administration interface

##### Limiting administration to specific IP addresses

The Confluence administration interface is a critical part of the application; anyone with access to it can potentially compromise not only the Confluence instance but the entire machine. As well as limiting access to users who really need it, and using strong passwords, you should consider limiting access to it to certain machines on the network or internet. If you are using an [Apache web server](#), this can be done with Apache's **Location** functionality as follows:

##### 1. Create a file that defines permission settings

This file can be in the Apache configuration directory or in a system-wide directory. For this example we'll call it "sysadmin\_ips\_only.conf". The file should contain the following:

```
Order Deny,Allow
Deny from All

Mark the Sysadmin's workstation
Allow from 192.168.12.42
```

##### 2. Add the file to your Virtual Host

In your Apache Virtual Host, add the following lines to restrict the administration actions to the Systems Administrator:

This configuration assumes you've installed Confluence under '/confluence'. If you have installed under '/' or elsewhere, adjust the paths accordingly.

```
<Location /confluence/admin>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/consumers/list>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/view-consumer-info>
 Include sysadmin_ips_only.conf
</Location>
```

```
<Location /confluence/plugins/servlet/oauth/service-providers/list>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/service-providers/add>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/consumers/add>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/consumers/add-manually>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/update-consumer-info>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/pages/templates/listpagetemplates.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/pages/templates/createpagetemplate.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/spacepermissions.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/pages/listpermissionpages.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/removespace.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/importtmbbox.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/viewmailaccounts.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/addmailaccount.action?>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/importpages.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/flyingpdf/flyingpdf.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/exportspacehtml.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/exportspacexml.action>
 Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/embedded-crowd>
 Include sysadmin_ips_only.conf
</Location>
```

```
<Location /confluence/plugins/servlet/upm>
 Include sysadmin_ips_only.conf
</Location>
```


## Trackback and External Referrers

### Trackback

When Trackback is enabled, any time you link to an external webpage that supports Trackback Autodiscovery, Confluence will send a trackback ping to that page to inform it that it has been linked to.

Confluence pages also support Trackback Autodiscovery and when Trackback is enabled, can receive trackback pings from other sites.

To enable trackback:

1. Go to  **> General Configuration > Further Configuration**
2. Choose **Edit**
3. Select the **Trackback** checkbox then **Save**

### External referrers

An external referrer is any site that links to your Confluence instance. Each time someone clicks on the external link, your Confluence site can record the click as a referral.

By default, external referrers for a page are listed under '**Hot Referrers**' on the '**Info**' screen of the page. Confluence shows a maximum of 10 referrers. If there are more than 10, confluence shows the 10 with the highest number of hits.

Note that you do *not* need to enable trackback in order to have external referrers enabled.

*Screenshot: hot referrers on the page information screen.*

### On this page:

- [Trackback](#)
- [External referrers](#)
  - [Excluding external referrers](#)
  - [Hiding external referrers](#)
  - [Ignoring external referrers](#)

## Page Information

Title:	The new logo
Author:	<a href="#">Sophie Staunton</a> Feb 15, 2013
Last Changed by:	<a href="#">Sophie Staunton</a> Feb 15, 2013
Tiny Link: (useful for email)	<a href="http://localhost:8090/x/HwA8">http://localhost:8090/x/HwA8</a>
Export As:	<a href="#">Word</a> · <a href="#">PDF</a>
Operations:	<a href="#">Copy</a>

### Labels

There are no labels assigned to this page.

### Hot Referrers

(2)	<a href="https://mail.google.com/mail/u/1/">https://mail.google.com/mail/u/1/</a>	<a href="#">[ban]</a>
(1)	<a href="https://atlassian.hipchat.com/chat">https://atlassian.hipchat.com/chat</a>	<a href="#">[ban]</a>

### Excluding external referrers

Confluence Administrators can exclude external referrers to prevent them from being recorded or displayed anywhere on your site. Once you have specified your list of blocked URLs, any incoming links from URLs that match the list will no longer be recorded.

Referrer URLs are blocked if they start with any of the URLs in the exclusion list. So <http://evilspamsite.blogspot.com> will also match <http://evilspamsite.blogspot.com/nastypage.html>

You might want to do this if:

1. **You're running a Confluence installation that is open to public**

In a site that is open to public, one unfortunate problem is that malicious sites can spam the display of a page's incoming links statistics. This is usually done to get the site's URL to appear in the sidebar. By adding these sites to the 'excluded referrers' list, you can prevent them from being listed on your site.

2. **Confluence is installed on a server with multiple domain names or IP addresses**

Confluence will consider any URL originating from the domain name where Confluence is installed as an internal link. However, if Confluence is installed on a server with multiple domain names or IP addresses, you will need to add the other domain name prefixes to this list to let Confluence know that any links from these domains should not be considered external links.

To add a URL to the excluded referrers list:

1. Go to



> **General Configuration > Manage Referrers**

2. Enter the URL in the **Referrer URL Prefix** field (you must include http://)
3. Choose **Add**

You can add multiple URLs to the list.

### Exclude External Referrers

Referrer URL Prefix

Add a URL Prefix that will no longer be recorded in the External Referrers. All URLs that start with this prefix will be excluded.


### Excluded Referrer URL Prefixes

URL Prefix	Operations
	<a href="#">Purge All</a>
<a href="http://evilspamsite.blogspot.com">http://evilspamsite.blogspot.com</a>	<a href="#">Delete</a> · <a href="#">Purge</a>

#### Hiding external referrers

By default, Confluence lists the external referrers as '**Hot Referrers**' on the page information screen for a page. If you turn this option off, external referrers will not be listed on the page.

To hide external referrers:

- Go to  > **General Configuration** > **Manage Referrers**.
- Deselect **Show Referrers in Page Info**.

Screenshot: Managing external referrers

### External Referrer Settings

Record External Referrers ☒ [More about External Referrers](#)  
 Allow Confluence to record the external URLs that link to Confluence pages.


Show Referrers in Page Info ☐  
 Show the top 10 external links pointing to that page on the 'Info' screen.

[Cancel](#)

#### Ignoring external referrers

An external referrer is any site that links to your Confluence instance. Each time someone clicks on the external link, your Confluence site can record the click as a referral. By default, Confluence records the number of hits made to a page from any link on an external site. If you turn this option off, Confluence will not record the hits.

To ignore external referrers:

- Go to  > **General Configuration** > **Manage Referrers**
- Deselect **Record External Referrers**

Screenshot: Managing external referrers

## External Referrer Settings

Record External Referrers

☐

Allow Confluence to record the external URLs that link to Confluence pages.  
[More about External Referrers](#)

Show Referrers in Page Info

☐

Show the top 10 external links pointing to that page on the 'Info' screen.

Submit

Cancel

## Best Practices for Configuring Confluence Security

The best way to harden a system is to look at each of the involved systems individually. Contact your company's security officer or department to find out what security policies you should be using. There are many things to consider, such as the configuration of your underlying operating systems, application servers, database servers, network, firewall, routers, etc. It would be impossible to outline all of them here.

This page contains guidelines on good security practices, to the best of our knowledge.

### Configuring the Web Server

Please refer to the following guides for system administrators:

- How to configure Apache to lock down the administration interface to those people who really need it: [Using Apache to limit access to the Confluence administration interface](#).
- How to reduce the risk of brute force attacks: [Using Fail2Ban to limit login attempts](#).

### Configuring the Application Server

See the following system administrator guide for general hints on the application server level:

- [Tomcat security best practices](#)

### Configuring the Application

The way you set up Confluence roles, permissions and processes makes a big difference in the security of your Confluence site.

Below are some more Confluence-specific items to consider. None of these provides 100% security. They are measures to reduce impact and to slow down an intruder in case your system does become compromised.

- Keep the number of Confluence administrators extremely low. For example, 3 system administrator accounts should be the maximum.
- Similarly, restrict the number of users with powerful roles or group memberships. If only one department should have access to particularly sensitive data, then do restrict access to the data to those users. Do not let convenience over-rule security. Do not give all staff access to sensitive data when there is no need.
- The administrators should have separate Confluence accounts for their administrative roles and for their day to day roles. If John Doe is an administrator, he should have a regular user account without administrator access to do his day to day work (such as writing pages in the wiki). This could be a 'john.doe' account. In addition, he should have an entirely separate account (that cannot be guessed by an outsider and that does not even use his proper name) for administrative work. This account could be 'jane smith' – using a username that is so obscure or fake that no outsider could guess it. This way, even if an attacker singles out the actual person John Doe and gets hold of his password, the stolen account would most likely be John's regular user account, and the attacker cannot perform administrative actions with that account.

- Lock down administrative actions as much as you can. If there is no need for your administrators to perform administrative actions from outside the office, then lock down access to those actions to known IP addresses, for example. See [Using Apache to limit access to the Confluence administration interface](#).
- Put documented procedures in place for the case of employees leaving the company.
- Perform security audits regularly. Know who can help in case a security breach occurs. Perform 'what if' planning exercises. ('What is the worst thing that could happen if a privileged user's password were stolen while he's on vacation? What can we do to minimise damage?').
- Make sure the Confluence database user (and all datasource database users) only has the amount of database privileges it really needs.
- Monitor your binaries. If an attacker compromises an account on your system, he will usually try to gain access to more accounts. This is sometimes done by adding malicious code, such as by modifying files on the system. Run routine scripts that regularly verify that no malicious change has been made.

As another precaution:

- Regularly monitor the above requirements. There are many things that could start out well, but deteriorate over time:
  - A system may start out with just 3 administrators, but over the course of a year this could grow to 30 administrators if no one prevents expansion.
  - Apache administration restrictions may be in place at the start of the year, but when the application server is migrated after a few months, people may forget to apply the rules to the new system.

Again, keep in mind that the above steps may only be a fraction of what could apply to you, depending on your security requirements. Also, keep in mind that none of the above rules can guarantee anything. They just make it harder for an intruder to move quickly.

### Hiding the People Directory

The [People Directory](#) provides a list of all users in your Confluence system.

If you need to disable the People Directory [set the following system properties](#) on your application server command line:

- **To disable the People Directory for anonymous users:**

```
-Dconfluence.disable.peopledirectory.anonymous=true
```

- **To disable the People Directory entirely:**

```
-Dconfluence.disable.peopledirectory.all=true
```

This workaround will prevent the People directory from appearing on the dashboard, but if you navigate to the profile of a user, and then click on the "People" in the breadcrumb link (Dashboard >> **People** >> FullName >> Profile) or you go to the URL directly `<CONFLUENCE_INSTALL>/browsepeople.action`, you will be able to access the people directory.

To workaround this, set up your Apache webserver in front of Confluence and redirect requests to this URL.

### Configuring Captcha for Spam Prevention

If your Confluence site is open to the public (you allow anonymous users to add comments, create pages etc) you may find that automated spam is being added, in the form of comments or new pages.

You can configure Confluence to deter automated spam by asking users to prove that they are human before they are allowed to:

- Sign up for an account.
- Add a comment.
- Create a page.

#### Related pages:

- [Configuring Confluence Security](#)

- Edit a page.
- Send a request to the Confluence administrators.

Captcha is a test that can distinguish a human being from an automated agent such as a web spider or robot. When Captcha is switched on, users will see a distorted picture of a word, and must enter it in a text field before they can proceed.

*Screenshot: Example of a Captcha test*



By default, Captcha is disabled. When enabled, the default is that only anonymous users will have to perform the Captcha test when creating comments or editing pages. You can also choose to enforce Captcha for all users or members of particular groups.

You need System Administrator permissions to configure Captcha for spam prevention in Confluence.

#### To enable Captcha for spam prevention in Confluence:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Spam Prevention** in the left-hand panel
3. Choose **ON** to turn on Captcha
4. If you want to disable Captcha for certain groups:
  - Select **No one** if you want everyone to see Captchas.
  - Select **Signed in users** if you want only anonymous users to see Captchas.
  - If you want everyone to see Captchas except members of specific groups, select **Members of the following groups** and enter the group names in the text box.  
You can click the magnifying-glass icon to search for groups. Search for all or part of a group name and click the **Select Groups** button to add one or more groups to the list.
  - To remove a group from the list, delete the group name
5. Choose **Save**

#### Hiding External Links From Search Engines

Hiding external links from search engines helps to discourage spammers from posting links on your site. If you turn this option on, any URLs inserted in pages and comments will be given the 'nofollow' attribute, which prevents search engines from following them.

**i** Shortcut links (e.g. CONF-2622@JIRA) and internal links to other pages within Confluence are not tagged.

To hide external links from search engines:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Click **'Security Configuration'** in the left panel.
3. This will display the **'Security Configuration'** screen. Click **'Edit'**.
4. Check the **'Hide External Links From Search Engines'** checkbox.
5. Click the **'Save'** button.

#### Background to the nofollow attribute

As part of the effort to combat the spamming of wikis and blogs (Confluence being both), Google came up with [some markup](#) which instructs search engines not to follow links. By removing the main benefit of wiki-spamming it's hoped that the practice will stop being cost-effective and eventually die out.

