



Documentation for Confluence 5.5

Contents

Confluence Administrator's Guide	7
Getting Started as Confluence Administrator	8
Managing Confluence Users	11
Adding and Inviting Users	12
Deleting or Deactivating Users	16
Searching For and Administering Users	18
Editing User Details	20
Resetting the Login Count for a User	21
Changing Usernames	22
Restoring Passwords To Recover Admin User Rights	23
Managing Site-Wide Permissions and Groups	27
Global Groups Overview	27
Adding or Removing Users in Groups	28
Global Permissions Overview	30
Setting Up Public Access	35
Configuring User Directories	36
Configuring the Internal Directory	38
Connecting to an LDAP Directory	39
Configuring the LDAP Connection Pool	48
Configuring an SSL Connection to Active Directory	49
Connecting to an Internal Directory with LDAP Authentication	60
Connecting to Crowd or JIRA for User Management	66
Reverting from Crowd or JIRA to Internal User Management	74
Connecting to JIRA 4.2 or Earlier for User Management	78
Managing Multiple Directories	79
Managing Nested Groups	81
Synchronising Data from External Directories	85
Diagrams of Possible Configurations for User Management	87
User Management Limitations and Recommendations	92
Requesting Support for External User Management	96
Disabling the Built-In User Management	98
Managing Add-ons and Macros	99
About Add-ons	99
Add-on loading strategies in Confluence	100
Removing Malfunctioning Add-ons	102
Enabling and Configuring Macros	104
Configuring the User List Macro	104
Enabling HTML macros	105
Enabling the html-include Macro	105
Troubleshooting the Gallery Macro	107
Adding, Editing and Removing User Macros	107
Writing User Macros	108
Best Practices for Writing User Macros	112
Examples of User Macros	113
Guide to User Macro Templates	123
Configuring the Office Connector	129
Customising your Confluence Site	135
Changing the Look and Feel of Confluence	135
Customising the Confluence Dashboard	135
Changing the Site Logo	137
Customising Colour Schemes	139
Working with Themes	142
Applying a Theme to a Site	142
Creating a Theme	143
Customising Site and Space Layouts	144
Upgrading Customised Site and Space Layouts	145

Working With Decorator Macros	147
Custom Decorator Templates	148
Customising a Specific Page	149
Customising the Login Page	150
Modify Confluence Interface Text	150
Customising the eMail Templates	152
Changing the Default Behaviour and Content in Confluence	152
Administering Site Templates	153
Importing Templates	153
Changing the Site Title	155
Choosing a Default Language	155
Configuring the Administrator Contact Page	156
Configuring the Site Home Page	158
Configuring the What's New Dialog	159
Customising Default Space Content	160
Customising the Getting Started Guide on the Dashboard	161
Editing the Site Welcome Message	163
Integrating Confluence with Other Applications	165
Linking to Another Application	165
Configuring Workbox Notifications	166
Integrating JIRA and Confluence	170
Installing Confluence and JIRA Together	171
Setting Up Trusted Communication between JIRA and Confluence	172
Registering External Gadgets	177
Managing your Confluence License	180
Viewing and Editing License Details	180
Getting a Confluence License	183
Reducing the User Count for your Confluence License	183
Finding Your Confluence Support Entitlement Number (SEN)	184
Managing Confluence Data	186
Database Configuration	186
Database JDBC Drivers	188
Database Setup for Oracle	189
Configuring an Oracle Datasource in Apache Tomcat	194
Database Setup for SQL Server	196
Configuring a SQL Server Datasource in Apache Tomcat	199
Database Setup For MySQL	200
Configuring a MySQL Datasource in Apache Tomcat	204
Database Setup for PostgreSQL	206
Configuring a PostgreSQL Datasource in Apache Tomcat	209
Embedded HSQLDB Database	211
Connecting to HSQLDB using DBVisualizer	211
Troubleshooting the Embedded HSQLDB Database	214
Migrating to Another Database	215
Configuring Database Character Encoding	220
Configuring database query timeout	223
Troubleshooting External Database Connections	224
Improving Database Performance	225
Creating a Lowercase Page Title Index	226
Surviving Database Connection Closures	228
Site Backup and Restore	229
Production Backup Strategy	230
Configuring Backups	232
User Submitted Backup & Restore Scripts	233
Manually Backing Up the Site	235
Restoring a Site	236
Restoring a Space	237
Restoring a Test Instance from Production	239
Restoring Data from other Backups	240
Retrieving File Attachments from a Backup	241
Troubleshooting failed XML site backups	243
Migrating from HSQLDB to MySQL	246

Troubleshooting XML backups that fail on restore	251
Attachment Storage Configuration	255
Configuring Attachment Size	257
Hierarchical File System Attachment Storage	258
Confluence Data Model	262
Finding Unused Spaces	269
Data Import and Export	270
Configuring a Confluence Environment	270
Important Directories and Files	271
Confluence Data Directory Configuration	273
Confluence Home Directory	274
Confluence Installation Directory	276
Application Server Configuration	276
Configuring URL Encoding on Tomcat Application Server	276
Managing Application Server Memory Settings	277
Switching to Apache Tomcat	277
Java Policy Settings for Enterprise or Webhosting Environments	280
Web Server Configuration	281
Configuring Web Proxy Support for Confluence	281
Running Confluence behind Apache	283
General Apache Configuration Notes	284
Using Apache with mod_proxy	285
Using Apache with virtual hosts and mod_proxy	288
Using Apache with mod_jk	290
Using mod_rewrite to Modify Confluence URLs	292
Configuring Apache to Cache Static Content via mod_disk_cache	293
Starting Confluence Automatically on System Startup	294
Start Confluence Automatically on Linux	294
Start Confluence Automatically on Windows as a Service	298
Configuring Confluence	300
Viewing System Information	301
Live Monitoring Using the JMX Interface	301
Tracking Customisations Made to your Confluence Installation	304
Viewing Site Statistics	304
Viewing System Properties	307
Configuring the Server Base URL	307
Configuring the Confluence Search and Index	308
Configuring Indexing Language	308
Configuring Quick Navigation	309
Content Index Administration	309
Enabling OpenSearch	312
Rebuilding the Ancestor Table	312
Setting Up Confluence to Index External Sites	312
Setting Up an External Search Tool to Index Confluence	313
Configuring Mail	313
Configuring a Server for Outgoing Mail	314
Setting Up a Mail Session for the Confluence Distribution	315
Configuring the Recommended Updates Email Notification	316
The Mail Queue	316
Configuring Character Encoding	317
Troubleshooting Character Encodings	319
"€" Euro character not displaying properly	321
MySQL 3.x Character Encoding Problems	322
Other Settings	322
Configuring a WebDAV client for Confluence	322
Configuring HTTP Timeout Settings	327
Configuring Number Formats	328
Configuring Shortcut Links	328
Configuring Time and Date Formats	329
Enabling the Remote API	330
Enabling Threaded Comments	330
Enabling Trackback	330

Installing a Language Pack	331
Installing Patched Class Files	332
Configuring System Properties	333
Recognised System Properties	347
Working with Confluence Logs	358
Configuring Logging	360
log4j Logging Levels	362
Troubleshooting SQL Exceptions	362
Configuring Confluence Security	363
Confluence Security Overview and Advisories	364
Confluence Cookies	366
Configuring Secure Administrator Sessions	368
Using Fail2Ban to limit login attempts	369
Securing Confluence with Apache	371
Using Apache to limit access to the Confluence administration interface	372
Managing External Referrers	374
Excluding external referrers	375
Hiding external referrers	376
Ignoring External Referrers	376
Best Practices for Configuring Confluence Security	377
Hiding the People Directory	378
Configuring Captcha for Spam Prevention	379
Hiding External Links From Search Engines	379
Configuring Captcha for Failed Logins	380
Configuring XSRF Protection	382
User Email Visibility	383
Anonymous Access to Remote API	383
Running Confluence Over SSL or HTTPS	384
Connecting to LDAP or JIRA or Other Services via SSL	389
Configuring RSS Feeds	390
Preventing and Cleaning Up Spam	391
Scheduled Jobs	393
Configuring the Whitelist	398
Operating Large or Mission-Critical Confluence Installations	399
Configuring a Large Confluence Installation	408
Confluence Clustering Overview	410
Technical Overview of Clustering in Confluence	411
Cluster safety mechanism	415
Changing Datasources Manually in a Cluster	416
Cluster Troubleshooting	416
Multicast Test	421
Clustering for Scalability vs Clustering for High Availability (HA)	421
Recommended network topology	424
Cluster Administration page	425
Cluster Checklist	426
Performance Tuning	431
Cache Performance Tuning	435
Cache Performance Tuning for Specific Problems	446
Cache Statistics	448
Confluence Cache Schemes	450
Memory Usage and Requirements	451
Requesting Performance Support	453
Access Log Scripts	457
Troubleshooting Slow Performance Using Page Request Profiling	457
Identifying Slow Performing Macros	460
Compressing an HTTP Response within Confluence	461
Garbage Collector Performance Issues	462
Confluence Installation and Upgrade Guide	463
System Requirements	464
Server Hardware Requirements Guide	465
Example Size and Hardware Specifications From Customer Survey	469
Running Confluence in a Virtualised Environment	471

Confluence Installation Guide	472
Installing Confluence	472
Installing Confluence on Windows	473
Installing Confluence on Windows from Zip File	476
Uninstalling Confluence from Windows	479
Installing Confluence on Linux	479
Installing Confluence on Linux from Archive File	482
Uninstalling Confluence from Linux	485
Change listen port for Confluence	486
Installing the Confluence EAR-WAR Edition	487
Known Issues for Apache Tomcat	490
Installing Java for Confluence	491
Setting the JAVA_HOME Variable in Windows	492
Confluence Cluster Installation	493
Confluence Cluster Installation with Existing Data	496
Apache and Tomcat load balancing	498
Upgrading a Confluence Cluster	500
Creating a Dedicated User Account on the Operating System to Run Confluence	502
Confluence Setup Guide	503
Configuring JIRA Integration in the Setup Wizard	510
Upgrading Confluence	517
Upgrading Beyond Current Licensed Period	520
Confluence Post-Upgrade Checks	521
Upgrading Confluence EAR-WAR Distribution	522
Migration from Wiki Markup to XHTML-Based Storage Format	527
Migration of Templates from Wiki Markup to XHTML-Based Storage Format	530
Upgrading Confluence Manually	532
Supported Platforms	537
End of Support Announcements for Confluence	539
Supported Platforms FAQ	549
Migrating Confluence Between Servers	550
From Confluence Evaluation through to Production Installation	553
Migrating from Confluence OnDemand to a Confluence Installed Site	555

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 using creating and administering space refer to the [Confluence User's Guide](#).

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](#)
- [Managing Confluence Users](#)
 - [Adding and Inviting Users](#)
 - [Deleting or Deactivating Users](#)
 - [Searching For and Administering Users](#)
 - [Managing Site-Wide Permissions and Groups](#)
 - [Configuring User Directories](#)
- [Managing Add-ons and Macros](#)
 - [About Add-ons](#)
 - [Enabling and Configuring Macros](#)
 - [Adding, Editing and Removing User Macros](#)
 - [Configuring the Office Connector](#)
- [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](#)
- [Managing your Confluence License](#)
 - [Viewing and Editing License Details](#)
 - [Getting a Confluence License](#)
 - [Reducing the User Count for your Confluence License](#)
 - [Finding Your Confluence Support Entitlement Number \(SEN\)](#)
- [Managing Confluence Data](#)
 - [Database Configuration](#)
 - [Site Backup and Restore](#)
 - [Attachment Storage Configuration](#)
 - [Confluence Data Model](#)
 - [Finding Unused Spaces](#)
 - [Data Import and Export](#)
- [Configuring a Confluence Environment](#)
 - [Important Directories and Files](#)
 - [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](#)

Downloads



Download the Confluence documentation in PDF format.

Other Resources

[Confluence User's Guide](#)

[Confluence Installation and Upgrade Guide](#)

[Confluence Knowledge Base](#)

[Atlassian Answers](#)

- [and Index](#)
- [Configuring Mail](#)
- [Configuring Character Encoding](#)
- [Other Settings](#)
- [Configuring System Properties](#)
- [Working with Confluence Logs](#)
- [Configuring Confluence Security](#)
- [Scheduled Jobs](#)
- [Configuring the Whitelist](#)
- [Operating Large or Mission-Critical Confluence Installations](#)
 - [Configuring a Large Confluence Installation](#)
 - [Confluence Clustering Overview](#)
 - [Performance Tuning](#)

Getting Started as Confluence Administrator

This page is an introduction for people just starting out as Confluence administrators. You will find this page useful if your Confluence site is brand new, or you are 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](#)

Related pages:

- [Getting Help and Support](#)
- [Confluence Administrator's Guide](#)

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 'GG': Press 'G' twice on your keyboard then continue typing the action you want.

Notes about finding administrative options via the search box:

- Pressing 'GG' puts your cursor into the search box.
- The 'GG' combination is familiar to JIRA users, because the same shortcut opens the JIRA administration search dialog.
- System administration, Confluence administration and space administration options may appear in the search results.
- Confluence permissions determine the administrative options that appear in the search results. You will only see the options that you have permission 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 administrative functions are available from the Confluence Administration Console, which is part of the Confluence web interface. If you have [administrative permissions](#), you will 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 the **cog icon** , then choose **General Configuration** under Confluence Administration.

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 [important directories and files](#). 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 [Creating a Space](#).
- Invite some users to your site. See [Adding and Inviting 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 [Giving People Access to Content](#).
- Get to know the power of add-ons (also called plugins), for extending and customising your Confluence site. See [About Add-ons](#).
- Investigate more ways of customising Confluence. See [Customising your Confluence Site](#).

Now you are ready to dive into the [Confluence Administrator's Guide](#).

Managing Confluence 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:

- [Adding and Inviting Users](#)
- [Deleting or Deactivating Users](#)
- [Searching For and Administering Users](#)
- [Managing Site-Wide Permissions and Groups](#)

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](#) 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](#). *Not applicable to Confluence OnDemand.*

On this page:

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

Related pages:

- [Configuring Confluence Security](#)
- [Confluence Administrator's Guide](#)

 *The information on this page does not apply to Atlassian OnDemand sites with multiple apps. If you are using Confluence-only OnDemand, the information does apply.*

Authentication

Seraph

Almost all authentication in Confluence (and JIRA) 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 OnDemand.*
- **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.

Adding and Inviting 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. See below.
- **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. See below.
- **By adding users manually:** Administrators with Confluence Administrator or System Administrator [permissions](#) can add new users. See below.
- **Via an external user directory:** See [Configuring User Directories](#). *Not applicable to Confluence OnDemand.*

You may also be interested in information about allowing anonymous users access to your site. Anonymous users do not count against your Confluence license totals. See [Setting Up Public Access](#).

Note: If you are using Confluence OnDemand with multiple applications, please refer to the following guide for information on adding and inviting users: [Managing Users and Groups](#).

Allowing 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.

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 **Invite Users** on the dashboard, then choose **User Signup Options**.
Or take the longer route: Choose the **cog icon** , then choose **General Configuration** under Confluence Administration. Then choose **Users > User Signup Options**.
2. Choose **Allow people to sign up to create their account**.
3. 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, a text box will appear. 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.
4. 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.

On this page:

- [Enabling and disabling notifications about user signup](#)
- [Inviting people to sign up](#)
- [Resetting the invitation link](#)
- [Adding users manually](#)
- [Notes](#)

Related pages:

- [Managing Confluence Users](#)
- [Setting Up Public Access](#)
- [Configuring a Server for Outgoing Mail \(Not applicable to Confluence OnDemand.\)](#)
- [Confluence Administrator's Guide](#)

 *The information on this page does not apply to Atlassian OnDemand sites with multiple apps. If you are using Confluence-only OnDemand, the information does apply.*

Enabling and disabling notifications about user signup

By default, Confluence will send an email notification to all Confluence administrators whenever someone signs up to the 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 **Invite Users** on the dashboard, then choose **User Signup Options**.
Or take the longer route: Choose the **cog icon** , then choose **General Configuration** under

- Confluence Administration. Then choose **Users > User Signup Options**.
- Remove the tick from **Notify administrators by email when an account is created**.
- Choose **Save**.

Screenshot: User signup options

Users

List Users Add Users Invite Users **User Signup Options**

The user signup mode defines if and how your users are able to create their own accounts without relying on an administrative action. [Read more about Confluence's signup options.](#)

☒ **Allow people to sign up to create their account**

☒ **Restricted by Domain(s)**
 Only people with an email address that matches the specified domain(s) will be able to create their own account.
 Separate multiple domains with commas, semicolons or white spaces

☐ **No restrictions**
 Everyone will be able to create their own account.

☐ **Notify administrators by email when an account is created**

Save

Inviting 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:

- Choose **Invite Users** on the dashboard.
 Or take the longer route: Choose the **cog icon** , then choose **General Configuration** under Confluence Administration. Then choose **Users > Invite Users**.
- 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:
 - 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
 - Optional: Change the **Message** if you want to.
 - Choose **Send**.

Resetting 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 do not 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 become unusable. People will no longer be able to use the previous link to sign up. If they try, they will see an error message that the signup token has expired.

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](#)

Adding users manually

To add a new user:

1. Choose **Invite Users** on the dashboard, then choose **Add Users**.
Or take the longer route: Choose the **cog icon** , then choose **General Configuration** under Confluence Administration. Then choose **Users > Add Users**.
2. Enter the user's details: username, name, password, and email address.
3. 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.

4. Choose **Create**.

Screenshot: Adding users

Users

List Users Add Users Invite Users User Signup Options

Add a User

Username*

Username 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 may define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you may use the default Confluence **internal directory** and also connect to an **LDAP** directory server. In such cases, 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](#). *Not applicable to Confluence OnDemand.*
- **Confluence OnDemand:** If you are using Confluence OnDemand with multiple applications, please refer to the following guide for information on adding and inviting users: [Managing Users and Groups](#).

Deleting or Deactivating Users

If you are a [Confluence Administrator](#), you can delete and deactivate users.

You can **delete** a user from Confluence if they have not yet added or edited any content on the site. Such

content includes pages and blog posts, and edits and comments on existing pages.

You can **deactivate**, or disable, a user, including one who has contributed content.

- Deactivated users can no longer log in to Confluence.
- Deactivating a user will not remove the content created by them.
- Deactivated users do not count towards your license count. (See the notes below.)

To remove a user:

1. Go to the user's [profile](#) and choose **Administer User**.
2. Choose **Delete**.

To deactivate a user:

1. Go to the user's [profile](#) and choose **Administer User**.
2. Choose **Disable**.

Related pages:

- [Managing Confluence Users](#)
- [Configuring User Directories](#) (Not applicable to Confluence OnDemand.)
- [Confluence Administrator's Guide](#)

 The information on this page *does not apply* to Atlassian OnDemand sites with multiple apps. For information on managing access for users if you use multiple applications in OnDemand, see [Managing Application Access](#).

If you are using Confluence-only OnDemand, the information does apply.

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 may define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you may use the default Confluence **internal directory** and also connect to an **LDAP** directory server. In such cases, 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. See [Viewing and Editing License Details](#). The number of registered users includes only users who have the 'Can Use' global permission. Deactivated users, as described above, are not included. Choose **Refresh** to make sure you see the latest count.

Searching For and Administering Users

If you have [Confluence Administrator](#) permissions, you can view users, edit their user details, reset their passwords, and assign them to groups.

Accessing the user management screen

There are two ways to do this.

Option 1: Administer a known user:

- Go to a user's [profile](#)
- Choose **Administer User**.

Option 2: Find the user first:

- Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
- Choose **Users** in the left-hand panel.
- The 'Users' screen appears. You can now list all users or search for a specific user.

Listing all users

To list all users:

1. Choose **Show all users**. All members of the 'confluence-users' or 'users' group are listed in alphabetical order, by username. If there are more users than can fit on one page, the results will be divided into multiple pages.
2. To move to another page of results, choose the numbered links, **Next** or **Previous** near the top or bottom of the page.
3. To specify how many results should be shown per page, choose a number **10**, **20**, **50** or **100** near the top of the page.

On this page:

- [Accessing the user management screen](#)
- [Listing all users](#)
- [Using the simple user search](#)
- [Using the advanced user search](#)
- [Notes](#)

Related pages:

- [Adding and Inviting Users](#)
- [Giving People Access to Content](#)
- [Confluence Administrator's Guide](#)

 *The information on this page **does not apply** to Atlassian OnDemand sites with multiple apps. If you are using Confluence-only OnDemand, the information does apply.*

Using the simple user search

To search for a user via the simple user search:

1. If the **Simple** link is showing, choose it. (If you see the 'Advanced' link and no 'Simple' link, then the simple search is already active.)
2. Type some information about the user into the 'Find User' text box. You can type all or part of their username, full name or email address.
3. Choose **Search**.
4. Confluence will display a list of matching users. Click the link on a username to see and edit the details for that user.

Using the advanced user search

The advanced user search allows you to specify the field in which your search term appears: username, full name or email address. This is useful if you need to limit the number of users appearing in the search results.

To search via the advanced user search:

1. If the **Advanced** link is showing, choose it. (If you see the 'Simple' link and no 'Advanced' link, then the advanced search is already active.)
2. Complete one or more of the following fields:
 - **Username** — Enter all or part of the person's username. This is their login ID, such as 'joe', or 'bloggs'.
 - **Full Name** — Enter all or part of the person's name. For example, 'joe bloggs', or 'bloggs', or 'joe'.
 - **Email** — Enter all or part of the person's email address. For example, 'acme'.
3. Choose **Search**.
4. Confluence will display a list of matching users. Click the link on a username to see and edit the details for that user.

Notes

- **Multiple user directories:** You may define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you may use the default Confluence **internal directory** and also connect to an **LDAP** directory server. In such cases, 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](#).

- **Crowd and the user search:** If you are using Atlassian's [Crowd](#) for user management, you will need [Crowd 1.5.1 or later](#) to use the 'Simple' option in the user search. If your version of Crowd does not support the simple user search, you will see only the 'Advanced' search form.

Screenshot: The user management screen

Users

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

Find User

Find User

Search for users by their user details eg: bob, Bob Smith, bob@example.com

[Search](#) [Show all users](#)

SHOW: 10

User	Username	Email
Ewan User	ewan	sample@email.com.au
Josh User	josh	sample@email.com.au

Editing User Details

You need [Confluence administrator](#) permissions to be able to edit the details of a user. The details include the person's name, password, email address, group membership, and ability to access Confluence.

To update a user's details:

- First, 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 click the '**Administer User**' link on the user's profile screen.
 - Or, Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
 - Select the link '**Manage Users**' in the left-hand panel.
 - Locate the user by doing a search on the username or the groups to which they belong.
 - Click the user link.
- Now you should be able to see the user'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 [Changing Usernames](#) 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.

Screenshot: User details

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](#)

Related pages:

- [Searching For and Administering Users](#)
- [Deleting or Deactivating Users](#)
- [Adding or Removing Users in Groups](#)
- [Adding and Inviting Users](#)
- [Confluence Administrator's Guide](#)

Notes

Multiple user directories: You may define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you may use the default Confluence **internal directory** and also connect to an **LDAP** directory server. In such cases, 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](#).

Resetting the Login Count for a User

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

If you are a [Confluence Administrator](#), you can manually set the failed login count for a user back to zero.