## Configuring Captcha for Failed Logins

If you have Confluence administrator permissions, you can configure Confluence to impose a maximum number of repeated login attempts. After a given number of failed login attempts (the default is three) Confluence will display a Captcha form asking the user to enter a given word when attempting to log in again. This will prevent brute force attacks on the Confluence login screen.

Similarly, after three failed login attempts via the XML-RPC or SOAP API, an error message will be returned instructing the user to log in via the web interface. Captcha will automatically be activated when they attempt this login.

'Captcha' is a test that can distinguish a human being from an automated agent such as a web spider or robot. When Captcha is activated, users will need to recognise a distorted picture of a word, and must type the word into a text field. This is easy for humans to do, but very difficult for computers.

*Screenshot: example of a Captcha test*



### On this page:

- [Enabling, Disabling and Configuring Captcha for Failed Logins](#)
- [Notes](#)

## Enabling, Disabling and Configuring Captcha for Failed Logins

By default, Captcha for failed logins is enabled and the number of failed login attempts is set to three.

**To enable, disable and configure Captcha for failed logins:**

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose '**Security Configuration**' from the left menu.
3. Choose '**Edit**'.
4. To enable Captcha:
  - Select the '**Enable**' checkbox next to '**CAPTCHA on login**'.
  - Set the maximum number of failed logins next to '**Maximum Authentication Attempts Allowed**'. You must enter a number greater than zero.
5. To **disable** Captcha, deselect the '**Enable**' checkbox.
6. Choose '**Save**'.

*Screenshot: Configuring Captcha for failed logins*

## Security and Privacy

Settings for user management, site security and user privacy.

☐ External user management

Delegate user management to JIRA. [More about User Management](#)

☒ Append wildcards to user and group searches

☒ Hide External Links From Search Engines

This helps discourage spammers from posting malicious links by preventing search engines to follow the site. [More about External Links](#)

☐ Anonymous Access to Remote API

Enabling this will allow 'anonymous' to access Confluence remotely.

☐ Custom Stylesheets for Spaces

☐ Show system information on the 500 page

User email visibility

Maximum RSS Items

Limit the maximum number of items an RSS Feed can request.

RSS timeout  seconds

The time in seconds allowed to create each RSS Feed. Any items rendered within the timeout will still be returned.

Page timeout  seconds

The time in seconds allowed to render the content of each wiki Page. Pages taking longer to render will display a timeout error to the user. The default is 120 seconds.

CAPTCHA on login ☒ Enable

maximum authentication attempts allowed

Secure administrator sessions ☒ Enable

minutes before automatic invalidation

Requires Confluence Administrator login to access administration functions. [More about Secure Administrator Sessions](#)

### Notes

- **Disabling all password confirmation requests, including Captcha on login.** Confluence installations that use a custom authentication mechanism may run into problems with the Confluence security measure that requires password confirmation. If necessary, you can set the `password.confirmation.disabled` system property to disable the password confirmation functionality on [administrative actions](#), change of email address and Captcha for failed logins. See [Recognised System Properties](#).

### Configuring XSRF Protection

Confluence requires an XSRF token to be present on comment creation, to prevent users being tricked into unintentionally submitting malicious data. All the themes bundled with Confluence have been designed to use this feature. However, if you are using a custom theme that does not support this security feature, you can disable it.

**⚠ Please carefully consider the security risks before you disable XSRF protection for comments in your Confluence installation.**

Read more about [XSRF \(Cross Site Request Forgery\)](http://cgisecurity.com) at cgisecurity.com.

#### To configure XSRF protection for comments:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Security Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Uncheck the **Adding Comments** checkbox in the **XSRF Protection** section, to disable XSRF protection.
5. Choose **Save**.

##### Related pages:

- [Configuring Confluence Security](#)
- [Confluence Administrator's Guide](#)
- [Developer documentation on XSRF protection in Confluence](#)

#### User Email Visibility

Confluence provides three options for email address privacy which can be configured by a Confluence administrator from the **Administration Console**:

- **Public**: email addresses are displayed publicly.
- **Masked**: email addresses are still displayed publicly, but masked in such a way to make it harder for spam-bots to harvest them.
- **Only visible to site administrators**: only [Confluence administrators](#) can see the email addresses. Note that, if you select this option, email addresses will not be available in the 'User Search' popup (e.g. when setting [Page Restrictions](#)).

#### To configure user email visibility:

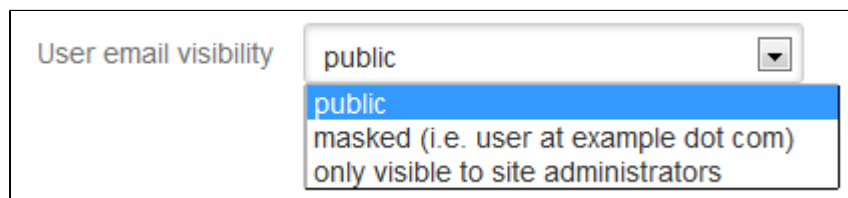
1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **'Security Configuration'**.
3. Choose **'Edit'**. The fields on the **'Security Configuration'** screen will be editable.
4. Select one of the options from the **'User email visibility'** dropdown: **'public'**, **'masked'**, or **'only visible to site administrators'**.
5. Choose **'Save'**.

Screenshot: *Email Visibility*



#### Anonymous Access to Remote API

Administrators may wish to disable anonymous access to the Confluence remote API, to make it harder for malicious users to write 'bots' that perform bulk changes to the site.

#### To disable anonymous access to the remote API:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Security Configuration** in the left-hand panel. The **Security Configuration** screen will appear.
3. Choose **Edit**.
4. Uncheck the **Anonymous Access to API** check box.
5. Choose **Save**.

## Notes

This page is about access to the remote API. If you are looking for information about preventing anonymous users from accessing Confluence, see [Global Permissions Overview](#).

## Running Confluence Over SSL or HTTPS

Atlassian applications allow the use of SSL within our products, however Atlassian Support does not provide assistance for configuring it. Consequently, Atlassian **cannot guarantee providing any support for it**.

- If assistance with conversions of certificates is required, please consult with the vendor who provided the certificate.
- If assistance with configuration is required, please raise a question on [Atlassian Answers](#).

This document tells you how to configure Confluence to enable access via HTTPS (HTTP over SSL), so that your Confluence logins and data are encrypted during transport to and from Confluence. SSL encryption is a good way to safeguard your Confluence data and user logins from being intercepted and read by outsiders.

These instructions apply to the following platforms:

- **Confluence** - Apache Tomcat is the application server shipped with Confluence, and is the only supported application server. If you are using a different application server or Apache HTTP Server ("httpd"), see the page on [Apache with mod\\_proxy](#) for instructions on how to terminate an SSL connection at the Apache web server.
- **Java 8** - JDK 1.8 is the [supported](#) Java version for Confluence. Note that you need the JDK, since it includes the `keytool` utility used in the instructions below. The JRE is not enough.

**i** The default connector port for Confluence is 8090.

## On this page:

- [Step 1. Create or Request a New SSL Certificate](#)
- [Step 2. Modify the Server Configuration File in your Confluence Installation](#)
- [Step 3. Specify the Location of your Certificate](#)
- [Step 4. Change your Confluence Base URL to HTTPS](#)
- [Step 5. Add a Security Constraint to Redirect All URLs to HTTPS](#)
- [Notes](#)
- [Troubleshooting](#)

## Related Topics

- [SSL Configuration HOW-TO in the Apache Tomcat 8.0 documentation](#)
- [keytool - Key and Certificate Management Tool in the Java SE documentation](#)
- [Connecting to LDAP or JIRA applications or Other Services via SSL](#)

## Step 1. Create or Request a New SSL Certificate

You will need a valid SSL certificate before you can enable HTTPS. If you already have a certificate prepared, skip to step 2 [below](#).

You can choose to create a self-signed certificate or to use a certificate issued by a certificate authority (CA, sometimes also called a 'certification authority'). We described both options below.

### Certificate Option 1 – Create a Self-Signed Certificate

Self-signed certificates are useful if you require encryption but do not need to verify the identity of the

requesting website. In general, you might use a self-signed certificate on a test environment and on internal corporate networks (intranets).

Because the certificate is not signed by a certificate authority (CA), users may receive a message that the site is not trusted and may have to perform several steps to accept the certificate before they can access the site. This usually will only occur the first time they access the site.

Follow the steps below to generate a certificate using Java's `keytool` utility. This tool is included in the JDK.

1. Use Java's `keytool` utility to generate the certificate:

Many SSL issuers (including but not limited to GoDaddy and RapidSSL) are now requiring a 2048-bit key size. To generate a key with 2048-bit encryption, add `'-keysize 2048'` to these queries.

- On Windows, run the following command at the command prompt:

```
"%JAVA_HOME%\bin\keytool" -genkeypair -keysize 2048 -alias tomcat
-keyalg RSA -sigalg SHA256withRSA
```

- On OS X or UNIX-based systems, run the following command at the command prompt:

```
$JAVA_HOME/bin/keytool -genkeypair -keysize 2048 -alias tomcat
-keyalg RSA -sigalg SHA256withRSA
```

2. When asked for a **password**:

- Specify the password you want to use for the certificate (private key). Note that the password text will not appear as you type it.
- Make a note of the password you choose, because you will need it in the next step when editing the configuration file.
- The default password is `'changeit'`.

Tomcat has a known issue with passwords containing special characters. You should use a password that only contains alphanumeric characters.

3. Follow the prompts to specify your name, organisation and location. This information is used to construct the X.500 Distinguished Name (DN) of the entity. The CN ("What is your first and last name?") must match the fully-qualified hostname of the server running Confluence, otherwise Tomcat will not be able to use the certificate for SSL. For example for a Confluence running on a server named `"confluence.example.com"`:

```
CN=confluence.example.com, OU=Java Software Division, O=Sun Microsystems
Inc, C=US
```

4. Enter `'y'` to confirm the details.

5. When asked for the **password** for `'tomcat'` (the alias you entered in the `keytool` command above), press the `'Enter'` key. This specifies that your keystore entry will have the **same password** as your private key. You **MUST** use the same password here as was used for the keystore password itself. This is a restriction of the Tomcat implementation.

6. Your certificate is now ready. Go to step 2 [below](#).

### **Certificate Option 2 – Use a Certificate Issued by a Certificate Authority**

When running Confluence in a production environment, you will need a certificate issued by a certificate authority (CA, sometimes also called a 'certification authority') such as [VeriSign](#), [Thawte](#) or [TrustCenter](#). The instructions below are adapted from the [Tomcat documentation](#).

First you will generate a local certificate and create a 'certificate signing request' (CSR) based on that certificate. You will submit the CSR to your chosen certificate authority. The CA will use that CSR to generate a certificate for you.

1. Use Java's `keytool` utility to generate a local certificate, as described in the previous section.
2. Use the `keytool` utility to generate a CSR, replacing the text `<MY_KEYSTORE_FILENAME>` with the path to and file name of the `.keystorefile` generated for your local certificate:

```
keytool -certreq -keyalg RSA -alias tomcat -file certreq.csr -keystore
<MY_KEYSTORE_FILENAME>
```

3. Submit the generated file called `certreq.csr` to your chosen certificate authority. **Refer to the documentation on the CA's website to find out how to do this.**
4. The CA will send you a certificate.
5. Import the new certificate into your local keystore:

```
keytool -importcert -alias tomcat -keystore <MY_KEYSTORE_FILENAME> -file
<MY_CERTIFICATE_FILENAME>
```

**Please note that some CAs require you to install an intermediate certificate before importing your certificate. Please refer to your CA documentation to successfully install your certificate.**

If you receive an error, and you use Verisign or GoDaddy, you may need to export the certificate to PKCS12 format along with the private key.

1. First, remove the certificate added above from the keystore:

```
keytool -delete -alias tomcat -keystore <MY_KEYSTORE_FILENAME>
```

2. Then export to PKCS12 format:

```
openssl pkcs12 -export -in <MY_CERTIFICATE_NAME> -inkey
<MY_PRIVATEKEY_NAME> -out <MY_PK12_KEYSTORE_NAME> -name tomcat
-CAfile
<MY_ROOTCERTIFICATE_NAME-alsoCalledBundleCertificateInGoDaddy>
-caname root
```

3. Then import from PKCS12 to jks:

```
keytool -importkeystore -deststorepass
<MY_DESTINATIONSTORE_PASSWORD> -destkeypass
<MY_DESTINATIONKEY_PASSWORD> -destkeystore <MY_KEYSTORE_FILENAME>
-srckeystore <MY_PK12_KEYSTORE_NAME> -srcstoretype PKCS12
-srcstorepass <MY_PK12_KEYSTORE_PASSWORD> -alias tomcat
```

## Step 2. Modify the Server Configuration File in your Confluence Installation

1. Edit the server configuration file at this location: `{CONFLUENCE-INSTALLATION}/conf/server.xml`.
2. Uncomment the following lines:

```
<Connector port="8443" maxHttpHeaderSize="8192"
 maxThreads="150" minSpareThreads="25"
maxSpareThreads="75"
 enableLookups="false" disableUploadTimeout="true"
 acceptCount="100" scheme="https" secure="true"
 clientAuth="false" sslProtocol="TLS" SSLEnabled="true"
 URIEncoding="UTF-8"
keystorePass="<MY_CERTIFICATE_PASSWORD>" />
```

3. Replace the text `<MY_CERTIFICATE_PASSWORD>` with the password you specified for your certificate.
4. Make sure that the attribute-value pair `SSLEnabled="true"` is part of the `Connector` element, as shown above. If this attribute is not present, attempts to access Confluence will time out.
5. Save the server configuration file.