To reset the failed login count for a user:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **Manage Users** in the left-hand panel.
3. Search for the required user and click the user in the search results. The 'View User' screen will appear.
4. Choose **Reset Failed Login Count** for the user. The 'Current Failed Login Count' will be reset to 0.

Related pages:

- [Configuring Captcha for Failed Logins](#) (Not applicable to Confluence OnDemand.)
- [Confluence Administrator's Guide](#)

 The information on this page *does not apply* to Atlassian OnDemand sites with multiple apps. If you are using Confluence-only OnDemand, the information does apply.

Screenshot: Resetting the failed login count for a user

View User: josh

« [Back to Users](#)

User: josh

Full Name: Josh User

Email: sample@email.com.au

Directory: Confluence Internal Directory

Created: Feb 11, 2013 09:50

Last Updated: Feb 11, 2013 15:09

Login: CAPTCHA required at next login Last Login: Feb 18, 2013 13:45 Last Failed Login: Feb 18, 2013 15:43 Total Failed Login Count: 5 Current Failed Login Count: 5 ([Reset Failed Login Count](#))

Groups: ☒ [confluence-users](#)
☒ [developers](#)

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

Changing Usernames

A username is the name used to log in to Confluence, for example jsmith. If you are a Confluence administrator you can change a user's username, for example if they change their name.

 The information on this page *does not apply* to Confluence OnDemand.

All active users must have a unique username. This means that two active users cannot 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. Refer to [Configuring User Directories](#) for more information.

Changing the username of a user managed in Confluence

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

To change a username:

1. Navigate to the user's record:
 - a. go to the user's profile and choose **Adminster User**, or
 - b. go to **Confluence Admin > Users** and search for the user

2. Choose **Edit Details**.
3. Enter the new username and choose **Submit**.

The user's username has been changed. The user 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 user mentions.

Changing the username of a user managed in an external directory

If you do not manage your users in the Confluence internal directory, you may still be able to change a user's username. Confluence cannot **update** external users but it will detect changes in usernames that are coming from some external directories.

The following table indicates 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 the user's original username. For example if jsmith's username is changed to jbrown, their personal space key will remain ~jsmith.

Restoring Passwords To Recover Admin User Rights

Use this document if you are unable to log in to Confluence as administrator. The most common reason for using these instructions is if you have lost the administration password for your Confluence site.

Before you Start

Please note the following 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.

New user management in Confluence 3.5 and later

- Confluence now uses the CWD_USER table in the database to store and refer to its users.
- During an upgrade from Confluence 3.4.9 or earlier, the upgrade process copied the users from the OS_USER table (for upgrades from versions older than 2.7) or the USERS table (for versions 2.7 to 3.4) into the CWD_USER table.

- The new user management framework also introduced user directories. Making modifications to users in the database will only fully work for users in Confluence's Internal Directory. The instructions below include extra steps for instances in which the user management has been delegated to external sources (via LDAP, Crowd or JIRA).

Please refer to the older documentation if you are still using [OSUser](#) or [AtlassianUser](#).

Using Crowd for SSO

- If Confluence is configured for SSO through Crowd, you will only be able to authenticate as users from the Crowd server.
- This document covers how to recover administration rights from the local 'Confluence Internal Directory' only. However, you will not be able to authenticate as a local Confluence administrator while Crowd SSO is enabled. Please refer to [Integrating Crowd with Atlassian Confluence](#) for details on how to configure or disable Crowd SSO.

On this page:

- [Before you Start](#)
- [Step 0. Get access to the database](#)
- [Step 1. Identify Administrator](#)
- [Step 2. Replace Administrator Password](#)
- [Step 3. Put the Internal Directory in First Position](#)
- [Step 4. Clean Up](#)
- [Notes](#)

Step 0. Get access to the database

If you are using the embedded HSQL database, you can find the files containing your database in `<confluence-home-directory>/database`. When you shut down Confluence, the SQL will be written to a '.script' or '.log' file in that directory to which you can append the SQL described below.

If you are using a proper production database, connect to the database with your normal tools. You will 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](#). Please download a database admin tool now if you do not 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 is 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 user name 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+XURPwX10Nx
p3Y3pB37A==');

insert into user_mapping values ('2c9681954172cf560000000000000001', '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', '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'));
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 using an Oracle database, use **sysdate** instead of a string for the **created_date** column.

Step 2. Replace Administrator Password

Confluence does not store passwords in plain text in the database, but uses hashes computed from the original password. You will 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`

```
x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPwX10Nx
p3Y3pB37A==
```

For an External Database**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 =  
'x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPwx10Nx  
p3Y3pB37A=='  
where id=<id from Stage 1>;
```

For the Evaluation Embedded HSQL Database**To change the password to admin for a given username:**

1. Shut down Confluence.
2. Open <confluence-home>/database/confluencedb.script, or confluencedb.log if the .script file looks empty.
3. Search for:

```
INSERT INTO CWD_USER VALUES(
```

4. Keep searching until you find the appropriate user, then replace their password with the hash value above.
5. Save the file.
6. Restart Confluence.

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 does not 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. **Do not leave your password as admin, or your instance will not 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 Confluence Administration > User Directories and rearrange your directories so they are correctly configured again.

Notes

- Learn more about the [password hash algorithm Confluence is using](#).

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. For example, it is quicker to give group 'X' access to Confluence, rather than giving every team member access individually. You can also set the access levels for [anonymous users](#).

Related pages:

- [Confluence Administrator's Guide](#)
- [Confluence Security Overview and Advisories](#)

Global Groups Overview

A group is a collection of users. Administrators create groups so that the administrator can assign permissions to a number of people at once. For example, it is quicker to give group 'X' access to Confluence, rather than giving every team member access individually. You need [Confluence Administrator permissions](#) to view and update groups.

Groups are available at the space and page levels to allow for flexible access control. A user in a group will automatically be granted all permissions granted to the group.

Special groups

There are two special default groups in Confluence:

1. **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.
2. **confluence-users**: This is the default group for all new users. Permissions you assign to this group will be assigned to all newly signed-up users of Confluence.

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.

On this page:

- [Special groups](#)
- [Anonymous users](#)
- [Updating groups](#)
- [Notes](#)

Related pages:

- [Managing Confluence Users](#)
- [Global Permissions Overview](#)
- [Confluence Administrator's Guide](#)

 *The information on this page **does not apply** to Atlassian OnDemand sites with multiple apps. If you are using Confluence-only OnDemand, the information does apply.*

Anonymous users

Confluence treats all users who do not log in when they access Confluence as being 'anonymous'. You can grant anonymous 'Use Confluence' permission via the Global Permissions screen. See [Setting Up Public Access](#). 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](#).

Updating 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 are 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.

Notes

- **Multiple user directories:** You may define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you may use the default Confluence **internal directory** and also connect to an **LDAP** directory server. In such cases, 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 [Users and Groups](#) and [Managing Confluence Users](#).

You can edit group membership in two places:

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

Both methods are described below.

Adding and removing members via the group management screen

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. The 'Groups' screen appears, showing a list of groups. Choose the group to which you want to add users.
4. The 'Group Members' screen appears, showing the users who belong to the selected group. 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, as described in [Searching for Users](#).
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. The 'Manage Groups' screen appears, showing a list of groups. Choose the group from which you want to remove the user.
4. The 'Group Members' screen appears, showing the users who belong to the selected group. Choose the 'Delete user from group' icon next to the user whose group membership you want to remove.

On this page:

- [Adding and removing members via the group management screen](#)
- [Editing group membership from the user management screen](#)
- [Notes](#)

Related pages:

- [Managing Confluence Users](#)
- [Global Permissions Overview](#)
- [Confluence Administrator's Guide](#)

 *The information on this page **does not apply** to Atlassian OnDemand sites with multiple apps. If you are using Confluence-only OnDemand, the information does apply.*

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](#)

Editing group membership from the user management screen

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:

- 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.
 - Choose **Users** in the left-hand panel.
 - The 'Users' screen appears. You can now choose to 'Show all users' or you can search for a specific user by entering all or part of the person's username, full name or email address. For more details about the user search, see [Searching For and Administering Users](#).
 - Choose the username you want to edit.
- The 'View User' screen appears. Choose **Edit Groups**.
- 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](#)

Notes

You may define multiple user directories in Confluence, so that Confluence looks in more than one place for its users and groups. For example, you may use the default Confluence **internal directory** and also connect to an **LDAP** directory server. In such cases, 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.

Overview of the 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.
Update User Status	This allows the user to update their user status message, which can be seen on the user's profile, pages in their personal space and on various activity streams accessible to other Confluence users.
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: • Server Base URL • Public Signup • Connection Timeouts
Further Configuration	The following functionality is disallowed: • Remote API plugin
Security Configuration	The following functionality is disallowed: • 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: • 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**.
5. 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.)

1. First [add the user](#) to Confluence, if you have not already done so.
2. Choose **Edit Permissions**. The 'Edit Global Permissions' screen appears.
3. Enter the username in the **Grant browse permission to** box in the 'Individual Users' section. You can search for the username.
4. Choose **Add**.
5. 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:

1. 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.
2. To allow anonymous access to your Confluence site, select the 'Use Confluence' and 'View User Profile' options in the 'Anonymous Access' section.
For more information about these permissions, refer to [Setting Up Public Access](#).
3. Choose **Save All** to save your changes.

Screenshot: Editing global permissions

Edit Global Permissions

You can edit the current space permissions here. Permissions can be granted to specific users or groups. You can also grant permissions to anonymous users. Anonymous users refers to all users that are not logged in. For information about the individual *permissions* that can be assigned, please see the [permissions guide](#).

Groups

These are the global permissions currently assigned to groups.

	Attach Files to User Profile [?]	Update User Status [?]	Personal Space [?]	Create Space(s) [?]	Confluence Administrator [?]	System Administrator [?]
confluence-administrators	☒ can use	☐	☒	☒	☒	☒
confluence-users	☒ can use	☐	☒	☒	☐	☐

Group test-user could not be found. Please confirm the existence of the group with a Confluence Administrator.

Grant browse permission to

Individual Users

These are the global permissions currently assigned to individual users.

	Attach Files to User Profile [?]	Update User Status [?]	Personal Space [?]	Create Space(s) [?]	Confluence Administrator [?]	System Administrator [?]
Rach Admin (admin)	☒ can use	☒	☒	☒	☒	☐

Grant browse permission to

Anonymous Access

When a user is using Confluence while not logged in, they are using it anonymously.
For example: Enabling anonymous 'USE' permission, allows non-logged-in users to browse pages and spaces in Confluence.

	Use Confluence [?]	View User Profiles [?]
Anonymous	☐ can use	☐

[Cancel](#)

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**.

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](#)
- [Adding and Inviting Users](#)
- [Global Permissions Overview](#)
- [Confluence Administrator's Guide](#)

 *The information on this page does not apply to Atlassian OnDemand sites with multiple apps. If you are using Confluence-only OnDemand, the information does apply.*

Disabling anonymous access to the site

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 [Adding and Inviting 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** as directory managers.

On this page:

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

 *The information on this page **does not apply** to Confluence OnDemand.*

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.

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. See [Connecting to Crowd or JIRA for User Management](#).
- Atlassian JIRA 4.3 or later. See [Connecting Confluence to JIRA for User Management](#).
- Atlassian JIRA 4.2 or earlier, using the legacy database connection. See [Connecting to JIRA 4.2 or Earlier 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

RELATED TOPICS

- [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](#)
- [Connecting to JIRA 4.2 or Earlier 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](#)
- [Adding and Inviting Users](#)
- [Managing Site-Wide Permissions and Groups](#)

[Administrators Guide Home](#) [Confluence Documentation Home](#)

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.

On this page:

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

 *The information on this page **does not** apply to Confluence OnDemand.*

Diagram of Possible Configuration




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

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](#)
- [Connecting to JIRA 4.2 or Earlier 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](#)

[How to Reenable the Internal Directory](#) (Knowledge base article)

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:

- 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.

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:
 - 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](#).

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](#)
- [Notes](#)

Related pages:[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](#)
- [Connecting to JIRA 4.2 or Earlier 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](#)

 *The information on this page does not apply to Confluence OnDemand.*

Server Settings

Setting	Description
Name	Enter a meaningful name to help you identify the LDAP directory server. Examples: • 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: • Microsoft Active Directory • OpenDS • And more.

Hostname	The host name of your directory server. Examples: • <code>ad.example.com</code> • <code>ldap.example.com</code> • <code>opens.example.com</code>
Port	The port on which your directory server is listening. Examples: • 389 • 10389 • 636 (for example, for SSL)
Use SSL	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.
Username	The distinguished name of the user that the application will use when connecting to the directory server. Examples: • <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>
Password	The password of the user specified above.

Schema Settings

Setting	Description
Base DN	The root distinguished name (DN) to use when running queries against the directory server. Examples: • <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	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: • <code>ou=Users</code>
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: • <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	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. i 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
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	<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. <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. <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. 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. Examples: • confluence-users • confluence-users,jira-users,jira-developers

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.
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 <code>'javax.naming.PartialResultException: Unprocessed Continuation Reference(s)'</code> error.
Naive DN Matching	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. This setting determines how your application will compare DNs to determine if they are equal. • 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 DNs. • If this checkbox is not selected, the application will parse the DN and then check the parsed version.
Enable Incremental Synchronisation	Enable incremental synchronisation if you only want changes since the last synchronisation to be queried when synchronising a directory.  Please be aware that when using this option, the user account configured for synchronisation must have read access to: • 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). 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 JIRA. This setting is only available if the directory type is set to "Microsoft Active Directory".
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)	This setting affects two actions. The default value is 0. • 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	This is the name of the class used for the LDAP user object. Example: • user
User Object Filter	The filter to use when searching user objects. Example: • (&(objectCategory=Person)(sAMAccountName=*)) More examples can be found here and here.
User Name Attribute	The attribute field to use when loading the username. Examples: • cn • sAMAccountName NB: In Active Directory, the 'sAMAccountName' is the 'User Logon Name (pre-Windows 2000)' field. The User Logon Name field is referenced by 'cn'.
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: • cn
User First Name Attribute	The attribute field to use when loading the user's first name. Example: • givenName
User Last Name Attribute	The attribute field to use when loading the user's last name. Example: • sn
User Display Name Attribute	The attribute field to use when loading the user's full name. Example: • displayName

User Email Attribute	The attribute field to use when loading the user's email address. Example: mail
User Password Attribute	The attribute field to use when loading a user's password. Example: unicodePwd
User Unique ID Attribute	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. 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.

Group Schema Settings

Setting	Description
Group Object Class	This is the name of the class used for the LDAP group object. Examples: - groupOfUniqueNames - group
Group Object Filter	The filter to use when searching group objects. Example: (&(objectClass=group)(cn=*))
Group Name Attribute	The attribute field to use when loading the group's name. Example: cn
Group Description Attribute	The attribute field to use when loading the group's description. Example: description

Membership Schema Settings

Setting	Description
Group Members Attribute	The attribute field to use when loading the group's members. Example: member
User Membership Attribute	The attribute field to use when loading the user's groups. Example: memberOf

Use the User Membership Attribute, when finding the user's group membership	Check this if your directory server supports the group membership attribute on the user. (By default, this is the 'memberOf' attribute.) • 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	Check this if your directory server supports the user membership attribute on the group. (By default, this is the 'member' attribute.) • 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.

Notes

Currently there is a bug which causes a system error if the username and password are not correct. This also happens if you are accessing anonymously, but the directory server does not support anonymous access. If you get a system error message, try checking the username and password credentials. You can watch this issue to see updates on this bug: [CONF-25961](#) - NPE when defining LDAP directory and having wrong password RESOLVED

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.

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: - 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: - 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.

RELATED TOPICS

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

 [Administrators Guide Home](#)  [Confluence Documentation Home](#)

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**.



The screenshot shows the 'Specify Setup Type' page of the 'Add Roles Wizard'. The left sidebar contains a tree view with the following items: 'Before You Begin', 'Server Roles', 'AD CS' (expanded), 'Role Services', 'Setup Type' (selected), 'CA Type', 'Private Key', 'Cryptography', 'CA Name', 'Validity Period', 'Certificate Database', 'Confirmation', 'Progress', and 'Results'. The main content area has a title 'Specify Setup Type' with a green plus icon. Below the title, there is a paragraph: '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.' There are two radio button options: 'Enterprise' (selected) and 'Standalone'. The 'Enterprise' option has a description: 'Select this option if this CA is a member of a domain and can use Directory Service to issue and manage certificates.' The 'Standalone' option has a description: '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.' At the bottom of the main area, there is a link: '[More about the differences between enterprise and standalone setup](#)'. At the bottom of the wizard, there are four buttons: '< Previous', 'Next >', 'Install', and 'Cancel'.

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



The screenshot shows the 'Specify CA Type' page of the 'Add Roles Wizard'. The left sidebar contains a tree view with the following items: 'Before You Begin', 'Server Roles', 'AD CS' (expanded), 'Role Services', 'Setup Type', 'CA Type' (selected), 'Private Key', 'Cryptography', 'CA Name', 'Validity Period', 'Certificate Database', 'Confirmation', 'Progress', and 'Results'. The main content area has a title 'Specify CA Type' with a green plus icon. Below the title, there is a paragraph: '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.' There are two radio button options: 'Root CA' (selected) and 'Subordinate CA'. The 'Root CA' option has a description: 'Select this option if you are installing the first or only certification authority in a public key infrastructure.' The 'Subordinate CA' option has a 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, there is a link: '[More about public key infrastructure \(PKI\)](#)'. At the bottom of the wizard, there 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

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

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**.



The screenshot shows the 'Add Roles Wizard' window with the title bar 'Add Roles Wizard' and a close button. The main window has a left-hand navigation pane and a main content area. The navigation pane lists the following steps: 'Before You Begin', 'Server Roles', 'AD CS', 'Role Services', 'Setup Type', 'CA Type', 'Private Key', 'Cryptography', 'CA Name' (which is highlighted with a blue background), 'Validity Period', 'Certificate Database', 'Confirmation', 'Progress', and 'Results'. The main content area is titled 'Configure CA Name' and contains the following text: '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.' Below this text are three text input fields. The first field is labeled 'Common name for this CA:' and contains the text 'dc-one-DC-ONE-CA'. The second field is labeled 'Distinguished name suffix:' and contains the text 'DC=dc-one,DC=crowd,DC=sydney,DC=atlassian,DC=com'. The third field is labeled 'Preview of distinguished name:' and contains the text 'CN=dc-one-DC-ONE-CA,DC=dc-one,DC=crowd,DC=sydney,DC=atlassian,DC=com'. At the bottom of the main content area is a blue hyperlink that reads 'More about configuring a CA name'. At the bottom of the window are four buttons: '< Previous', 'Next >', 'Install', and 'Cancel'.

Add Roles Wizard

Configure CA Name

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

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
```

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

```
keytool -import -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 use the '**Secure SSL**' option when connecting your application to your directory server.

UNIX

1. Navigate to the directory in which Java is installed. cd \$JAVA_HOME will usually get you there.
2. Run the command below, where server-certificate.crt is the name of the file from your directory

server:

```
sudo keytool -import -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 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 `/Library/Java/Home`.
2. Run the command below, where `server-certificate.crt` is the name of the file from your directory server:

```
sudo keytool -import -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 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](#)

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:
 - 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](#).
7. Add your users and groups in Confluence. See [Adding and Inviting 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: • 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: • Microsoft Active Directory • OpenDS • And more.
Hostname	The host name of your directory server. Examples: • ad.example.com • ldap.example.com • opens.example.com
Port	The port on which your directory server is listening. Examples: • 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: • 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	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 you check this box the following additional fields will appear on the screen, which are described in more detail below: • Default Group Memberships • Synchronise Group Memberships • User Schema Settings (described in a separate section below)
Default Group Memberships	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. 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. Examples: • confluence-users • bamboo-users, jira-users, jira-developers
Synchronise Group Memberships	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. If you check this box the following additional fields will appear on the screen, both described in more detail below: • 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	The root distinguished name (DN) to use when running queries against the directory server. Examples: • o=example,c=com • cn=users,dc=ad,dc=example,dc=com • For Microsoft Active Directory, specify the base DN in the following format: dc=domain1,dc=local. You will need to replace the domain1 and local 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	The attribute field to use when loading the username. Examples: • cn • sAMAccountName

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.
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.

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: • ou=Users

User Object Class	This is the name of the class used for the LDAP user object. Example: • user
User Object Filter	The filter to use when searching user objects. Example: • (&(objectCategory=Person)(sAMAccountName=*))
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: • cn
User First Name Attribute	The attribute field to use when loading the user's first name. Example: • givenName
User Last Name Attribute	The attribute field to use when loading the user's last name. Example: • sn
User Display Name Attribute	The attribute field to use when loading the user's full name. Example: • displayName
User Email Attribute	The attribute field to use when loading the user's email address. Example: • mail

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: • ou=Groups
Group Object Class	This is the name of the class used for the LDAP group object. Examples: • groupOfUniqueNames • group
Group Object Filter	The filter to use when searching group objects. Example: • (objectCategory=Group)

Group Name Attribute	The attribute field to use when loading the group's name. Example: cn
Group Description Attribute	The attribute field to use when loading the group's description. Example: description

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: member
User Membership Attribute	The attribute field to use when loading the user's groups. Example: memberOf
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.) - 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.

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](#)
- [Connecting to JIRA 4.2 or Earlier 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](#)

[Administrators Guide Home](#) [Confluence Documentation Home](#)

Connecting to Crowd or JIRA for User Management

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

On this page:

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

 The information on this page does not apply to Confluence OnDemand.

Connecting Confluence to Crowd for User Management

Atlassian Crowd is an application security framework that handles authentication and authorisation for your web-based applications. With Crowd you can integrate multiple web applications and user directories, with support for single sign-on (SSO) and centralised 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.

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:
 - 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](#).
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: • Crowd Server • Example Company Crowd
Server URL	The web address of your Crowd console server. Examples: • http://www.example.com:8095/crowd/ • http://crowd.example.com
Application Name	The name of your application, as recognised 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 .

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 for User Management

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

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

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

If you are running JIRA 4.2 or earlier, please see [Connecting to JIRA 4.2 or Earlier for User Management](#).

To connect Confluence to JIRA 4.3 or later:

- Go to your **JIRA administration screen** and define the Confluence application to JIRA:
 - For JIRA 4.3.x, select '**Other Applications**' from the '**Users, Groups & Roles**' section of the 'Administration' menu.
 - For JIRA 4.4 or later, select '**Users**' > '**JIRA User Server**' in Administration mode.
 - 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 JIRA:
 - a. Add the confluence-users and confluence-administrators groups in JIRA.
 - 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 JIRA, 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	A meaningful name that will help you to identify this JIRA server amongst your list of directory servers. Examples: • JIRA Server • My Company JIRA
Server URL	The web address of your JIRA server. Examples: • http://www.example.com:8080 • http://jira.example.com
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 4.3 or later is running on the given url	This may be because: • 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: • 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 JIRA 4.3 or later.	A runtime exception has occurred. Refer to the Confluence log files for further troubleshooting information.

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](#)
- [Connecting to JIRA 4.2 or Earlier 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](#)

 [Administrators Guide Home](#)  [Confluence Documentation Home](#)

Reverting from Crowd or JIRA to Internal User Management

If your Confluence site currently uses JIRA or Crowd 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.



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.

On this page:

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

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 JIRA, you must create a group in Confluence with the same name.
 - If a user who exists in JIRA 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 Users and Groups to the Confluence Database

Use this option to migrate External Application (Crowd or JIRA) 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 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 [Restoring 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.

RELATED TOPICS

Configuring User Directories

Connecting to JIRA 4.2 or Earlier for User Management

Atlassian JIRA is an issue and project tracking tool. Like Confluence, JIRA offers the ability to store its users and groups in its database. You can configure Confluence to look for its users and groups in the JIRA database. This page describes the **legacy JIRA database connector**, which provides a direct connection to the JIRA database.

When to use this option: Choose the legacy JIRA database connector if your JIRA server is **JIRA 4.2 or earlier**, for backwards compatibility with the already-existing option for Confluence to use JIRA for user management.

*If you are using **JIRA 4.3 or later**, you cannot use the legacy JIRA database connector. Instead, choose the 'Atlassian JIRA' directory type.*

On this page:

- [Connecting Confluence to JIRA](#)
- [JIRA Settings in Confluence](#)

Connecting Confluence to JIRA

To connect Confluence to JIRA 4.2 or earlier:

1. Edit the Confluence server .xml file, to construct the datasource location, as described below.
2. Restart Confluence.
3. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
4. Click **User Directories** in the left-hand panel.
5. **Add** a directory and select type **Legacy Atlassian JIRA (4.2 and earlier)**. Enter the settings as described below.
6. Save the directory settings.
7. 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](#).
8. 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 JIRA:
 - a. Add the confluence-users and confluence-administrators groups in JIRA.
 - 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 JIRA, 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.

JIRA Settings in Confluence

Setting	Description

Name	A meaningful name that will help you to identify this JIRA server amongst your list of directory servers. Examples: • JIRA • Example Company JIRA
Datasource Location	The JNDI name of the JIRA datasource configured in your application server. Example: java:comp/env/jdbc/YourJiraDatasource In JIRA standalone distributions (using the default application server, Tomcat 6) you can construct the datasource location as follows: 1. Open your <jira_install>/conf/server.xml file in a text editor. 2. Look for the database setup section in that file. It looks something like this: <pre><Resource auth="Container" driverClassName="com.mysql.jdbc. Driver" maxActive="20" name="*jdbc/JiraDS*" password="jirauser" type="javax.sql.DataSource" url="jdbc:mysql://localhost/jira db?useUnicode=true&characterEnco ding=UTF8" username="jirauser" validationQuery="select 1"/></pre> 3. Copy the above lines (the 'Resource' section) and paste it to your Confluence's server.xml file (located at <confluence_install>/conf/server.xml), under the Context path. This will then expose the value of the name attribute as the JNDI resource locator. 4. Copy the JNDI name from the name parameter. In this example, the datasource location is: java:comp/env/jdbc/JiraDS

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](#)
- [Connecting to JIRA 4.2 or Earlier 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 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.

On this page:

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

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

The directory order is significant when granting the user permissions based on group membership. 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](#)
- [Connecting to JIRA 4.2 or Earlier 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](#)



[Administrators Guide Home](#)



[Confluence Documentation Home](#)

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.

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](#)

 *The information on this page does not apply to Confluence OnDemand.*

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 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.

- 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 JIRA 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
```