### Step 3. Specify the Location of your Certificate

By default, Tomcat expects the keystore file to be named `.keystore` and to be located in the user home directory under which Tomcat is running (which may or may not be the same as your own home directory). This means that, by default, Tomcat will look for your SSL certificates in the following location:

- On Windows: `C:\Documents and Settings\%CURRENT_USER%\keystore`
- On OS X and UNIX-based systems: `~/keystore`

You may decide to move the certificate to a custom location. If your certificate is not in the default location, you will need to update your server configuration file as outlined below, so that Tomcat can find the certificate.

1. Edit the server configuration file at this location: `{CONFLUENCE-INSTALLATION}/conf/server.xml`
2. Add the attribute `keystoreFile="<MY_CERTIFICATE_LOCATION>"` to the `Connector` element, so that the element looks like this:

```
<Connector port="8443" maxHttpHeaderSize="8192"
 maxThreads="150" minSpareThreads="25"
maxSpareThreads="75"
 enableLookups="false" disableUploadTimeout="true"
 acceptCount="100" scheme="https" secure="true"
 clientAuth="false" sslProtocol="TLS" SSLEnabled="true"
 URIEncoding="UTF-8"
keystorePass="<MY_CERTIFICATE_PASSWORD>"
 keystoreFile="<MY_CERTIFICATE_LOCATION>" />
```

3. Replace the text `<MY_CERTIFICATE_LOCATION>` with the path to your certificate, including the path and the name of the `.keystore` file.
4. Save the server configuration file.

### Step 4. Change your Confluence Base URL to HTTPS

1. In your browser, go to the Confluence Administration Console.
2. Change the Server Base URL to HTTPS. See the documentation on [configuring the server base URL](#).
3. Restart Tomcat and access Confluence on `https://<MY_BASE_URL>:8443/`.

### Step 5. Add a Security Constraint to Redirect All URLs to HTTPS

Although HTTPS is now activated and available, the old HTTP URLs (<http://localhost:8090>) are still available.

Now you need to redirect the URLs to their HTTPS equivalent. You will do this by adding a security constraint in `web.xml`. This will cause Tomcat to redirect requests that come in on a non-SSL port.

1. Check whether your Confluence site uses the **RSS macro**. If your site has the RSS macro enabled, you may need to configure the URL redirection with a firewall rule, rather than by editing the `web.xml` file. Skip the steps below and follow the steps on the [RSS Feed Macro](#) page instead.
2. Otherwise, Edit the file at `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`.
3. Add the following declaration to the end of the file, **before** the `</web-app>` tag:

```
<security-constraint>
 <web-resource-collection>
 <web-resource-name>Restricted URLs</web-resource-name>
 <url-pattern>/</url-pattern>
 </web-resource-collection>
 <user-data-constraint>
 <transport-guarantee>CONFIDENTIAL</transport-guarantee>
 </user-data-constraint>
</security-constraint>
```

4. Restart Confluence and access <http://localhost:8090>. You should be redirected to <https://localhost:8443/login.action>.

**i** Confluence has two `web.xml` files. The other one is at `<CONFLUENCE_INSTALLATION>/conf/web.xml`. Please only add the security constraints to `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`, as described above.

#### Notes

- **Background information on generating a certificate:** The `'keytool -genkeypair'` command generates a key pair consisting of a public key and the associated private key, and stores them in a keystore. The command packages the public key into an X.509 v3 self-signed certificate, which is stored as a single-element certificate chain. This certificate chain and the private key are stored in a new keystore entry, identified by the `alias` that you specify in the command. The [Java SE documentation](#) has a good overview of the utility.
- **Custom SSL port:** If you have changed the port that the SSL connector is running on from the default value of 8443, you must update the `redirectPort` attribute of the standard HTTP connector to reflect the new SSL port. Tomcat needs this information to know which port to redirect to when an incoming request needs to be secure.
- **Multiple instances on the same host:** When running more than one instance on the same host, it is important to specify the `address` attribute in the `<CONFLUENCE_INSTALLATION>/conf/server.xml` file because by default the connector will listen on **all available network interfaces**, so specifying the address will prevent conflicts with connectors running on the same default port. See the Tomcat Connector documentation for more about setting the address attribute: <https://tomcat.apache.org/tomcat-8.0-doc/config/http.html>

```
<Connector port="8443" address="your.confluence.url.com"
maxHttpHeaderSize="8192"
 maxThreads="150" minSpareThreads="25"
maxSpareThreads="75"
 enableLookups="false" disableUploadTimeout="true"
 acceptCount="100" scheme="https" secure="true"
 clientAuth="false" sslProtocol="TLS" SSLEnabled="true"
 URIEncoding="UTF-8"
keystorePass="<MY_CERTIFICATE_PASSWORD>"
 keystoreFile="<MY_CERTIFICATE_LOCATION>" />
```

- **Protection for logins only or for individual spaces:** As of Confluence 3.0, Atlassian does not support HTTPS for logins only or for specific pages. We support only site-wide HTTPS. To see the

reasoning behind this decision, please see [CONF-18120](#) and [CONF-4116](#).

### Troubleshooting

- Check the Confluence knowledge base articles on [troubleshooting SSL](#)
- Problems with **Internet Explorer being unable to download attachments**: Applying SSL site wide can prevent IE from downloading attachments correctly. To fix this problem, edit `<CONFLUENCE_INSTANCE_HOME>/conf/server.xml` and add the following line within the `<Context ... />` element:

```
<Valve className="org.apache.catalina.authenticator.NonLoginAuthenticator"
 disableProxyCaching="true" securePagesWithPragma="false" />
```

### Connecting to LDAP or JIRA applications or Other Services via SSL

This page documents configuration of SSL, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with SSL. Please be aware that this material is provided for your information only, and that you use it at your own risk.

#### Related pages:

- [Configuring an SSL Connection to Active Directory](#)
- [Configuring Web Proxy Support for Confluence](#)
- [Running Confluence Over SSL or HTTPS](#)

This page describes how to get Confluence connecting to external servers over SSL, via the various SSL-wrapped protocols.

Here are some examples of when you may need to connect to an external server over SSL/HTTPS:

- You need to connect to an LDAP server, such as Active Directory, if the LDAP server is running over SSL.  
**For specific instructions for Active Directory, see [Configuring an SSL Connection to Active Directory](#).**
- You want to set up your JIRA application as a trusted application in Confluence, when JIRA is running over SSL.
- You want to refer to an `https://...` URL in a Confluence macro.

If you want to run Confluence itself over SSL, see [Running Confluence Over SSL or HTTPS](#).

There's a [Confluence SSL plugin](#) that facilitates this process.

### Importing SSL Certificates

The following commands apply to JDK 1.5. For commands/syntax relevant to JDK 1.6, please refer to [this document from Oracle](#).

1. Add the root certificate to your default Java keystore with the following command. This is the certificate that was used to authorise the LDAP server's certificate. It will be either the one that was used for signing it, or will come from further up in the trust chain, possibly the root certificate. This is often a self-signed certificate, when both ends of the SSL connection are within the same network. Again, the exact alias is not important.

```
keytool -import -alias serverCert -file RootCert.crt -keystore
%JAVA_HOME%/jre/lib/security/cacerts (Windows)
keytool -import -alias serverCert -file RootCert.crt -keystore
$JAVA_HOME/jre/lib/security/cacerts (Linux/Unix/Mac)
```

2. Import your LDAP or JIRA server's public certificate into the JVM Keystore. This is the certificate that the LDAP server will use to set up the SSL encryption. You can use any alias of your choosing in place of "JIRAorLDAPServer.crt".

```
keytool -import -alias ldapCert -file JIRAorLDAPServer.crt
-keystore %JAVA_HOME%/jre/lib/security/cacerts (Windows)
keytool -import -alias ldapCert -file JIRAorLDAPServer.crt
-keystore $JAVA_HOME/jre/lib/security/cacerts (Linux/Unix/Mac)
```

3. Verify that the certificate has been added successfully by entering the following command:

```
keytool -list -keystore %JAVA_HOME%/jre/lib/security/cacerts
(Windows)
keytool -list -keystore $JAVA_HOME/jre/lib/security/cacerts
(Unix/Linux)
keytool -list -keystore /Library/Java/Home/lib/security/cacerts
(Mac)
```

4. Ensure that you have updated JAVA\_OPTS to specify the path to the keystore, as specified in [Connecting to SSL services](#) before restarting Tomcat/Confluence.  
There is no need to specify an alias for Confluence to use. On connecting to the LDAP server, it will search through the keystore to find a certificate to match the key being presented by the server.

### Troubleshooting

Check the following knowledge base articles:

- [Unable to Connect to SSL Services due to PKIX Path Building Failed](#)
- [SSL troubleshooting articles](#)

### Configuring RSS Feeds

A Confluence System Administrator can configure the following aspects of RSS feeds:

- The maximum number of items that Confluence returns to an RSS feed request.
- The maximum time period that Confluence allows to respond to an RSS feed request.

Both of these are set in the 'Edit Security Configuration' screen.

#### To configure RSS feeds:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Security Configuration**.
3. Choose **Edit**.
4. Enter a value for **Maximum RSS Items**. The default value is 200.
5. Enter a value for **RSS timeout**.
6. Choose **Save**.

*Screenshot: Configuring RSS feeds*

#### On this page:

- [Notes](#)

#### Related pages:

- [The RSS Feed Builder](#)

Maximum RSS Items	<input type="text" value="200"/>
Limit the maximum number of items an RSS Feed can request.	
RSS timeout	<input type="text" value="60"/> seconds
The time in seconds allowed to create each RSS Feed. Any items rendered within the timeout will still be returned.	

#### Notes

- When using the RSS [Feed Builder](#), a user could potentially enter such a large value for the number of feed items returned that Confluence would eventually run out of memory.
- When using the [Feed Builder](#), if a users a value greater than this setting (or less than 0) they will get a validation error.
- If any pre-existing feeds are set to request more than the configured maximum, they will be supplied with only the configured maximum number of items. This is done silently - there is no logging and no message is returned to the RSS reader.
- If Confluence times out when responding to an RSS feed request, any items already rendered are returned.

## Preventing and Cleaning Up Spam

If your Confluence site is public-facing you may be affected by spammers.

### Stopping Spammers

#### To prevent spammers:

1. Enable Captcha. See [Configuring Captcha for Spam Prevention](#).
2. Run Confluence behind an [Apache webserver](#) and create rules to block the spammer's IP address.

### Blocking Spam at Apache or System Level

If a spam bot is attacking your Confluence site, they are probably coming from one IP address or a small range of IP addresses. To find the attacker's IP address, follow the [Apache access logs](#) in real time and filter for a page that they are attacking.

For example, if the spammers are creating users, you can look for `signup.action`:

```
$ tail -f confluence.atlassian.com.log | grep signup.action
1.2.3.4 - - [13/Jan/2010:00:14:51 -0600] "GET /signup.action HTTP/1.1"
200 9956 "-" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1)"
37750
```

Compare the actual spam users being created with the log entries to make sure you do not block legitimate users. By default, Apache logs the client's IP address in the first field of the log line.

Once you have the offender's IP address or IP range, you can add it to your firewall's blacklist. For example, using the popular [Shorewall](#) firewall for Linux you can simply do this:

```
echo "1.2.3.4" >> /etc/shorewall/blacklist
/etc/init.d/shorewall reload
```

To block an IP address at the Apache level, add this line to your Apache vhost config:

```
Deny from 1.2.3.4
```

You can restart Apache with a "graceful" command which will apply the changes without dropping any current sessions.

If this still does not stop the spam, then consider turning off [public signup](#).

### Deleting Spam

#### Profile Spam

By 'profile spam', we mean spammers who create accounts on Confluence and post links to their profile page.

If you have had many such spam profiles created, the easiest way to delete them is via SQL.