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](#)
- [Connecting to JIRA 4.2 or Earlier 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](#)

[Administrators Guide Home](#)[Confluence Documentation Home](#)

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](#)

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.

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](#)
- [Connecting to JIRA 4.2 or Earlier 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

[Administrators Guide Home](#) [Confluence Documentation Home](#)

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](#)

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.

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](#)
- [Connecting to JIRA 4.2 or Earlier 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](#)

[Administrators Guide Home](#) [Confluence Documentation Home](#)

User Management Limitations and Recommendations

This page describes the optimal configurations and limitations that apply to user management 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](#)
 - [Optimal Number of Users and Applications](#)
 - [JIRA OnDemand not supported](#)
 - [Recommendations](#)

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.

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 JIRA 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, 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.

Optimal Number of Users and Applications

Please consider the following limitations when connecting to a JIRA server for user management:

- Maximum 500 users.
- Maximum 5 connected applications.

JIRA OnDemand not supported

You cannot use JIRA OnDemand to manage FishEye standalone users. OnDemand users and your FishEye users need to be managed separately.

Recommendations

Your environment	Recommendation
If all the following are true: • You have fewer than 500 users. • You want to share user and group management across just a few applications, such as one JIRA server and one Confluence server, or two JIRA servers. • You do not need single sign-on (SSO) between JIRA 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 JIRA.	Your environment meets the optimal requirements for using JIRA for user management.
If one or more of the following are true: • You have more than 500 users. • 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.	We recommend that you install Atlassian Crowd for user management and SSO.

If you are considering creating a custom directory connector to define your own storage for users and groups...

Please see if one of the following solutions will work for you:

- 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](#).

RELATED TOPICS

[Connecting to an LDAP Directory](#)

[Connecting to Crowd or JIRA for User Management](#)

[Configuring User Directories](#)

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 Atlassian JIRA 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.

On this page:

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

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** , 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 or exploded WAR directory, then zip the full `/logs` subdirectory into a single file for us to examine. The logs

subdirectory is located in your Confluence Home directory.

RELATED TOPICS

Troubleshooting Problems and Requesting Technical Support 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](#)
- [Connecting to JIRA 4.2 or Earlier 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](#)

 [Administrators Guide Home](#)  [Confluence Documentation Home](#)

Disabling the Built-In User Management

By selecting the '**External user management**' option in Confluence, you can disable the group and user management screens in Confluence. You need [system administrator permissions](#) to set this option.

You will find it useful to select external user management under the following circumstances:

- When [Crowd's directory permissions](#) are configured so that Confluence cannot update the Crowd directories, then Confluence's external user management setting **must** be turned on. Otherwise, a 'System Error' will occur when Confluence attempts to write data into Crowd. For more information about integrating Crowd with Confluence, see [Connecting to Crowd or JIRA for User Management](#).
- If you are using JIRA for user management, we recommend that you turn on Confluence's external user management setting. This centralises user management in JIRA. See [Connecting to Crowd or JIRA for User Management](#) and [Connecting to JIRA 4.2 or Earlier for User Management](#).

To disable management of users and groups within Confluence:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Click '**Security Configuration**' in the left-hand panel.
3. The '**Edit Security Configuration**' screen will appear. Click '**Edit**'.
4. Tick the '**External user management**' check box.
5. Click '**Save**'.

Notes

- Please refer to the following bugs and improvement requests:
 - [CONF-16709](#) – When the External User Management check box is ticked, the group and user management screens are still functional.
 - [CONF-21158](#) – Enabling both public signup and external user management renders a blank screen during signup.
 - [CONF-9830](#) – This is a request to rename this feature to better reflect its functionality.

RELATED TOPICS

Content by label

There is no content with the specified labels



 [Administrators Guide Home](#)  [Confluence Documentation Home](#)

Managing Add-ons and Macros

An **add-on** is a separately installed component that enhances or modifies Confluence. Some add-ons are shipped with Confluence, others are available for you to install yourself. An add-on that specifically plugs into the architecture of an Atlassian application such as Confluence is sometimes called a **plugin**, although the terms 'plugin' and 'add-on' are often used interchangeably.

A **macro** allows a developer to perform programmatic functions within a page, and gives the Confluence user access to more complex content structures. Many macros are made available by plugins.

You need [System Administrator](#) permissions in order to install and configure plugins.

Installing and configuring add-ons and macros

- [About Add-ons](#)
 - [Add-on loading strategies in Confluence](#)
 - [Removing Malfunctioning Add-ons](#)
 - [Enabling and Configuring Macros](#)
 - [Configuring the User List Macro](#)
 - [Enabling HTML macros](#)
 - [Enabling the html-include Macro](#)
 - [Troubleshooting the Gallery Macro](#)
 - [Adding, Editing and Removing User Macros](#)
 - [Writing User Macros](#)
 - [Best Practices for Writing User Macros](#)
 - [Examples of User Macros](#)
 - [Hello World Example of User Macro](#)
 - [Error Box Macro - Example of a User Macro](#)
 - [Colour and Size Macro - Example of a User Macro](#)
 - [NoPrint Example of a User Macro](#)
 - [Panel Preformatted with Specific Colours - Example of a User Macro](#)
 - [Preformatted Table - Example of a User Macro](#)
 - [Guide to User Macro Templates](#)
- [Configuring the Office Connector](#)

About Add-ons

An add-on is an installable component that supplements or enhances the functionality of Confluence in some way. For example, the [Team Calendars for Confluence](#) is an add-on that lets users embed team calendars into Confluence pages. Other Confluence add-ons are available for creating charts, tracking usage and modifying the Confluence visual theme.

Confluence comes with many pre-installed add-ons (called system add-ons). You can install more add-ons either by acquiring an add-on from the [Atlassian Marketplace](#) or by uploading an add-on from your file system. This means that you can install add-ons that you have developed yourself. For information about developing your own add-ons for Confluence, see the [Confluence Developer documentation](#).

About the Universal Plugin Manager

The Universal Plugin Manager (UPM) is itself an add-on that you use to administer add-ons from the Confluence Administration Console. UPM works across Atlassian applications, providing a consistent interface for administering add-ons in Confluence, Crucible, Fisheye, JIRA, Stash or Bamboo.

UPM comes pre-installed in recent versions of all Atlassian applications, so you do not normally need to install it yourself. However, like other add-ons, the UPM software is subject to regular software updates. Before administering add-ons in Confluence, therefore, you should [verify your version](#) of the UPM and update it if needed.

Administering Add-ons in Confluence

You can update UPM, or any add-on, from the UPM's own add-on administration pages. Additionally, you can perform these tasks from the administration pages:

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

If the add-on request feature is enabled in your Atlassian application, non-administrative users can also discover add-ons on the Atlassian Marketplace. Instead of installing the add-ons, however, the users have the option of requesting the add-ons from you, the administrator of the Atlassian application. For an end-user's view of the add-on request feature in Confluence, see [Requesting Add-ons](#).

For more information on administering this feature and performing other add-on administration tasks, see the [Universal Plugin Manager documentation](#).

For add-on information specific to Confluence, see these pages:

- [Add-on loading strategies in Confluence](#)
- [Removing Malfunctioning Add-ons](#)

Related pages:

- [Confluence Plugin Guide for Developers](#)
- [Adding, Editing and Removing User Macros](#)

Add-on loading strategies in Confluence

The categories

Confluence add-ons have different behaviour based on how they are loaded by Confluence. The add-ons themselves are the same, but based on how they are loaded, they may or may not be upgraded, or may not be disabled, or may not be uninstalled. This chart should explain how plugins can be loaded by Confluence, and the ramifications for each choice.

The category any *particular* add-on is in can vary with Confluence version or circumstance. The examples mentioned here describe the way particular add-ons are loaded by default in Confluence 2.8.

 The information on this page *does not apply* to Confluence OnDemand.