#### To delete a spam profile:

1. Shut down Confluence and back up your database.  
**Note:** This step is essential before you run any SQL commands on your database.
2. Find the last real profile:

```
SELECT bodycontentid,body FROM bodycontent WHERE contentid IN
(SELECT contentid FROM content WHERE contenttype='USERINFO')
ORDER BY bodycontentid DESC;
```

3. Look through the bodies of the profile pages until you find where the spammer starts. You may have to identify an number of ranges.
4. Find the killset:

```
CREATE TEMP TABLE killset AS SELECT bc.bodycontentid,c.contentid,c.username
FROM
 bodycontent bc JOIN content c ON bc.contentid=c.contentid WHERE
 bodycontentid >= BOTTOM_OF_SPAM_RANGE AND bodycontentid <= TOP_OF_SPAM_RANGE
 AND c.contenttype='USERINFO';

DELETE FROM bodycontent WHERE bodycontentid IN (SELECT bodycontentid FROM
killset);

DELETE FROM links WHERE contentid IN (SELECT contentid FROM killset);

DELETE FROM content WHERE prevver IN (SELECT contentid FROM killset);

DELETE FROM content WHERE pageid IN (SELECT contentid FROM killset);

DELETE FROM content WHERE contentid IN (SELECT contentid FROM killset);

DELETE FROM os_user_group WHERE user_id IN (SELECT id FROM killset k JOIN
os_user o ON o.username=k.username);

DELETE FROM os_user WHERE username IN (SELECT username FROM killset);
```

If you're using Confluence 5.6 or earlier use the SQL commands below:

▼ [For Confluence 5.6 and earlier...](#)

```

CREATE TEMP TABLE killset AS SELECT
bc.bodycontentid,c.contentid,c.username FROM
 bodycontent bc JOIN content c ON bc.contentid=c.contentid WHERE
 bodycontentid >= BOTTOM_OF_SPAM_RANGE AND bodycontentid <=
TOP_OF_SPAM_RANGE
 AND c.contenttype='USERINFO';

DELETE FROM bodycontent WHERE bodycontentid IN (SELECT bodycontentid FROM
killset);

DELETE FROM links WHERE contentid IN (SELECT contentid FROM killset);

DELETE FROM content WHERE prevver IN (SELECT contentid FROM killset);

DELETE FROM attachments WHERE pageid IN (SELECT contentid FROM killset);

DELETE FROM content WHERE contentid IN (SELECT contentid FROM killset);

DELETE FROM os_user_group WHERE user_id IN (SELECT id FROM killset k JOIN
os_user o ON o.username=k.username);

DELETE FROM os_user WHERE username IN (SELECT username FROM killset);


```

5. Once the spam has been deleted, restart Confluence and [rebuild the index](#). This will remove any references to the spam from the search index.

## Scheduled Jobs

The administration console allows you to schedule various administrative jobs in Confluence, so that they are executed at regular time intervals. The types of jobs which can be scheduled cover:

- Confluence site backups
- Storage optimisation jobs to clear Confluence's temporary files and caches
- Index optimisation jobs to ensure Confluence's search index is up to date
- Mail queue optimisation jobs to ensure Confluence's mail queue is maintained and notifications have been sent.

 You'll need [System Administrator](#) permissions in order to edit and manually run jobs.

### On this page:

- [Accessing Confluence's Scheduled Jobs Configuration](#)
- [Running a job manually](#)
- [Changing a job's schedule](#)
- [Disabling/Re-enabling a Job](#)
- [Viewing a Job's Execution History](#)
- [Types of Jobs](#)
- [Cron Expressions](#)

### Related pages:

- [Trigger Module](#)
- [Configuring Backups](#)

## Accessing Confluence's Scheduled Jobs Configuration

To access Confluence's Scheduled Jobs configuration page:



1. > **General Configuration > Scheduled Jobs**

2. All scheduled jobs are listed with:

- **Status** - the job's status, which is either 'Scheduled' (if it is currently enabled) or 'Disabled'.
- **Last Execution** - the date and time when the job was last executed. This field will be empty of the job was never executed.
- **Next Execution** - the date and time when the job is next scheduled to be executed. This field will contain dash symbol ('-') if the job is disabled.
- **Avg. Duration** - the length of time (in milliseconds) that it took to complete the job (the last time it ran).
- **Actions** - Options to edit the job's schedule, run it manually, view the history or disable the job.

*Screenshot: Scheduled Jobs*

Scheduled Jobs					
Job	Status	Last Execution	Next Execution	Avg. Duration	Actions
Back Up Confluence	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 02:00	17450	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Edit</a> · <a href="#">Disable</a>
Check Cluster Safety	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:39	16	<a href="#">History</a> · <a href="#">Run</a>
Clean Journal Entries	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 02:00	882	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Edit</a> · <a href="#">Disable</a>
Clean Temporary Directory	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 04:00	1107	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Edit</a>
Clear Expired Mail Errors	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 03:00	653	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Edit</a>
Clear Expired Remember Me Tokens	Scheduled	Feb 20, 2015 08:17	Mar 20, 2015 00:00	300	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Edit</a>
Cluster Cache Compaction	Scheduled	Mar 11, 2015 08:10	Mar 11, 2015 12:00	0	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Disable</a>
EhCache Cache Compaction	Scheduled	Mar 11, 2015 11:34	Mar 11, 2015 11:39	10	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Disable</a>
Email Daily Reports	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 00:00	6952	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Edit</a> · <a href="#">Disable</a>
Flush Edge Index Queue	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:38	37	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Edit</a> · <a href="#">Disable</a>
Flush Index Queue	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:38	20	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Edit</a>
Flush Local Task Queue	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:39	65	<a href="#">History</a>
Flush Mail Queue	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:39	89	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Edit</a> · <a href="#">Disable</a>
Flush Task Queue	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:39	47	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Disable</a>
JIRA Metadata Cache Configuration	Scheduled	Mar 11, 2015 11:19	Mar 11, 2015 12:19	37	<a href="#">History</a> · <a href="#">Run</a>
scheduledjob.desc.conversionQueueMonitor	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:38	1	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Edit</a> · <a href="#">Disable</a>
scheduledjob.desc.purgeHistoryJob	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 00:00	1444	<a href="#">History</a> · <a href="#">Edit</a>
Send Recommended Updates Email	Scheduled	Mar 11, 2015 08:10	Mar 11, 2015 12:00	857	<a href="#">History</a> · <a href="#">Run</a> · <a href="#">Disable</a>

## Running a job manually

To run a job manually head to the Scheduled Jobs list and choose **Run** next to the job. It will run immediately.

Not all jobs can be run manually.

## Changing a job's schedule

To change a job's schedule:

1. Choose **Edit** next to the job you want to change.
2. Enter the new day or time to run the job as a cron expression - there's more info about cron expressions [below](#).
3. **Save** your changes to the job's schedule, or **Revert** back to the default setting.

Not all jobs' schedules are configurable.

*Screenshot: Configuring a Job Schedule*

### Edit Schedule for Clean Index Queue

#### Advanced Configuration

Cron Expression

Define schedule using a [cron expression](#).

#### Future Schedules:

20-Feb-2013	02:00:00
21-Feb-2013	02:00:00
22-Feb-2013	02:00:00
23-Feb-2013	02:00:00
24-Feb-2013	02:00:00
25-Feb-2013	02:00:00
26-Feb-2013	02:00:00
27-Feb-2013	02:00:00
28-Feb-2013	02:00:00
01-Mar-2013	02:00:00

Default Save Cancel

### Disabling/Re-enabling a Job

By default, all jobs in Confluence are enabled.

Use the **Disable** / **Enable** links in the action column to disable and re-enable each job.

Not all jobs in Confluence can be disabled.

### Viewing a Job's Execution History

To see when a job was last run, and how long the job took to run, click the **History** link beside the job.

If a job has not run at least once the History link won't appear.

*Screenshot: Job Execution History*

### History for Clean Temporary Directory

Started	Ended	Duration (ms)
19-Feb-2013 04:00:00	19-Feb-2013 04:00:01	1901
18-Feb-2013 04:00:00	18-Feb-2013 04:00:00	42
17-Feb-2013 04:00:00	17-Feb-2013 04:00:00	334

### Types of Jobs

Job Name	Description	Execution Behaviour	Default Schedule
Back Up Confluence	Performs a <a href="#">backup</a> of your entire Confluence site.	Per cluster	At 2am every day
Check Cluster Safety	For clustered Confluence installations, this job ensures that only one Confluence instance in the cluster writes to the database at a time. For standard (non-clustered) editions of Confluence, this job is useful for alerting customers who have accidentally connected a second Confluence instance to a Confluence database which is already in use.	Per cluster	Every 30 seconds
Clean Journal Entries	Periodically clears journal entries that have already been processed to ensure that its size does not grow indefinitely.	Per node	At 2am every day
Clean Temporary Directory	Cleans up temporary files generated in the <code>&lt;confluence-home&gt;/temp</code> directory. This temp directory is created by exports etc. This doesn't include the temp directory located in the confluence install directory.	Per node	At 4am every day
Clear Expired Mail Errors	Clears notification errors in the <a href="#">mail error queue</a> . A notification error is sent to the mail error queue whenever the notification fails to be sent due to an error.	Per cluster	At 3am every day
Clear Expired Remember Me Tokens	Clears all expired 'Remember Me' tokens from the Confluence site. Remember Me tokens expire after two weeks.	Per cluster	On the 20th of each month
Email Daily Reports	Emails a daily summary report of all Confluence changes <a href="#">to all subscribers</a> . Since each email report only records changes from the last 24-hour period, it is recommended that you only change the time of this job whilst keeping the job's frequency to 24 hours.	Per cluster	At 12am every day
Flush Edge Index Queue	Flushes the Edge Index Queue so Confluence's search results stay up to date.	Per node	Every 30 seconds
Flush Index Queue	Flushes changes to Confluence's index so that Confluence's search results are up to date. Confluence records each content update in its search index.	Per node	Every 5 seconds
Flush Local Task Queue	Flushes the local task queue. (These are internal Confluence tasks that are typically flushed at a high frequency.)	Per node	Every minute
Flush Mail Queue	Sends notifications that have been queued up in the <a href="#">mail queue</a> .	Per cluster	Every minute
Flush Task Queue	Flushes the task queue. (These are internal Confluence tasks that are typically flushed at a high frequency.)	Per node	Every minute
Send Recommended Updates Email	Triggers sending recommended update emails to users. The job runs hourly, but users will receive the notification weekly or daily, depending on the setting in their profile, at a time that matches their timezone.	Per cluster	Hourly

## Cron Expressions

A cron expression is a string of 6-7 'time interval' fields that defines the frequency with which a job is executed. Each of these fields can be expressed as either a numerical value or a special character and each field is separated by at least one space or tab character.

The table below shows the order of time interval fields in a cron expression and each field's permitted numerical values.

You can specify a special character instead of a numerical value for any field in the cron expression to provide flexibility in defining a job's frequency. Common special characters include:

- '\*' — a 'wild card' that indicates 'all permitted values'.
- '?' — indicates 'ignore this time interval' in the cron expression. That is, the cron expression will not be bound by the time interval (such as 'Month', 'Day of week' or 'Year') to which this character is specified.

For more information about cron expressions, please refer to the [Cron Trigger tutorial on the Quartz website](#).

Order in cron expression	Time interval field	Permitted values*	Required?
1	Seconds	0-59	Yes
2	Minutes	0-59	Yes
3	Hours	0-23	Yes
4	Day of month	1-31	Yes
5	Month	1-12 or JAN-DEC	Yes
6	Day of week	1-7 or SUN-SAT	Yes
7	Year	1970-2099	No

\* Excluding special characters.

## Configuring the Whitelist

Confluence administrators can choose to allow incoming and outgoing connections and content from specified sources for use in the RSS macro, HTML Include macro and gadgets, by adding URLs to the whitelist.

Confluence will display an error if content has been added that is not from an allowed source, and prompt the user to add the URL to the whitelist.

[Application links](#) are automatically added to the whitelist. You don't need to manually add them.

Note: The [HTML Include macro](#) is disabled by default.

### Add allowed URLs to the whitelist

To add a URL to the whitelist:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Whitelist**.
3. Enter the URL or expression you want to allow.
4. Choose the **Type** of expression (see [below](#) for examples of the types available).
5. Choose **Allow Incoming** if you need to allow CORS requests (see [below](#)).
6. Choose **Add**.

Your URL or expression appears in the whitelist.

To test that your whitelisted URL is working as expected you can enter a URL in the **Test a URL** field. Icons will

indicate whether incoming and / or outgoing traffic is allowed for that URL.

## Expression Types

When adding a URL to the whitelist, you can choose from a number of expression types.

Type	Description	Example
Domain name	Allows all URLs from the specified domain.	<code>http://www.example.com</code>
Exact match	Allows only the specified URL.	<code>http://www.example.com/thispage</code>
Wildcard Expression	Allows all matching URLs. Use the wildcard * character to replace one or more characters.	<code>http://*example.com</code>
Regular Expression	Allows all URLs matching the regular expression.	<code>http(s)?://www\.example\.com</code>

## Allow Incoming

**Allow Incoming** enables [CORS](#) requests from the specified origin. The URL must match the format `scheme://host[:port]`, with no trailing slashes (`:port` is optional). So `http://example.com/` would not allow CORS requests from the domain `example.com`.

## Disabling the whitelist

The whitelist is enabled by default. You can choose to disable the whitelist however this will allow all URLs, including malicious content, and is not recommended.

To disable the whitelist:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose **Whitelist**.
3. Choose **Turn off whitelist**.
4. Choose **Confirm**.

All URLs will now be allowed. We strongly recommend **not** disabling the whitelist.

## Configuring the Time Interval at which Drafts are Saved

By default, Confluence saves a draft of your page once every thirty seconds. Confluence administrators can configure how often drafts are saved.

As a Confluence administrator, you can set the time interval at which drafts are saved as follows:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Click **Further Configuration** in the left-hand panel.
3. Edit the setting for **Draft Save Interval**.

### Related pages:

- [Drafts](#)

## Performance Tuning

This document describes tuning your application for improved performance. It is not a guide to troubleshooting Confluence outages. Check [Trouble](#)

[shooting Confluence hanging or crashing](#) for help if Confluence is crashing.

Like any server application, Confluence may require some tuning as it is put under heavier use. We do our best to make sure Confluence performs well under a wide variety of circumstances, but there's no single configuration that is best for everyone's environment and usage patterns.

If you are having problems with the performance of Confluence and need our help resolving them, you should read [Requesting Performance Support](#).

### Use the latest version of your tools

Use the latest versions of your application servers and Java runtime environments. Newer versions are usually better optimized for performance.

#### On this page:

- [Use the latest version of your tools](#)
- [Avoid swapping due to not enough RAM](#)
- [Being aware of other systems using the same infrastructure](#)
- [Choice of database](#)
- [Database connection pool](#)
- [Database in general](#)
- [Database statistics and query analysers](#)
- [Cache tuning in Confluence and Apache](#)
- [Antivirus software](#)
- [Enabling HTTP compression](#)
- [Performance testing](#)
- [Access logs](#)
- [Built-in profiler](#)
- [Application server memory settings](#)
- [Web server configuration](#)
- [Troubleshooting possible memory leaks](#)

### Avoid swapping due to not enough RAM

Always watch the swapping activity of your server. If there is not enough RAM available, your server may start swapping out some of Confluence's heap data to your hard disk. This will slow down the JVM's garbage collection considerably and affect Confluence's performance. In clustered installations, swapping can lead to a [Cluster Panic due to Performance Problems](#). This is because swapping causes the JVM to pause during [Garbage Collection](#), which in turn can break the inter-node communication required to keep the clustered nodes in sync.

### Being aware of other systems using the same infrastructure

It may sound tempting: Just have one powerful server hosting your database and/or application server, and run **all** your crucial programs on that server. **If** the system is set up perfectly, then you might be fine. Chances are however that you are missing something, and then one application's bug might start affecting other applications. So if Confluence is slow every day around noon, then maybe this is because another application is using the shared database to generate complicated reports at that time? Either make sure applications can't harm each other despite sharing the same infrastructure, or get these systems untangled, for example by moving them to separate instances that can be controlled better.

### Choice of database

The **embedded H2 database** is provided for evaluating Confluence, not for production Confluence sites. After the evaluation finishes, you must switch to a [supported external database](#). We recommend using what you are familiar with, because your ability to maintain the database will probably make far more difference to what you get out of it than the choice of database itself.

### Database connection pool

If load on Confluence is high, you may need more simultaneous connections to the database.

- If you are using JNDI data-sources, you will do this in your application server's configuration files.
- If you have configured Confluence to access the database directly, you will need to manually edit the `hibernate.c3p0.max_size` property in the `confluence.cfg.xml` file in your `confluence.home` directory. After you have changed the URL in this file, restart Confluence.

To assess whether you need to tune your database connection pool, take [thread dumps](#) during different times (including peak usage). Inspect how many threads have concurrent database connections.

### Database in general

If Confluence is running slowly, one of the most likely cause is that there is some kind of bottleneck in (or around) the database.

The first item you should check is the "**Database Latency**" field in the System Information tab in the admin console.

<b>Database Connection</b>	Transaction isolation	Read committed
<b>Database Latency</b>		0
<b>Confluence Usage</b>		

The latency is calculated by sending a trivial request to the database, querying a table which is known to have only one column and one row. ("select \* from CLUSTERSAFETY"). Obviously this query should be blazing fast, and return within 1 or 2 milliseconds. If the value displayed is between 3 and 5 milliseconds, you might already have an issue. If the value is above 10ms, then you **definitely** need to investigate and improve something! A few milliseconds may not sound so bad, but consider that Confluence sends quite a few database queries per page request, and those queries are a lot more complex too! High latency might stem from all sorts of problems (slow network, slow database, connection-pool contention, etc), so it's up to you to investigate. Don't stop improving until latency is below 2ms on average.

Obviously, latency is just the very first thing to look at. You may get zero latency and still have massive database problems, e.g. if your tables are poorly indexed. **So don't let a low latency fool you either.**

### Database statistics and query analysers

Modern databases have query optimisers based on collecting statistics on the current data. Using the SQL EXPLAIN statement will provide you information on how well the query optimiser is performing. If the cost estimate is wildly inaccurate then you will need to run statistics collection on the database. The exact command will depend on your database and version. In most cases you can run statistics collection while Confluence is running, but due to the increased load on the database it's best to do this after normal hours or on a week-end.

### Cache tuning in Confluence and Apache

To reduce the load on the database, and speed up many operations, Confluence keeps its own cache of data. Tuning the size of this cache may speed up Confluence (if the caches are too small), or reduce memory (if the caches are too big).

Please have a look at our documentation on [Cache Performance Tuning](#) for information on how to tune Confluence caches.

To improve performance of a large Confluence site, we recommend that you move the caching of static content from the JVM into Apache. This will prevent the JVM from having a number of long running threads serving up static content. See [Configuring Apache to Cache Static Content via mod\\_disk\\_cache](#).

### Antivirus software

Antivirus software greatly decreases the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental, and may even cause errors with Confluence. You should configure your antivirus software to ignore the Confluence home directory, its index directory and any database-related directories.

### Enabling HTTP compression

If bandwidth is responsible for bottlenecking in your Confluence installation, you should consider [enabling HTTP compression](#). This may also be useful when running an external facing instance to reduce your

bandwidth costs.

Take note of the [known issues with HTTP compression](#) in versions of Confluence prior to 2.8, which may result in high memory consumption.

## Performance testing

You should try out all configuration changes on a demo system. Ideally, you should run and customize loadtests that simulate user behaviour.

## Access logs

You can find out which pages are slow and which users are accessing them by [enabling Confluence's built-in access logging](#).

## Built-in profiler

You can identify the cause of page delays using Confluence's built-in profiler according to [Troubleshooting Slow Performance Using Page Request Profiling](#).

## Application server memory settings

See [How to fix out of memory errors by increasing available memory](#).

## Web server configuration

For high-load environments, performance can be improved by using a web server such as Apache in front of the application server. There is a configuration guide to [Running Confluence behind Apache](#).

When configuring your new web server, make sure you configure sufficient threads/processes to handle the load. This applies to both the web server and the application server connector, which are typically configured separately. If possible, you should enable connection pooling in your web server connections to the application server.

## Troubleshooting possible memory leaks

Some external plugins, usually ones that have been written a long time ago and that are not actively maintained anymore, have been reported to consume memory and never return it. Ultimately this can lead to a crash, but first this manifests as reduced performance. The [Troubleshooting Confluence hanging or crashing](#) guide is a good place to start. Some of the known causes listed there could result in performance issues short of a crash or hang.

## Cache Performance Tuning

Confluence performance can be significantly affected by the performance of its caches. It is essential for the administrator of a large production installation of Confluence to tune the caches to suit its environment.

There are several configurable parameters for each of the cache regions, most notably cache size, cache expiry delay and eviction policy. In most cases, cache size is the parameter you would want to change.

To change the size of a cache:

1. Go to



> **General Configuration > Cache Management.**

### On this page:

- [Cache tuning example](#)
- [Finding the configuration file](#)
- [Cache key mappings](#)
- [Caching in Confluence Data Center](#)
- [Important caches](#)
- [Monitoring the contents of a cache](#)
- [Notes](#)

2. Choose **Show Advanced View**.
3. Choose **Adjust Size** next to the cache you want to change.

To modify other parameters you can modify the cache configuration files manually.

### Cache tuning example

As an example of how to tune Confluence's caches, let's have a look at the following table:

Caches	% Used	% Effectiveness	Objects/Size	Hit/Miss/Expiry
Attachments	87%	29%	874/1000	78226/189715/187530
Content Attachments	29%	9%	292/1000	4289/41012/20569
Content Bodies	98%	81%	987/1000	28717/6671/5522
Content Label Mappings	29%	20%	294/1000	4693/18185/9150
Database Queries	96%	54%	968/1000	105949/86889/83334
Object Properties	27%	18%	279/1000	5746/25386/8102
Page Comments	26%	11%	261/1000	2304/17178/8606
Users	98%	5%	982/1000	6561/115330/114279


The maximum size of the caches above is 1000 (meaning that it can contain up to 1000 objects). You can tell when a cache size needs to be increased because the cache has both:

- a high usage percentage (above 75%)
- a low effectiveness percentage.

Check the 'effectiveness' versus the 'percent used'. A cache with a low percent used need not have its size lowered; it does not use more memory until the cache is filled.

Based on this, the sizes of the "Attachments", "Database Queries", and "Users" caches should be increased to improve their effectiveness.

As the stored information gets older or unused it will expire and be eliminated from the cache. Cache expiry may be based on time or on frequency of use.

 There is not much that you can do with a cache that has both a low percentage of usage and effectiveness. Over time, as the cache is populated with more objects and repeat requests for them are made, the cache's effectiveness will increase.

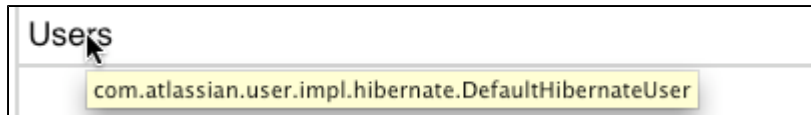
### Finding the configuration file

Cache configurations are stored in `<confluence-home>/shared-home/config/cache-settings-overrides.properties`

For Confluence Data Center (clustered) it can be found in `<confluence-shared-home>/config/cache-settings-overrides.properties` (in the shared home directory for the cluster).

### Cache key mappings

The cache configuration file configures caches by their keys. To find out a cache key hover your mouse over the cache name in the **Cache Management** screen.



## Caching in Confluence Data Center

In Confluence Data Center (clustered) you have a distributed cache and a cluster node-local cache. The Cluster Management page will indicate cluster distributed cache and cluster node-local cache.

The cache configuration file is stored in the shared home directory for the cluster.

## Important caches

The following suggestions are general guidelines. In cases of large databases, 20-30% of the size of the table may be unnecessarily large. Check the effectiveness and percent used categories in the cache for more specific assessments.

- **Content Objects cache** (`com.atlassian.confluence.core.ContentEntityObject`) should be set to at least 20-30% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the query `select count(*) from CONTENT where prevver is null`.
- **Content Body Mappings cache** (`com.atlassian.confluence.core.ContentEntityObject.bodyContents`) should be set to at least 20% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the query `select count(*) from CONTENT where prevver is null`.
- **Embedded Crowd Internal User cache** (`com.atlassian.crowd.model.user.InternalUser`) should be set to the number of users you have in the internal directory. You can discover this number by using the following SQL:

```
SELECT
 COUNT(*)
FROM
 cwd_user u
JOIN
 cwd_directory d
ON
 u.directory_id = d.id
AND d.directory_name = 'Confluence Internal Directory';
```

- **Embedded Crowd Users cache** (`com.atlassian.confluence.user.crowd.CachedCrowdUserDao.USER_CACHE`) should be set to the number of rows in the `cwd_user` table.

```
SELECT
 COUNT(*)
FROM
 cwd_user u;
```

- **Space permissions by ID cache** (`com.atlassian.confluence.security.SpacePermission`) should be set to the number of space permissions in your deployment (a good rule of thumb is 20 times the number of spaces). You can find the number of space permissions using the query `select count(*) from SPACEPERMISSIONS`.

## Monitoring the contents of a cache

To monitor what is in a cache:

1. In your browser go to `<confluence-URL>/admin/cachecontents.jsp`  
All caches that contain items will appear.
2. Select a cache from the list.

### Notes

- To improve performance of a large Confluence site, we recommend that you move the caching of static content from the JVM into Apache. This will prevent the JVM from having a number of long running threads serving up static content. See [Configuring Apache to Cache Static Content via mod\\_disk\\_cache](#).

## Cache Statistics

Confluence provides statistics about its internal caches that allow you to track the size and hit ratio of each cache and tune it for better performance (if necessary). See [Performance Tuning](#) for more information.

### Configurable Caches

System administrators can change the sizes of Confluence's internal caches through the Administration Console and these changes will take effect without the need to first shut down and then restart Confluence. The maximum number of units for any of the defined cache regions can be adjusted individually.

Note that larger cache sizes will require more memory at runtime, so you should review the memory allocation of the Confluence Java process and the physical memory available on your server.