Category	Description	Example
Static	cannot be installed or upgraded without a Confluence restart	
Core	Included with Confluence and cannot be uninstalled. The classes and plugin.xml are not bundled into add-on distribution JAR files, but are mixed in with Confluence source on the main classpath. Additionally, the plugin.xml definitions are not called "atlassian-plugin.xml" as they are everywhere else, but are named for the add-on, e.g., "basic-macros.xml". We would like to separate some of them out and turn them into <i>Bundled</i> add-ons.	Admin Sections

<i>WEB-INF/lib</i>	Confluence also places some add-on JAR files inside WEB-INF/lib . They are inserted during the build process by Maven. These add-ons, likewise, cannot be uninstalled. In ancient times, this was the only way to install add-ons, so users were also free to install add-ons here. We now discourage this installation method, however. As of version 3.0, most of the JAR files in this directory are library dependencies, not add-on files.	
Dynamic	the opposite of static, these can be installed/upgraded while Confluence is running	
<i>Bundled</i>	Bundled add-ons can be administered from the Manage Add-ons page in the application's Administration Console. You can upload or disable them there. <i>Bundled</i> add-ons are included in a ZIP archive of JAR files called atlassian-bundled-plugins.zip, which is on the main Confluence classpath, in a resources directory - <code><confluence-install>/confluence/WEB-INF/atlassian-bundled-plugins</code>. At Confluence startup, they are extracted and copied into the \$CONFLUENCE_HOME/bundled-plugins directory, from whence they are loaded. To remove a bundled add-on (you shouldn't normally have to do this), remove the add-on from the <code>atlassian-bundled-plugins.zip</code> file and the <code>bundled-plugins</code> directory, otherwise Confluence will just put it back in place on the next startup. In versions later than 2.6, you'll have to recreate the .jar file (if the jar file is from the lib folder) or recreate the zip folder (if its in the classes folder). Bundled add-ons can be upgraded or disabled.	Office Connector
<i>Uploaded</i>	Installed by the user via the plugin repository or the Universal Plugin Manager. These add-ons are stored in the database and then copied to the \$CONFLUENCE_HOME/plugins-cache folder on each Confluence node.	could be anything

To summarise the relationships of categories in the table, all add-ons are either *Static* or *Dynamic*. *Static* add-ons can be further categorised into *Core* or *WEB-INF/lib*. *Dynamic* add-ons are divided into *Bundled* and *Uploaded*.

Use of the categories in Confluence

Within Confluence, the *Core* and *WEB-INF/lib* categories are not actually named as such, and they don't map neatly to other names (though they do map, as will be explained). They are used here because of the logical distinction they provide.

In Confluence, some of the *Core* add-ons are called "System Add-ons". Add-ons can be designated as "System" by adding a flag to the add-on manifest file. To do this, `system=true` should be added to the top-level `atlassian-plugin` element of the manifest file. The manifest file is generally called `atlassian-plugin.xml`, but it could have another name; the *Core* add-on files do.

All the *Core* add-ons were once labeled "System", but the convention has faded over time. If an add-on is designated as "System", it cannot be enabled/disabled in the Manage Add-ons page. However, it *will* show up in the Plugin Repository Client, where it can be disabled; allowing disabling there is probably incorrect behavior.

Static add-ons that are not marked as "System" (any remaining *Core* and *WEB-INF/lib* plugins), are simply called *Static* in Confluence. There is no way to tell the *WEB-INF/lib* and *Core* add-ons apart from within Confluence. You just have to figure out where the classes are.

Members of the other specific categories—*Bundled* and *Uploaded*—can be determined. We can tell which add-ons are *Bundled* and which add-ons are *Uploaded*, so we know which add-ons are *Uploaded* though this specific term is never used in the Confluence UI. Instead, they are called *Dynamic*.

Updating add-on versions

- *Core* add-ons cannot be upgraded.
- *WEB-INF/lib* add-ons can be upgraded by replacing the JAR in *WEB-INF/lib* and restarting Confluence.
- *Bundled* add-ons can be upgraded using the Universal Plugin Manager or from the Plugin Repository Client. A new add-on JAR is uploaded and stored as an *Uploaded* add-on. Confluence compares the version number with the *Bundled* add-on and uses the newer.
- *Uploaded* add-ons are upgradable using the Universal Plugin Manager or from the Plugin Repository Client. When a new add-on JAR file is uploaded, the previous version is discarded from the database and the `$CONFLUENCE_HOME/plugin-cache` directory.

RELATED TOPICS

Removing Malfunctioning Add-ons

Removing Malfunctioning Add-ons

Confluence goes to some lengths to prevent itself being unusable due to a problematic add-on. However, sometimes an add-on will manage to do this anyway. This page describes what to do if an add-on cannot be disabled or deleted from the administration console (from **Administration > Manage Add-ons**).

 The information on this page *does not apply* to Confluence OnDemand.

Add-on Loading Strategies

1. Read through [Plugin loading strategies in Confluence](#).
2. Determine where your add-on file is located. The usual locations are:
 - a. The `PLUGINDATA` table on the database
 - b. The `<confluence-home>/bundled-plugins` folder
 - c. The `<confluence-home>/plugin-cache` folder
 - d. The `<confluence-home>/plugins-osgi-cache` folder
 - e. The `<confluence-home>/plugins-temp` folder
 - f. The `<confluence-install>/confluence/WEB-INF/lib` folder (deprecated approach)

Check these locations when troubleshooting add-on loading issues.

 Check the [How do I display the Confluence System Classpath?](#) FAQ and the [Knowledge Base Article on plugin loading problems](#) for more information on troubleshooting add-on loading issues.

Deleting an add-on from the Database

To remove an add-on from Confluence when Confluence is not running,

1. Connect to the Confluence database.
2. Run the following SQL statement in your database:

```
select pluginid, pluginkey, filename, lastmoddate from
PLUGINDATA;
```

3. After you have found the pluginid value for the offending add-on, run the following:

```
delete from PLUGINDATA where pluginid='XXXXXX';
```

where XXXXXX is the pluginid value.

4. Restart Confluence.

Disabling an add-on from the database

To disable the add-on in the database,

Run the following query on your Confluence database:

```
select BANDANAVALUE from BANDANA where BANDANAKEY =
'plugin.manager.state.Map'
```

This will return a value like:

```
<map>
  <entry>
    <string>com.atlassian.confluence.ext.usage</string>
    <boolean>true</boolean>
  </entry>
</map>
```

Edit the value boolean to have **false**:

```
<map>
  <entry>
    <string>com.atlassian.confluence.ext.usage</string>
    <boolean>>false</boolean>
  </entry>
</map>
```

Deleting a Bundled Add-on

Bundled add-ons can be administered from the Manage Add-ons page in the application's Administration Console. You can upload or disable them there.

Bundled add-ons are included in a ZIP archive of JAR files called **atlassian-bundled-plugins.zip**, which is on the main Confluence classpath, in a resources directory - <confluence-install>/confluence/WEB-INF/atlassian-bundled-plugins. At Confluence startup, they are extracted and copied into the **\$CONFLUENCE_HOME/bundled-plugins** directory, from whence they are loaded. To remove a bundled add-on (you shouldn't normally have to do this), remove the add-on from the atlassian-bundled-plugins.zip file and the bundled-plugins directory, otherwise Confluence will just put it back in place on the next startup. In versions later than 2.6, you'll have to [recreate the .jar file](#) (if the jar file is from the lib folder) or recreate the zip folder (if its in the classes folder). Bundled add-ons can be upgraded or disabled.

If you need to remove a bundled add-on, check to see if you have duplicates in the <confluence-home>/bundled-plugins or <confluence-home>/plugin-cache directory.

Usually, the problem is that an old add-on is getting loaded along with the properly bundled one, but if you need to remove a bundled add-on, check [Add-on loading strategies in Confluence](#).

Enabling and Configuring Macros

Macros allow you to perform programmatic functions within a page, and can be used for generating more complex content structures.

Generally speaking, a macro is simply a command wrapped inside curly braces {...}. To learn how to write your own macro, or use macros written by other people, read the [Confluence Plugin Guide](#).

 The information on this page *does not apply to* Confluence OnDemand.

RELATED TOPICS:

- [Configuring the User List Macro](#)
- [Enabling HTML macros](#)
 - [Enabling the html-include Macro](#)
- [Troubleshooting the Gallery Macro](#)

Configuring the User List Macro

The [User List macro](#) has an optional **Display Online** parameter. If the User Listener plugin is configured to allow this feature, then the page author can select **Display Online** to show a list of all online users.

 You need to have [System Administrator](#) permissions in order to perform this function.

To enable the Display Online filter in the User List macro:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Select **Plugins** in the left-hand panel. This will list the currently installed plugins.
3. Scroll down and click **User Listener**. The User Listener plugin panel will appear at the top of the screen.
4. Enable the User Log In Listener module by clicking **Enable** on its right.
5. Restart Confluence.

 The information on this page *does not apply to* Confluence OnDemand.

List of online users can be misleading

When the **Display Online** parameter is used, Confluence uses a context listener to generate the list of online users. A context listener is a J2EE term for something that listens for events in the application server. We listen for session open and close events, so a user is 'online' if they have a session on the application server. Some application servers don't correctly despatch close events for sessions – in these cases, the list of online users may be misleading.

Screenshot: Enabling the User Log In Listener

User Listener**Vendor:** [Atlassian Software Systems](#)**Plugin Version:** 2.1

A plugin which reports on Users, per group, within Confluence

□ [Disable plugin](#)**userlister**

Outputs lists of users, whether entirely or in specified groups

[Disable](#)**User Log in Listener**

Informs the UserListener macro when users log in or out of Confluence.

[Enable](#)**Related Topics**[User List Macro](#)[Enabling and Configuring Macros](#)**Enabling HTML macros**

The {html} macro allows you to use HTML code within a Confluence page.

The {html-include} macro allows you to include the contents of an HTML file in a Confluence page.

**Caution: Including unknown HTML inside a web page is dangerous.**

Because HTML can contain active scripting components, it would be possible for a malicious attacker to present a user of your site with script that their web browser would believe came from you. Such code could be used, for example, to steal a user's authentication cookie and give the attacker their Confluence login password.

By default, the HTML macros are disabled. You should only turn on these macros if you trust all your users not to attempt to exploit them.

You need [System Administrator](#) permissions in order to perform this function.**Related pages:**

- [Working with Macros](#)
- [Confluence Administrator's Guide](#)



The information on this page *does not apply* to Confluence OnDemand.

To enable the HTML macros:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **Manage Add-ons** in the left-hand panel. This will display the installed add-ons on this Confluence installation.
3. Choose **Show System Add-ons**.
4. Choose **Confluence HTML Macros**
5. Expand the **x of 11 modules enabled** to see the list of modules in the plugin.
6. Enable one or both of the following modules:
 - The **html (html-xhtml)** module for the [HTML Macro](#).
 - The **html-include (html-include-xhtml)** module for the [HTML Include Macro](#).

You will need to configure [one or more allowed sources](#) for this macro by adding them to the whitelist.**Enabling the html-include Macro**

The `{html-include}` macro allows you to include the content of an HTML file in a Confluence page. This page tells you how to enable the macro, so that it is available on your Confluence site. For help on using the macro, see [HTML Include Macro](#).

**CAUTION: Including unknown HTML inside a web page is dangerous.**

Because HTML can contain active scripting components, it would be possible for a malicious attacker to present a user of your site with script that their web browser would believe came from you. Such code could be used, for example, to steal a user's authentication cookie and give the attacker their Confluence login password.



*The information on this page **does not apply to** Confluence OnDemand.*

Enabling the HTML Macros

By default, the HTML macros are disabled. You should only turn on these macros if you trust all your users not to attempt to exploit them.



You need to have [System Administrator](#) permissions in order to perform this function.

To enable the HTML macros,

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Select '**Plugins**' in the left-hand panel. This will display the installed plugins active for this Confluence installation.
3. Click '**HTML macros**', then click '**Enable Plugin**'.



If the plugin is already enabled please click on the + sign to expand the modules. Ensure that each relevant module is enabled by hovering over the "disabled" dialog and clicking the enable button that appears.

To embed an external page,

Use the following syntax:

```
{html-include:url=http://www.example.com}
```

To include HTML inline,

Use the following syntax:

```
{html}  
<b>I like cheese</b>  
{html}
```

RELATED TOPICS

[HTML Include Macro](#)

Content by label

There is no content with the specified labels



 [Administrators Guide Home](#)  [Confluence Documentation Home](#)

Troubleshooting the Gallery Macro

For guidelines on using the macro, see [Gallery Macro](#).

Troubleshooting

If you encounter the following error message: System does not support thumbnails: no JDK image support then ensure that you have following system property available for your JVM:

```
JAVA_OPTS=-Djava.awt.headless=true
```

Also see [CONF-1737](#)

 Please note that gallery-ext.jar is available at [CONF-6620](#)

 The information on this page *does not apply to* Confluence OnDemand.

Adding, Editing and Removing User Macros

User macros are short pieces of code that perform an often-used function or add some custom formatting to a page. People can call the macro into action by adding the macro keyword to their Confluence pages. You can write a 'user macro' by adding code on a screen in the Confluence Administration Console.

Notes:

- You need [System Administrator](#) permissions in order to perform this function.
- See [Shared User Macros](#) for a list of community-donated macros.
-  Be careful when installing user macros from unknown authors.
- If you remove a user macro that is in use on Confluence pages, you will need to remove the macro from the pages manually. When you remove the user macro, the usage of the macro on the page will become invalid. *Hint:* Use the [Confluence search](#) to find all occurrences of the macro on pages and blog posts.

 The information on this page *does not apply to* Confluence OnDemand.

To add a user macro:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Click **User Macros** in the left-hand panel.
3. Click **Create a User Macro** at the top of the list of macros.
4. Enter the macro details as explained in the guide to [writing user macros](#).
5. Click **Add**.

To edit a user macro:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Select **User Macros** in the left-hand panel. This will list the currently configured user macros.
3. Click **Edit** next to the relevant macro.
4. Update the macro details as explained in the guide to [writing user macros](#).

5. Click **Save**.

To delete a user macro:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Select **User Macros** in the left-hand panel. This will list the currently configured user macros.
3. Click **Delete** next to the relevant macro.

Related Pages

- [Best Practices for Writing User Macros](#)
- [Examples of User Macros](#)
- [Confluence Administrator's Guide](#)

Writing User Macros

User macros are short pieces of code that perform an often-used function or add some custom formatting to a page. People can add the macro to a page by choosing it from the Macro Browser when editing a Confluence page. The macro is run when the page is loaded by the browser. You can write a user macro by adding code on a screen in the Confluence Administration Console.

You need to have [System Administrator](#) permissions in order to create user macros.

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](#).

On this page:

- [Creating a User Macro](#)
- [Examples and Best Practices](#)
- [Related Topics](#)

 *The information on this page **does not apply** to Confluence OnDemand.*

Creating a User Macro

To create a user macro:

1. Go to the **Confluence Administration Console** and click **User Macros** in the left-hand panel.
2. Click **Create a User Macro**.
3. Supply the information in the input fields as explained below, then click **Add**.

The sections below tell you about each of the input fields.

Macro Name

Enter the text that people will see when looking for the macro in the Macro Browser.

Visibility

Set the visibility options to specify who can see this macro when they are searching using the Macro Browser or Autocomplete.

 User macros must have parameters defined in order to appear in the Confluence 4.0 Macro Browser.

The options are as follows:

Visibility Option	Meaning

Visible to all users	All users will see this macro when searching for a macro using the Macro Browser or Autocomplete.
Visible only to system administrators	Choose this option if you want the macro to be 'hidden' from most users when the users are looking for a macro to add to a page. Note that this does not completely hide the macro. Instead, it is useful if you want to avoid cluttering the Macro Browser and Autocomplete with unnecessary macros. Specifically, if you are: • Editing a page and inserting a macro using the Macro Browser: Only system administrators will see this macro in the Macro Browser. For other users, the macro will not show up in the Macro Browser when the user searches for a macro to add to a page. • Editing a page and inserting a macro using Autocomplete: Only system administrators will see this macro in Autocomplete. For other users, the macro will not show up in the Autocomplete list when the user searches for a macro to add to a page. • Viewing the page: The macro output will be visible to all users who have permission to see the page. • Editing a page that already contains the macro: Provided a user has permission to edit the page, the macro will be visible to all users when editing the page, and all users who have permission to edit the page will also be able to edit or remove the macro. Please note that all the macro information will also be 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.

Macro Title

Enter the text that should appear in the Macro Browser and in Autocomplete, to identify this macro when people are looking for it to insert onto a page.

Description

Enter the text that should appear in the Macro Browser describing this macro. Note that the Macro Browser's search will pick up matches in the description as well as in the title.

Categories

Select one or more categories for your macro. To select more than one category, hold down the 'Ctrl' key while selecting. These are the categories that appear in the Macro Browser, helping users to choose a macro from a logical set.

Icon URL

If you would like the Macro Browser to display an icon for your macro, enter the URL here. You can enter an absolute URL or a path relative to the Confluence base URL. For example:

- Absolute URL:

```
http://mysite.com/mypath/status.png
```

- Relative URL:

```
/images/icons/macrobrowser/status.png
```

Documentation URL

Enter the URL pointing to the online help or other documentation for your macro.

Macro Body Processing

Specify how you want Confluence to process the body of your macro before passing it to your macro. Below is an explanation of the macro body and the options available.

What is the macro body?

The macro body is the content that is displayed on the wiki page. If the macro allows a body, users will be able to enter body content when configuring the macro in the Macro Browser.

How can I use the macro body?

If you specify that your macro has a body, you will be able to pass text to the macro when you invoke it from within a page.

If your macro has a body, any body content that the user enters will be available to the macro in the `$body` variable. See the section about the template [below](#). In addition, the options below allow you to tell Confluence to pre-process the body before it is placed in the macro output.

What are the options for macro body?

Body Processing Option	Meaning
No macro body	Select this option if your macro does not need a body.
Escaped	If your macro has a body, and you make use of the body as <code>\$body</code> in your template, Confluence will add escape characters to the HTML markup in the macro body. You could use this if you want to show the HTML markup in the rendered page. For example, if the body is: <pre>Hello World</pre> Then value of <code>\$body</code> will be: <pre>&lt;b&gt;Hello World&lt;/b&gt;</pre> This will render as: <pre>Hello World</pre>

Unrendered	If your macro has a body, and you make use of the body as \$body in your template, HTML in the body will be processed within the template before being output. Ensure that HTML is ultimately output by the template.
Rendered	If your macro has a body, and you make use of the body as \$body in your template, Confluence will recognise HTML in the macro body. For example, if the body is: <pre>Hello World</pre> Then value of \$body will be: <pre>Hello World</pre> This will render as: Hello World

Template

Enter code to specify what the macro will do. For example, to add a macro inside the macro you are writing, you would write:

```
<ac:macro ac:name="someOtherMacro" />
```

Quick guide:

- 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 '\$body'.
- Each parameter variable you use must have a matching metadata definition. Use @param to define metadata for your macro parameters.
- When using the information passed using parameters, refer to your parameters as \$paramXXX where 'XXX' is the parameter name that you specified in the @param metadata definition.
- Use @noparams if your macro does not accept parameters.

See our detailed guide to [writing a user macro template](#).

Examples and Best Practices

See:

- [Examples of User Macros](#)
- [Best Practices for Writing User Macros](#)

Related Topics

Developer documentation:

- [User Macro Module](#)
- [Macro Module](#)
- [Confluence Plugin Guide](#)

Community contributions

- Blog post: [On converting wiki markup based user macros for use with Confluence 4](#)
- Library of user-contributed user macros: [Shared User Macros](#)

 Be careful when installing user macros. Ideally use only macros from authors and sources that are well known to you.

Best Practices for Writing User Macros

This section contains tips and suggestions for best practice in macro coding. To see how to write a user macro and add it to your Confluence site, take a look at our guide to [writing user macros](#).

 *The information on this page **does not apply** to Confluence OnDemand.*

Add a Descriptive Header to your Macro Template

We recommend that you include a short description of your macro via comments at the top of the **Template** field as shown below. You can see an excellent example in the ['Image rollover' user macro](#).

```
## 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

Confluence offers great options for making your macro look good in the macro browser. 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.

In particular, read the documentation on [defining the macro parameters](#) to be displayed in the macro browser.

Supply Default Values for Macro Parameters

You cannot guarantee that a user will supply parameters, so 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. It 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
```

Related Topics

[Writing User Macros](#)

Examples of User Macros

Below are some sample user macros. To see how to write a user macro and add it to your Confluence site, take a look at our guide to [writing user macros](#).

Example 1: A macro that displays 'Hello World'

Take a look at an [example of a 'Hello World' macro](#).

Example 2: The 'Error' macro that creates a red box

Let's write a simple macro that creates a red box (using an existing Confluence style) around some text. See [Error or Box Macro - Example of a User Macro](#).

Example 3: A macro that demonstrates the use of parameters

See [Colour and Size Macro - Example of a User Macro](#).

Example 4: A macro that prevents text from being printed

See [NoPrint Example of a User Macro](#).

On this page:

- [Example 1: A macro that displays 'Hello World'](#)
- [Example 2: The 'Error' macro that creates a red box](#)
- [Example 3: A macro that demonstrates the use of parameters](#)
- [Example 4: A macro that prevents text from being printed](#)
- [Example 5: A macro that creates a preformatted panel](#)
- [Example 6: A macro that creates a preformatted table](#)
- [Community-contributed user macros](#)

Related pages:

- [Writing User Macros](#)
- [Guide to User Macro Templates](#)
- [Guide to User Macro TemplatesWorking with Macros](#)



The information on this page does not apply to Confluence OnDemand.

Example 5: A macro that creates a preformatted panel

This user macro creates a panel preformatted with specific colours. See [Panel Preformatted with Specific Colours - Example of a User Macro](#).

Example 6: A macro that creates a preformatted table

This user macro creates a table with predefined headings. See [Preformatted Table - Example of a User Macro](#).

Community-contributed user macros

You may want to take a look at the library of [shared user macros](#).

Be careful when installing user macros from unknown authors.

Hello World Example of User Macro

This page tells you 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. For full details about creating a user macro, see the guide to [writing user macros](#).

Defining the 'Hello World' user macro**To create the 'Hello World' user macro:**

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **User Macros** in the left-hand panel.
3. Choose **Create a User Macro** at the bottom of the list of macros.
4. Enter the macro attributes as follows:
 - 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 empty.
 - Documentation URL: You can leave this field empty.
 - Macro Body Processing: Rendered
 - Template:

```
## @noparams  
Hello World!  
$body
```

5. Choose **Save**.

Related pages:

- [Writing User Macros](#)
- [Guide to User Macro Templates](#)
- [Examples of User Macros](#)

 The information on this page *does not apply* to Confluence OnDemand.

Screenshot: Defining the 'Hello World' user macro

Macro Browser Information

Macro Name*

Macro names are converted to lower case

- Visibility ☒ Visible to all users in the Macro Browser
☐ Visible only to system administrators in the Macro Browser

Macro Title*

Description

Categories

Administration
Communication
Confluence Content
Development
External Content

Icon URL

Absolute URL (e.g. <http://mysite.com/myimage.png>), or relative to the Confluence base URL (e.g. [/images/logo/confluence_48_white.png](/images/logo/confluence_48_white.png))

Documentation URL

Definition of User Macro

- Macro Body ☐ No macro body
Processing ☐ Escaped
☐ Unrendered

You should use this option for bodies that are processed within the template before being output. Ensure that HTML is ultimately output by the template.

- ☒ Rendered

The body will be rendered so most HTML entered will be passed to the template unmodified but Confluence specific mark up such as macro definitions will be rendered.

Template*


See the [documentation](#).

Save


Cancel

Using the 'Hello World' macro on a page

Now you can add the macro to your Confluence page using the Macro Browser.



The result is:



You can also use autocomplete to add the macro onto your page: start typing '{hello}' in the editor, and select the 'Hello World' macro from the list of suggestions that appears.

Error Box Macro - Example of a User Macro

Let's write a simple macro that creates a red box (using an existing Confluence style) around some text. This may be useful for writing about error conditions, for example. For full details about creating a user macro, see the guide to [writing user macros](#).