### Viewing Cache Statistics and Modifying Cache Sizes

To view the cache statistics:

1. Go to  **> General Configuration > Cache Management.**
2. Choose **Show Advanced View.**

Here is an example for one of the most frequently used caches, the 'Content Object' cache.

Cache Statistics							
This page displays statistics about Confluence's internal caches. You may wish to tune cache sizes and expiration policies to improve memory usage or performance.							
Cache Name	Capacity Utilisation	Effectiveness	Current / Max Entries	Current Heap Size (MB)	Hit / Miss / Evicted	Adjust Max Entries	Flush Entries
Content Objects	36%	73%	3653 / 10000	Unknown	307748 / 112538 / 37051	Adjust Size	Flush

About the generated numbers:

<b>Capacity Utilisation</b>	$\text{=(Objects)/(Size)}$ For example Percent Used = $4023 / 5000 = 80\%$
<b>Effectiveness :</b>	$\text{=(Hits)/(Hits + Misses)}$ For example Effectiveness = $374550 / (374550 + 140460) = 73\%$
<b>Current / Max Entries</b>	The number of entries in the cache / the number of total possible entries allowed (this is the size of the cache).

<b>Current Heap Size</b>	Heap memory (in MB) allocated to this cache (if applicable)
<b>Hit / Miss / Evicted</b>	The number of reads accessing cache where required content was found / the number of reads accessing cache where required content was not found / the number of objects evicted from the cache.
<b>Adjust Size</b>	Use this option to specify a different maximum cache size.
<b>Flush</b>	Flushes the cache.

Changes to cache size configurations are saved in the `config/cache-settings-overrides.properties` file in your home directory (or shared home directory in a Confluence Data Center cluster).

## Memory Usage and Requirements

Managing Confluence's performance and memory usage really depends on what resources are available. Confluence will run faster if you give it lots of memory for its caches, but it should still be able to run quite well in low-memory environments, with the right tuning. Below are some tips on getting the most out of your Confluence site.

### Increasing the amount of memory available to Confluence

See [Increasing JIRA Memory](#) for details on how to increase the memory available to web application servers typically used to run Confluence.

### Embedded database

The embedded HSQL database that comes with Confluence essentially holds all your data in memory while the Confluence server is running. If you are running out of memory, you should consider [migrating Confluence to an external database](#).

### Caching

By default, Confluence keeps large in-memory caches of data to improve its responsiveness and the user experience. The trade off is an increase in memory requirements to support the cache. Administrators of larger Confluence sites may need to configure the size of their caches to improve performance.

To customise Confluence's cache to meet your needs, see [cache tuning](#).

To increase the amount of memory available to Confluence, see [How to fix out of memory errors by increasing available memory](#).

### Mail error queue

Confluence keeps a copy of all emails that it failed to send within an internal error queue. In the event of intermittent failures such as network connectivity issues, the emails in this queue can be manually resent when the problem is fixed. Under certain circumstances, the mail queue can fill up with large objects. The queue is regularly flushed, but if you get a *lot* of mail errors, you might get a spike in memory usage.

### Attachments

The indexing of large attachments requires that the attachment be loaded into memory. In the case of large attachments, this can cause a temporary strain on the systems resources, and may result in indexing failing

#### On this page:

- [Increasing the amount of memory available to Confluence](#)
- [Embedded database](#)
- [Caching](#)
- [Mail error queue](#)
- [Attachments](#)
- [System backup and restore](#)
- [Known issues that we do not have control over](#)
- [Confluence is taking long periods of time to respond to some actions](#)

#### Related pages:

- [Performance Tuning](#)
- [Requesting Performance Support](#)

because the attachment could not be fully loaded into memory.

## System backup and restore

The Confluence backup and restore process scales linearly with the size of data. This can have a significant impact on large Confluence instances where the amount of data exceeds the amount of available memory. If you are experiencing an `OutOfMemoryError` during either a backup or restore processes, then we strongly recommend that you choose and [Production Backup Strategy](#).

If you encounter an `OutOfMemoryError` while restoring a backup and wish to overcome this issue by increasing memory, how much more will you need to make this process work? A good rule of thumb is to have a look at the size of the `entities.xml` file in your backup. This file contains all of the data Confluence will be loading, so at least that much is required. Add another 64-128Mb to ensure that Confluence has enough memory to load and function and that should be enough. To increase the amount of memory available to Confluence, see [How to fix out of memory errors by increasing available memory](#).

## Known issues that we do not have control over

There are also some memory issues we don't have any control over. For example,

- There's a memory leak in the Oracle 10g JDBC drivers. Not much we can do about that.
- One customer found a rather nasty memory leak that appeared to originate inside Tomcat 5, but only using the IBM JDK on PowerPC.

If you are having problems that appear to result from a memory leak, log an issue on <http://support.atlassian.com>. Our memory profiler of choice is [YourKit](#). It would be helpful to us if you can provide us with a memory dump from that tool showing the leak.

## Confluence is taking long periods of time to respond to some actions

A common cause of random pauses in Confluence is the JVM running garbage collection. To determine if this is what is happening, enable verbose garbage collection and look at how long Java is taking to free up memory. If the random pauses match when Java is running its garbage collection, garbage collection is the cause of the pause.

Verbose garbage collection will generate log statements that indicate when Java is collecting garbage, how long it takes, and how much memory has been freed.

To enable gc (garbage collection) logging, start Confluence with the option `-XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc -Xloggc:gc.log`. Replace `gc.log` with an absolute path to a `gc.log` file.

For example, with a Windows service, run:

```
tomcat5 //US//Confluence ++JvmOptions="-XX:+PrintGCDetails
-XX:+PrintGCTimeStamps -verbose:gc -Xloggc:c:\confluence\logs\gc.log"
```

or in `bin/setenv.sh`, set:

```
export CATALINA_OPTS="$CATALINA_OPTS -XX:+PrintGCDetails
-XX:+PrintGCTimeStamps -verbose:gc
-Xloggc:${CATALINA_BASE}/logs/gc.log"
```

If you modify `bin/setenv.sh`, you will need to restart Confluence for the changes to take effect.

What can you do to minimise the time taken to handle the garbage collection? See <http://java.sun.com/docs/hotspot/gc1.4.2/> for details on tuning the JVM to minimise the impact that garbage collection has on the running application.

## Requesting Performance Support

### Basic performance troubleshooting steps

Begin with the following procedures:

1. Go through the [Troubleshooting Confluence hanging or crashing](#) page to identify the major known performance problems.
2. Proceed with the [Performance Tuning](#) tips to help optimise performance.

### Requesting basic performance support

If the above tips don't help or you're not sure where to start, open a [support ticket](#) starting with at least the basic information:

1. The [atlassian-confluence.log](#)
2. The `catalina.out` log (or your application server log), with a series of three [thread dumps](#) separated by 10 seconds
3. A description with as much detail as possible regarding:
  - a. What changes have been made to the system?
  - b. When did performance problems begin?
  - c. When in the day do performance issues occur?
  - d. What pages or operations experience performance issues?
  - e. Is there a pattern?

Continue with as much of the advanced performance troubleshooting information as you can.

### Advanced performance troubleshooting

Please gather **all** of the information listed below and include it in your support request, even if you think you have a good idea what's causing the problem. That way we don't have to ask for it later.

#### System information

##### Confluence server

- Take a screenshot of Confluence's [Administration](#) [System Information](#) (or save the page as HTML)
- Take a screenshot of Confluence's [Administration](#) [Cache Statistics](#) (or save the page as HTML)
- Find out the exact hardware Confluence is running on
  - How many CPUs? What make and model? What MHz?
  - How much memory is installed on the machine?
  - How much memory is assigned to Confluence's JVM? (i.e. what are the `-Xmx` and `-Xms` settings for the JVM?)
  - What other applications are being hosted on the same box?

##### Confluence content

- How many *users* are registered in Confluence?
- On average, to how many *groups* does each user belong?
- How many *spaces* (global and personal) are there in your Confluence server?
- How many of those spaces would be viewable by the average user?
- Approximately how many *pages*? (Connect to your database and perform `'select count(*) from content where prevver is null and contenttype = 'PAGE''`)
- How much data is being stored in Bandana (where plugins usually store data)? (Connect to your database and perform `'select count(*), sum(length(bandanavalue)) from bandana'`)

#### On this page:

- [Basic performance troubleshooting steps](#)
- [Requesting basic performance support](#)
- [Advanced performance troubleshooting](#)

#### Related pages:

- [Memory Usage and Requirements](#)
- [Confluence for Enterprise](#)

### The database

- What is the exact version number of Confluence's database server?
- What is the exact version number of the JDBC drivers being used to access it? (For some databases, the full filename of the driver JAR file will suffice)
- Is the database being hosted on the same server as Confluence?
- If it is on a different server, what is the network latency between Confluence and the database?
- What are the database connection details? How big is the connection pool? If you are using the standard configuration this information will be in your `confluence.cfg.xml` file. Collect this file. If you are using a Data source this information will be stored in your application server's configuration file, collect this data.

### User management

- Are you using external user management or authentication? (i.e. JIRA or LDAP user delegation, or single sign-on)
- If you are using external JIRA user management, what is the latency between Confluence and JIRA's database server?
- If you are using LDAP user management:
  - What version of which LDAP server are you using?
  - What is the latency between Confluence and the LDAP server?

## Diagnostics

### Observed problems

- Which pages are slow to load?
  - If it is a specific wiki page, attach the wiki source-code for that page
- Are they always slow to load, or is the slowness intermittent?

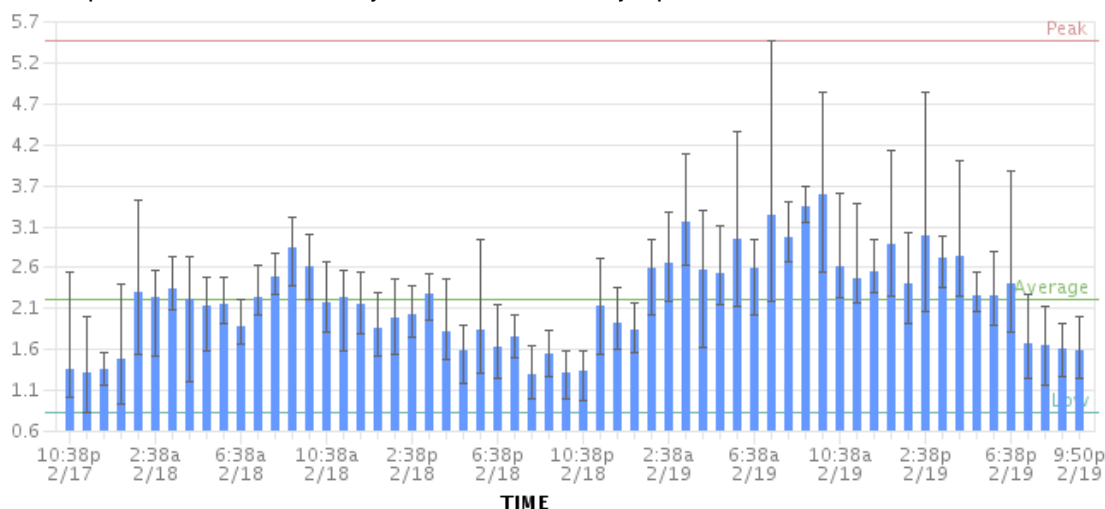
### Monitoring data

Before drilling down into individual problems, helps a lot to understand the nature of the performance problem. Do we deal with sudden spikes of load, or is it a slowly growing load, or maybe a load that follows a certain pattern (daily, weekly, maybe even monthly) that only on certain occasions exceeds critical thresholds? It helps a lot to have access to continuous monitoring data available to get a rough overview.

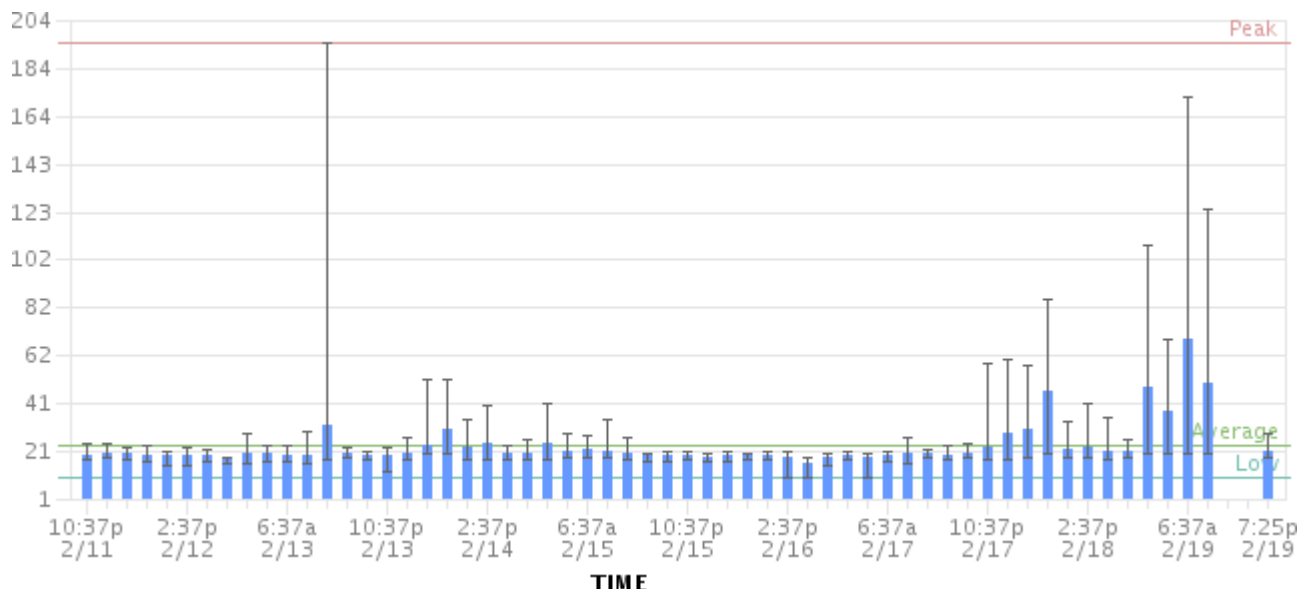
Here are sample graphs from the `confluence.atlassian.com` system, showing

### Load

This graph shows the load for two consecutive days. The obvious pattern is that the machine is under decent load, which corresponds to the user activity, and there is no major problem.

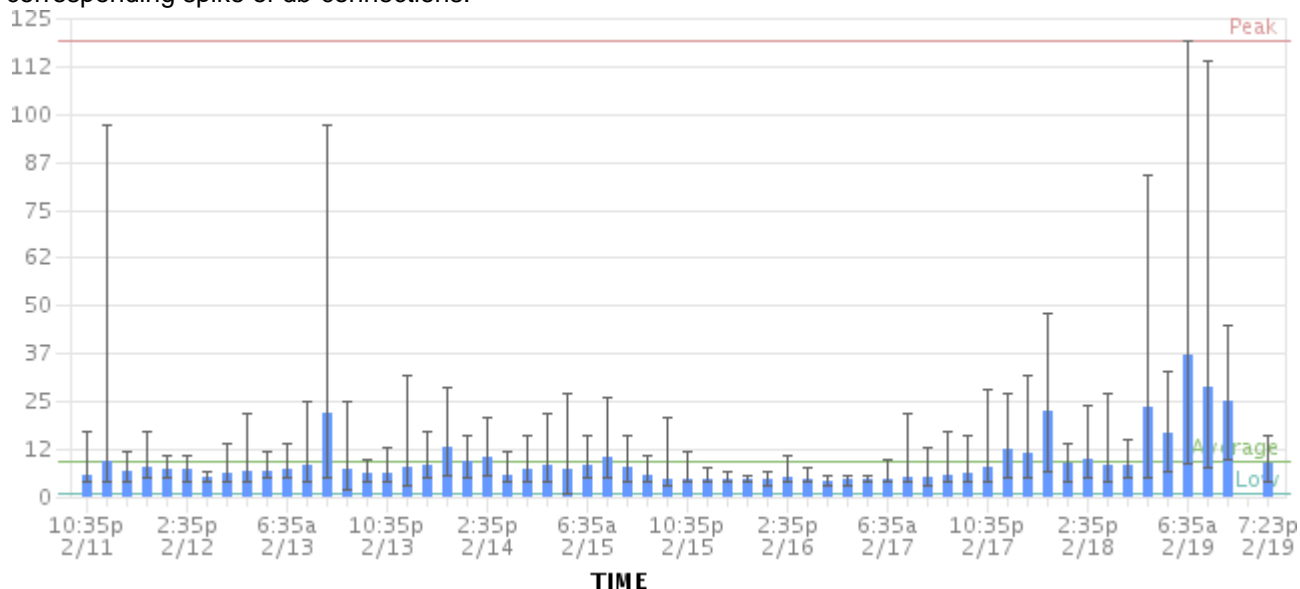


### Resin threads and database connections



#### Active number of Java Threads

These two charts show the active threads in the application server (first chart) and the size database connection pool (second chart). As you can see, there was a sudden spike of server threads and a corresponding spike of db-connections.



#### The database connection pool size

The database connection pool size peaked over 112, which happened to be more than the maximum number of connections the database was configured for (100). So it was no surprise that some requests to Confluence failed and many users thought it had crashed, since many requests could not obtain the crucial database connections.

We were able to identify this configuration problem quite easily just by looking at those charts. The next spikes were uncritical because more database connections were enabled.

The bottom line being: it helps a lot to monitor your Confluence systems continuously (we use Hyperic, for example), and it helps even more if you are able to send us graphs when you encounter problems.

#### Access logs

- [How to Enable User Access Logging](#), including redirecting the logs to a separate file
  - You can run this file through a log file analyser such as [AWSStats](#), or manually look through for pages which are slow to load.

#### Profiling and logs

- Enable Confluence's built-in profiling for long enough to demonstrate the performance problem using T

### Troubleshooting Slow Performance Using Page Request Profiling.

- If a single page is reliably slow, you should make several requests to that page
- If the performance problem is intermittent, or is just a general slowness, leave profiling enabled for thirty minutes to an hour to get a good sample of profiling times
- Find Confluence's standard output logs (which will include the profiling data above). Take a zip of the entire logs directory.
- [Take a thread dump](#) during times of poor performance

### CPU load

- If you are experiencing high CPU load, please [install the YourKit profile](#) and attach two profiler dumps taken during a CPU spike. If the CPU spikes are long enough, please take the profiles 30-60 seconds apart. The most common cause for CPU spikes is a virtual machine operating system.
- If the CPU is spiking to 100%, try [Live Monitoring Using the JMX Interface](#), in particular with the [Top threads plugin](#).

### Site metrics and scripts

- It is essential to understand the user access and usage of your instance. Please use the [access log scripts](#) and [sql scripts](#) to generate Usage statistics for your instance.

### Next step

Open a ticket on <https://support.atlassian.com> and attach all the data you have collected. This should give us the information we need to track down the source of your performance problems and suggest a solution. Please follow the progress of your enquiry on the support ticket you have created.

## Access Log Scripts

The access log scripts are [attached to this page](#). To use the scripts:

1. Unzip the 7z file.
2. Copy all the daily access logs to a folder called logs.
3. Run `Atlassian-processDailyLog.rb`. This will generate a `csv` file called `summary.csv` and several directories which contain the access logs of each defined user action.
4. Run the appropriate script `Atlassian-processDailyLog-hourly.rb` `<admin/comment/create/edit/search/rss>`. Each script will generate a different `csv` file. For example, `Atlassian-processDailyLog-hourly.rb admin` will process the admin logs extracted in step 3.
5. Import the `csv` files to `www-log-Analysis.xls` (`summary.csv` to 'raw stats - daily' sheet and `admin.csv` to 'admin -hours' sheet, etc) to generate the load profiles and graphs. You may need to modify the number of rows in each sheet depending on the number of logs.

### Note

All scripts are written in Ruby and assume the log file name contains the string 'confluence.atlassian.com-access.log'. Scripts need to be changed if another name is used. Modify the line: `filenameRegexp = Regexp.new('confluence.atlassian.com-access.log')`

## Troubleshooting Slow Performance Using Page Request Profiling

This page tells you how to enable page-request profiling. With profiling turned on, you will see a record of the time it takes (in milliseconds) to complete each action made on any Confluence page. If Confluence is responding slowly, an internal timing trace of the slow page request can help to identify the cause of the delay.

You will need access to the Confluence server to view a profile.

### On this page:

- [Enabling Page-Request Profiling](#)
- [Profiling an Activity](#)
- [Example of a Profile](#)
- [Start Confluence with Profiling Enabled](#)

### Related pages:

- [Requesting Performance Support](#)
- [Working with Confluence Logs](#)

## Enabling Page-Request Profiling

To see just the slow performing macros, see [Identifying Slow Performing Macros](#).

**i** You need to have [System Administrator](#) permissions in order to enable or disable profiling.

To enable page profiling:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose '**Logging and Profiling**' in the left-hand panel.
3. The '**Logging and Profiling**' screen appears. Choose '**Enable Profiling**'.



If profiling is already enabled, the button will be labelled 'Disable Profiling'.

To disable page profiling:

1. Choose the **cog icon**



, then choose **General Configuration** under Confluence Administration

2. Choose '**Logging and Profiling**' in the left-hand panel.
3. The '**Logging and Profiling**' screen appears. Choose '**Disable Profiling**'.



If profiling is already disabled, the button will be labelled 'Enable Profiling'.

*Screenshot: Changing Log Levels and Profiling*

## Logging and Profiling

Performance Profiling  
Profiling is currently OFF.

---

SQL Logging

---

Log4j Logging  
loglevel.profileDesc

OR:  
Customise specific logging settings  
**Add New Entry**  

Class/Package Name

New Level

INFO ▾

**Existing Levels**  

Class/Package Name	Current Level	New Level	
atlassian.plugin	INFO	INFO ▾	<a href="#">Remove</a>
com.atlassian.confluence.admin.actions	INFO	INFO ▾	<a href="#">Remove</a>

## Profiling an Activity

1. Enable profiling, using either of the methods described [above](#).  
Profiles for every page hit, for all users, will now be logged to your application server's default logs until Confluence is restarted. Note that each time a user visits a link, a single profile is printed.
2. Confirm that profiles are being written to the Confluence log file — see [Working with Confluence Logs](#) for location of the log files and other details.
3. Perform the activity that is resulting in unusually slow response time.
4. Copy the profile for that action. When deciding which profiles to copy, look for the links that took a long time to respond. If a single page is slow, only that profile is necessary. If Confluence is generally or intermittently slow, copy all profiles logged during the slowdown until a reasonable sample has been collected.
5. If you were instructed to profile your instance by Atlassian technical support, attach all relevant profiles to your support ticket.
6. Turn profiling off again, using either of the methods described [above](#).
7. Confirm that profiles are no longer being printed to the Confluence log file.

### Example of a Profile

Below are the first few lines of a normal profile for accessing a page called Confluence Overview.

```
[344ms] - /display/ds/Confluence+Overview
 [313ms] - SiteMesh: parsePage:
http://localhost:8080/display/ds/Confluence+Overview
 [313ms] - XW Interceptor: Before defaultStack:
/pages/viewpage.action (ViewPageAction.execute())
 [0ms] - SpaceAwareInterceptor.intercept()
 [16ms] - PageAwareInterceptor.intercept()
 [0ms] - AOP: PageManager.getPage()
 [16ms] - AOP: PermissionManager.hasPermission()
 [0ms] - AOP: SpacePermissionManager.hasPermission()
 [16ms] - AOP: SpacePermissionManager.hasPermission()
 [0ms] - AOP: SpacePermissionManager.hasPermission()
 [0ms] - AOP: SpacePermissionManager.hasPermission()
 [281ms] - XW Interceptor: After defaultStack:
/pages/viewpage.action (ViewPageAction.execute())
 [281ms] - XW Interceptor: After validatingStack:
/pages/viewpage.action (ViewPageAction.execute())
 ...
```

**i** Notice that each indented line is a recursive call that rolls up into the parent line. In the example above, the Confluence Overview page takes 344ms. Part of that, 313ms, is spent in sitemesh.

### Start Confluence with Profiling Enabled

There may be some situations where you may wish to have Confluence profiling enabled during startup. This may be useful if you restart often and may forget to enable profiling for Support/Trouble-shooting purposes.

Edit the file `CONFLUENCE_HOME\confluence\WEB-INF\web.xml`. You should see a stanza similar to the one below. Set the parameter value for **autostart** to **true**:

```
<filter>
 <filter-name>profiling</filter-name>

 <filter-class>com.atlassian.core.filters.ProfilingAndErrorFilter</filter-class>
 <init-param>
 <!-- specify the which HTTP parameter to use to turn the
filter on or off -->
 <!-- if not specified - defaults to "profile.filter" -->
 <param-name>activate.param</param-name>
 <param-value>profile</param-value>
 </init-param>
 <init-param>
 <!-- specify the whether to start the filter
automatically -->
 <!-- if not specified - defaults to "true" -->
 <param-name>autostart</param-name>
 <param-value>true</param-value>
 </init-param>
</filter>
```

Remember to turn it back to **false** or your logs will grow very large.

### Identifying Slow Performing Macros

[Page Profiling](#) gives good detail on what operations are slow in a page load. In addition, you can add debug level logging:

#### *Version 3.1 and Later*

Set the package name `com.atlassian.renderer.v2.components.MacroRendererComponent` to `DEBUG` in Administration >> [Logging and Profiling](#).

#### *Prior to version 3.1*

Download `WikiMarkupParser.class`, available from the [attachments to this page](#). This will result in logs like:

```
2009-04-23 10:27:54,789 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Entering macro rendering. Processed
text: {spaces}
2009-04-23 10:27:55,768 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Exiting macro text rendering. Total
time: 979ms
2009-04-23 10:27:55,785 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Entering macro rendering. Processed
text: {create-space-button}
2009-04-23 10:27:55,857 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Exiting macro text rendering. Total
time: 72ms
2009-04-23 10:27:55,862 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Entering macro rendering. Processed
text: {recently-updated-dashboard:dashboard|showProfilePic=true}
2009-04-23 10:27:56,704 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Exiting macro text rendering. Total
time: 842ms
2009-04-23 10:27:56,707 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Entering macro rendering. Processed
text: {favpages:maxResults=10}
2009-04-23 10:27:56,889 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Exiting macro text rendering. Total
time: 182ms
```

To add the class:

1. Add this line to the file <confluence-install>/confluence/WEB-INF/classes/log4j.properties:  
`log4j.logger.com.atlassian.renderer=DEBUG`
2. Add the appropriate [WikiMarkupParser.class](#) to /confluence/WEB-INF/classes/com/atlassian/renderer/v2.  
You'll have to make the renderer and v2 folders.

In combination with page profiling, this should give good specifics on the amount of time various plugins take. You can also use this utility to [Search Confluence for Uses of a Macro](#).

### Resolution

Experiment with the tips from the [performance tuning](#) page, or open an enhancement request about the specific macro. In some instances there is no resolution - you'll just be aware of the overhead of various macros.

## Compressing an HTTP Response within Confluence

Confluence supports HTTP GZip transfer encoding. This means that Confluence will compress the data it sends to the user, which can speed up Confluence over slow or congested Internet links, and reduce the amount of bandwidth consumed by a Confluence server.

Turn on Confluence's GZip encoding if:

- Users are accessing Confluence over the Internet, or a WAN connection with limited bandwidth.
- You wish to reduce the amount of data transfer between the Confluence server and client.

If you are accessing Confluence over a Local Area Network or over a particularly fast WAN, you may wish to leave GZip encoding disabled. If the network is fast enough that transferring data from Confluence to the user isn't a limiting factor, the additional CPU load caused by compressing each HTTP response may slow Confluence down.

### Enabling HTTP Compression

1. Choose the **cog icon**



- , then choose **General Configuration** under Confluence Administration
2. Select '**General Configuration**' in the left-hand panel.
  3. Enable '**Compress HTTP Responses**'.

It is possible to configure which types of content are compressed within Confluence. By default, the following mime types will be compressed:

- text/htmltext
- javascript
- text/css
- text/plain
- application/x-javascript
- application/javascript

If you wish to change the types of content to be compressed, add a replacement `urlrewrite-gzip-default.xml` file within the `WEB-INF/classes/com/atlassian/gzipfilter/` directory in your Confluence Installation Directory. A sample file is provided as an [attachment](#). It is unlikely that you will need to alter this file.

## Garbage Collector Performance Issues

This document relates broadly to memory management with Oracle's Hotspot JVM. These recommendations are based on Support's successful experiences with customers with large Confluence instances.

Please **do not** use the Concurrent Mark Sweep (CMS) Collector with Confluence, unless otherwise advised by Atlassian Support. It requires extensive manual tuning and testing, and is likely to result in degraded performance.

### Use a small heap

Keep your heap as small as possible, without the instance experiencing OutOfMemory errors. If you experience OutOfMemory errors and need to increase this, we recommend you do it in 512mb or 1gb allotments, and monitor the instance. If you continue to receive OutOfMemory errors, increase the heap by another 512mb or 1gb, and continue this process until you are operating stably with no OutOfMemory errors. Do not increase the heap further than required, as this will result in longer garbage collections.

### Remove any old tuning parameters

On every full GC, the JVM will resize the allocations of Eden, Survivor etc based on the throughput it is actually seeing. It will tune itself based on the real world data of the objects that are being created and collected. Most of the time simply allowing JVM to tune itself will give you better performance.

If you have added JVM parameters in the past and are experiencing difficulties with GC now, we'd recommend you remove all GC related parameters, unless you added them to solve a specific problem, and they did in fact solve that problem. You should also consider re-benchmarking now to ensure that they are still solving that problem, and are not causing you any other issues.

### Check your VM resources

If you run Confluence on a VM, check that it is not using the swap file. If it does, when the JVM garbage collects it has to load the objects from the swap file into memory to clean them, and this can cause significantly longer GC pauses. Instead of using swapping, ballooning and bursting, allocate adequate memory to the VM.

### Manual Tuning

If you find you are still experiencing difficulties with GC after following these recommendations and you would like to see if you can tune the JVM better to improve performance, we recommend following the instructions in our [Garbage Collection \(GC\) Tuning Guide](#). This document will take you through the process of choosing

performance goals (throughput/footprint/latency), and how to tune for those goals.

## Viewing your GC logs

[How to Enable Garbage Collection \(GC\) Logging](#), and use a tool like [Chewiebug's GCViewer](#) to view the resulting logs.


## Cluster Troubleshooting

This page covers troubleshooting for a clustered installation of Confluence.

- For information about clustering in general, refer to [Confluence Enterprise Resources](#).
- If you're experiencing Cluster Panic messages in non-clustered installation of Confluence, visit the Knowledge Base article '[Database is being updated by an instance which is not part of the current cluster](#)' Error Message.

### On this page:

- [Symptoms](#)
- [Didn't find a solution?](#)

 You must ensure the clocks on your cluster nodes don't diverge, as it can result in a range of problems with your cluster.

## Symptoms

Below is a list of potential problems with a Confluence cluster, and their likely solutions. The solutions are listed below.

Problem	Likely solutions
Database is being updated by an instance which is not part of the current cluster errors on a stand-alone	<a href="#">'Database is being updated by an instance which is not part of the current cluster' Error Message</a>
Database is being updated by an instance which is not part of the current cluster errors on a cluster	<a href="#">Add multicast route, Check firewall, Cluster Panic due to Multiple Deployments</a>
Cannot assign requested address on startup, featuring an IPv6 address	<a href="#">Prefer IPv4</a>
Error in log: The interface is not suitable for multicast communication	<a href="#">Change multicast interface, Add multicast route</a>
Multicast being sent, but not received	<a href="#">Check firewall, Check intermediate routers, Increase multicast TTL</a>
Add-on is unlicensed on some nodes after updating the license on one node.	Disable and re-enable the add-on in the Universal Plugin Manager.
After a add-on update, strings appear in the UI instead of buttons and icons on some nodes.	Restart the affected node.
Any issue not covered here	<a href="#">Contact support</a>

## Multicast

- Which multicast address?

The multicast address and port used by Confluence can be found on the Cluster Configuration page, or in `confluence.cfg.xml` in the Confluence home directory.

- Multicast address generation.

Confluence uses a hashing algorithm to take the inputted name during setup and it is then turned into a

multicast address stored in the config file. Thus, once the initial setup is completed, Confluence will use the address this is the reason why user can change the address if needed, without actually changing the name. Consequently the additional nodes using the same multicast address specified in the config file are able to join the cluster.

Each node has a multicast address configured in the `confluence-cfg.xml` file

```
name="confluence.cluster.address">xxx.xx.xxx.xxx</property>
```

A warning message is displayed when an user changes the address from the one that Confluence has generated by the hashing of the name. There is no way of eliminating the message any other way other than by returning the address to the one that matches the cluster name. Purpose of the warning message is to remind the user that the address has been changed - as it is not the hashed version any longer - consequently the node can not join the cluster just by using the name. It is also necessary to provide the correct address as well.

### Mapping interface to IP address.

To ensure that the interface name is mapped correctly, the following [tool](#) can be used. It shows the mapping of the interface name to the IP address.

```
C:\>java -jar list-interfaces.jar
interfaces.size() = 4
networkInterface[0] = name:lo (MS TCP Loopback interface) index: 1
addresses:
/127.0.0.1;

networkInterface[1] = name:eth0 (VMware Virtual Ethernet Adapter for
VMnet8) index: 2 addresses:
/192.168.133.1;

networkInterface[2] = name:eth1 (VMware Virtual Ethernet Adapter for
VMnet1) index: 3 addresses:
/192.168.68.1;

networkInterface[3] = name:eth2 (Broadcom NetXtreme 57xx Gigabit
Controller - Packet Scheduler Miniport) index: 4 addresses:
/192.168.0.101;
```

### Debugging tools

Listed below are some debugging tools that help determine what the status of the multicast traffic is:

Tool	Information provided
<code>netstat -gn</code>	Lists multicast groups. Does not work on Mac OS X.
<code>netstat -rn</code>	Lists system routing table.
<code>tcpdump -i interface</code>	Captures network traffic on the given interface. Most useful on an interface that only receives cluster traffic.

### Add multicast route

Multicast networking requirements vary across operating systems. Some operating systems require little

configuration, while some require the multicast address to be explicitly added to a network interface before Confluence can use it. If multicast traffic can't be sent or received correctly, adding a route for multicast traffic on the correct interface will often fix the problem. The example below is for a Ubuntu Linux system:

```
route add -net 224.0.0.0 netmask 240.0.0.0 dev eth0
```

To support multiple applications using multicast on different interfaces, you may need to specify a route specific to the Confluence multicast address.

### Check firewall

Ensure your firewall allows UDP traffic on the multicast address and port used by Confluence.

### Prefer IPv4

There are known issues relating to IPv6. You should configure your JVM to try binding to an IPv4 address first.

### Change multicast interface

Confluence might have selected the incorrect interface for multicast traffic, which means it cannot connect to other nodes in the cluster. To override the interface used for multicast traffic after initial setup, edit the `confluence.cluster.interface` property in `<local-home>/confluence.cfg.xml` and specify the network interface. For example to tell Confluence to use `eth1`:

```
<property name="confluence.cluster.interface">eth1</property>
```

### Overriding Hazelcast Configuration

If the solution to your problem involves changes to the Hazelcast configuration, these changes should **not** be made to the Confluence configuration files. Instead, to ensure your configuration survives upgrades, make your changes by creating a Hazelcast override file.

### Increase multicast TTL

The multicast time-to-live (TTL) specifies how many *hops* a multicast packet should be allowed to travel before it is discarded by a router. It should be set to the number of routers in between your clustered nodes: 0 if both are on the same machine, 1 if on two different machines linked by a switch or cable, 2 if on two different machines with one intermediate router, and so on.

To increase the multicast TTL by edit the `confluence.cluster.ttl` property in the `<local-home>/confluence.cfg.xml` file on each node. For example to set the TTL to 3:

```
<property name="confluence.cluster.ttl">3</property>
```

### Check intermediate routers

Advanced switches and routers have the ability to understand multicast traffic, and route it appropriately. Unfortunately sometimes this functionality doesn't work correctly with the multicast management information (IGMP) published by the operating system running Confluence.

If multicast traffic is problematic, try disabling advanced multicast features on switches and routers in between the clustered nodes. These features can prevent multicast traffic being transmitted by certain

operating systems.

### Didn't find a solution?

#### **Check Related Articles from the Confluence Knowledge Base**

- [Recovering from a Data Center cluster split-brain](#)
- [Port \[5801\] is already in use and auto-increment is disabled. Hazelcast cannot start.](#)
- ["Exception bootstrapping cluster:Shared home directory is not configured correctly" Error during Confluence Data Center startup](#)
- [Multicast communication works only one-way](#)
- [Cannot find "external\\_id" column when trying to upgrade to a Confluence CDC license after upgrading from a pre-5.5 Confluence Clustered installation](#)
- [Hazelcast CANNOT start on this node. No matching network interface found.](#)
- [Cluster Panic due to Multicast Traffic Communication Problem](#)
- [Configuration of Confluence Cluster Fails with 'Cannot assign requested address'](#)
- [How to suppress cluster warning messages in the Confluence log files](#)

#### **Contact Atlassian support**

We have dedicated staff on hand to support your installation of Confluence. Please follow the instructions for [raising a support request](#) and mention that you're having trouble setting up your Confluence cluster.

## Data Collection Policy

### Why does Confluence collect usage data?

We're proud that Confluence is one of the most versatile collaboration tools on the planet, and we will continue to deliver innovative new features as quickly as we can. In order to prioritize the features we deliver, we need to understand how our customers use Confluence, what's important, what's not, and what doesn't work well. The collection of usage data allows us to measure the user experience across many thousands of users and deliver features that matter.

### What data is collected?

The type of data we collect is covered in our [Privacy Policy](#). Please read it - we've tried to avoid legal jargon and made it as straightforward as possible.

To view a sample of data that might be collected from your specific installation, go to



> **General Configuration > Analytics.**

Data is always collected in Confluence Cloud.

### How is data collected from Confluence?

Older versions of Confluence (prior to Confluence 5.6 or Confluence Questions 1.0.618) didn't collect usage data. Analytics are collected using the Atlassian Analytics add-on. The add-on collects analytics events in a log file which is located in `<confluence-home>/analytics-logs`. The logs are periodically uploaded using an encrypted session and then deleted. If Confluence is unable to connect to the Internet, no logs are ever uploaded.

### Enabling/disabling data collection in Confluence

You can turn off analytics collection at any time. Go to



> **General Configuration** > **Analytics**.