Defining the 'Error' user macro

To create the 'Error' user macro:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **User Macros** in the left-hand panel.
3. Choose **Create a User Macro** at the bottom of the list of macros.
4. Enter the macro attributes as follows:
 - Macro Name: error
 - Visibility: Visible to all users in the Macro Browser
 - Macro Title: Error
 - Description: Displays a red box around some text
 - Categories: Confluence Content
 - Icon URL: You can leave this field empty.
 - Documentation URL: You can leave this field empty.
 - Macro Body Processing: Rendered
 - Template:

```
## @noparams
<div class="error">$body</div>
```

5. Choose **Save**.

Related pages:

- [Writing User Macros](#)
- [Guide to User Macro Templates](#)
- [Examples of User Macros](#)

 *The information on this page does not apply to Confluence OnDemand.*

Using the 'Error' macro on a page

To add the macro to a page, edit the page and choose **Insert > Other Macros** and find the 'Error' macro. (Or use autocomplete: start typing '{err}' in the editor, and select the 'Error' macro from the list of suggestions that appears.)

Your page will display an error box, like this:

(Write your error message here.)

Colour and Size Macro - Example of a User Macro

This example demonstrates how you can pass parameters into your macro. Let's say you want to write your own font colour macro, with a parameter allowing the user to specify the colour. Then perhaps you want to add another parameter, that allows the user to specify the font size.

For full details about creating a user macro, see the guide to [writing user macros](#).

Defining the 'Colour' user macro

This example uses a single parameter.

To create the 'Colour' user macro:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **User Macros** in the left-hand panel.
3. Choose **Create a User Macro** at the bottom of the list of macros.
4. Enter the macro attributes as follows:
 - Macro Name: colour
 - Visibility: Visible to all users in the Macro Browser
 - Macro Title: Colour
 - Description: Colours a block of text
 - Categories: Confluence Content
 - Icon URL: You can leave this field empty.
 - Documentation URL: You can leave this field empty.
 - Macro Body Processing: Rendered
 - Template:

```
## @param 0:title=colour|type=string
<span style="color: $param0">$body</span>
```

5. Choose **Save**.

Related pages:

- [Writing User Macros](#)
- [Guide to User Macro Templates](#)
- [Examples of User Macros](#)

 *The information on this page **does not apply** to Confluence OnDemand.*

Using the 'Colour' macro on a page

To add the macro to a page, edit the page and choose **Insert > Other Macros** and find the 'Colour' macro. (Or use autocomplete: start typing '{colo' in the editor, and select the 'Colour' macro from the list of suggestions that appears.)

Defining the 'Stylish' user macro

If your macro requires more than one parameter, you can use variables \$param0 to \$param9 to represent them. Let's say that you want to add a parameter that allows the user to specify the size of the text.

Enter the macro attributes as follows:

- 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 empty.
- Documentation URL: You can leave this field empty.
- Macro Body Processing: Rendered
- Template:

```
## @param 0:title=colour|type=string
## @param 1:title=size|type=string
<span style="color: $param0; font-size: $param1">$body</span>
```

Naming your parameters

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.

```
## @param Colour:title=colour|type=string
## @param Size:title=size|type=string
<span style="color: $paramColour; font-size: $paramSize">$body</span>
```

NoPrint Example of a User Macro

This page gives an example of a user macro, the 'NoPrint' macro, that you can use to prevent text from being printed. For full details about creating a user macro, see the guide to [writing user macros](#).

Defining the 'NoPrint' user macro

To create the 'NoPrint' user macro:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **User Macros** in the left-hand panel.
3. Choose **Create a User Macro** at the bottom of the list of macros.
4. Enter the macro attributes as follows:
 - Macro Name: noprint
 - Visibility: Visible to all users in the Macro Browser
 - Macro Title: NoPrint
 - Description: Hides text from printed output.
 - Categories: Confluence Content
 - Icon URL: You can leave this field empty.
 - Documentation URL: You can leave this field empty.
 - Macro Body Processing: Rendered
 - Template:

```
## @noparams
<div class="noprint">$body</div>
```

5. Choose **Save**.


Related pages:

- [Writing User Macros](#)
- [Guide to User Macro Templates](#)
- [Examples of User Macros](#)

 *The information on this page does not apply to Confluence OnDemand.*

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](#).

Panel Preformatted with Specific Colours - Example of a User Macro

This user macro creates a panel pre-formatted to specific colours. It will create a panel that looks like this:



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

Related pages:

- [Writing User Macros](#)
- [Guide to User Macro Templates](#)
- [Examples of User Macros](#)

 *The information on this page **does not** apply to Confluence OnDemand.*

Defining the 'Formatted Panel' user macro

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **User Macros** in the left-hand panel.
3. Choose **Create a User Macro** at the bottom of the list of macros.
4. Enter the macro attributes as follows:
 - 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 empty.
 - Documentation URL: You can leave this field empty.
 - Macro Body Processing: Escaped
 - Template:

```
## @param Title:title=Title|type=string|desc=Title

<ac: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:macro>
```

5. Choose **Save**.

Explanation of the code in the macro template

Item	Description

<pre>## @param Title:title=Title type=string desc=Title</pre>	@param defines the metadata for your macro parameters. When users select this macro, the macro will contain a parameter called "Title" where they can enter data. A macro dialog window appears when the user selects this macro using Insert > Other Macros or when a user clicks the macro placeholder and chooses Edit. The macro will, later on, use the data stored in this parameter to enter data in the title section of the Panel macro. @param Title This parameter is called "Title". title=Title defines the parameter title that will appear in the macro dialog window as "Title". type=string defines the field type for the parameter as a text field. desc=Title defines the description of the parameter.
<pre><ac:macro ac:name="panel"></pre>	This command activates the Confluence Panel macro . Hint: To discover the code name of a Confluence macro, see the documentation for each macro. If the macro you want is not documented there, follow these steps: 1. Create and save a page containing a Confluence macro you want to investigate. 2. Choose Tools > View Storage Format. This option is available to Confluence administrators only, and shows the XML source code for the page. (See Confluence Storage Format.) 3. A Confluence macro starts with the following string: <ac: macro ac:name=
<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>	Sets the parameters for the macro: the background colour, border style, border colour, border width and title colour. Hint: To discover the names of the parameters for a Confluence macro, see the documentation for each macro. If the macro you want is not documented there, follow these steps: 1. Create and save a page containing a Confluence macro you want to investigate. 2. Choose Tools > View Storage Format. This option is available to Confluence administrators only, and shows the XML source code for the page. (See Confluence Storage Format.) 3. The macro parameters start with the following string: <ac: parameter ac:name
<pre><ac:parameter ac:name="title">\${!paramTitle}</ac:parameter></pre>	Enters the value stored in the 'Title' parameter into the title section of the macro. The ! tells the macro to leave the title blank, when there is no data in the "Title" parameter.

<code><ac:rich-text-body>\$body</ac:rich-text-body></code>	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.
<code></ac:macro></code>	This command marks the end of the macro.

Preformatted Table - Example of a User Macro

This user macro creates a 2 x 2 table, with the headings defined as 'Parameter' and 'Description'. It will create a table that looks like this:

Parameter	Description

Note: As the macro is written, the user cannot amend the heading titles when using the macro on a Confluence page.

Related pages:

- [Writing User Macros](#)
- [Guide to User Macro Templates](#)
- [Examples of User Macros](#)

 The information on this page *does not apply to* Confluence OnDemand.

Defining the 'Formatted Table' user macro

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **User Macros** in the left-hand panel.
3. Choose **Create a User Macro** at the bottom of the list of macros.
4. Enter the macro attributes as follows:
 - Macro Name: formtable
 - Visibility: Visible to all users in the Macro Browser
 - Macro Title: Formatted Table
 - Description: Creates a simple 2 x 2 table with the column headings filled in
 - Categories: Formatting
 - Icon URL: You can leave this field empty.
 - Documentation URL: You can leave this field empty.
 - Macro Body Processing: Escaped
 - Template:

```
## @param Head1:type=string|desc=Heading
## @param Head2:type=string|desc=Heading
## @param Cell1:type=string|desc=cell
## @param Cell2:type=string|desc=cell
#set ($paramHead1 ="Parameter")
#set ($paramHead2 ="Description")

<div id="preformattedtable">
<table>
<tr>
<th>$paramHead1</th>
<th>$paramHead2</th>
</tr>
<tr>
<td>$!paramCell1</td>
<td>$!paramCell2</td>
</tr>
</table>
</div>
```

5. Choose **Save**.

Using the macro on a Confluence page

To add the macro to a page:

1. In the Confluence editor, choose **Insert > Other Macros**.
2. Find and select the 'Formatted Table' macro.
3. Enter the cell contents into the form.
4. Choose **Insert**.

Amending the contents of the table

To change the content in the cells of the table:

1. Edit the page.
2. Click the macro placeholder for the 'Formatted Table' macro, to see the properties panel.
3. Choose **Edit**.
4. Enter the cell contents into the form.
5. Choose **Save**

Note: Content entered into the body of the 'Formatted Table' macro will not appear on the page.

Guide to User Macro Templates

You can create a user macro in Confluence by typing it into a screen in the Confluence Administration Console. The 'template' is one of the fields that you define when writing a user macro. (See the rest of the guide to [writing user macros](#).) This page gives you guidelines about the code you can enter in a user macro template.

Quick guide to user macro templates:

- 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 '\$body'.
- Each parameter variable you use must have a matching metadata definition. Use @param to define metadata for your macro parameters.
- When using the information passed using parameters, refer to your parameters as \$paramXXX where 'XXX' is the parameter name that you specified in the @param metadata definition.
- Use @noparams if your macro does not accept parameters.

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](#)
- [Examples of User Macros](#)



The information on this page *does not apply* to Confluence OnDemand.

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 `helloWorld`. 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's 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 of your macro's parameters. The field type is determined by the parameter type you specify for each parameter.

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.

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: <pre>## @param colour:title=Colour type=enum enumValues=Grey, Red, Yellow, Green</pre> <i>Note about i18n:</i> Confluence does not support internationalisation of the enum values. The value the user sees is the one passed to the macro as the parameter value, with the capitalisation given. In this case 'Grey', 'Red', etc.
string	A text field. This is the default type. Example with a required field: <pre>## @param status:title=Status type=string required=true desc=Status to display</pre>
confluence-content	Offers a control allowing the user to search for a page or blog post. Example: <pre>## @param page:title=Page type=confluence-content required=true desc=Select a page do use</pre>
username	Search for user. <pre>## @param user:title=Username type=username desc=Select username to display</pre>
spacekey	Offers a list of spaces for selection. Passes the space key to the macro. Example: <pre>## @param space:title=Space type=spacekey</pre>

date	Confluence accepts this type, but currently treats it in the same way as 'string'. Example: <pre>## @param fromDate:title=From Date type=date desc=Date to start from. Format: dd/mm/YYYY</pre> <i>Note about dates:</i> A user can enter a date in any format, you should validate the date format in your user macro.
int	Confluence accepts this type, but currently treats it in the same way as 'string'. Example with a default value: <pre>## @param numPosts:title=Number of Posts type=int default=15 de sc=Number of posts to display</pre>
percentage	Confluence accepts this type, but currently treats it in the same way as 'string'. Example: <pre>## @param pcent:title=Percentage type= percentage desc=Number of 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. That will let Confluence know that it need not display a parameter input field in the macro browser.

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

Configuring the Office Connector

The Office Connector is a Confluence add-on that allows Confluence users to interact with Microsoft Office and Open Office in various ways. You can display content from Office documents on a wiki page and import content from an Office document into Confluence. Please refer to the [User Guide](#) for details of these interactions.

The Office Connector add-on is shipped with Confluence. A [System Administrator](#) can enable or disable parts of the Office Connector and can configure options as described below.

Enabling and Disabling the Office Connector and its Modules

The Office Connector is bundled with Confluence, so you should not need to install it. But you may wish to enable or disable some of its modules.

To enable or disable the Office Connector and its modules:

1. Select **Manage Add-ons** in the left-hand panel of the Confluence Administration Console.
2. Click **Show system add-ons** under 'System Add-ons'.
3. Enter 'Office Connector' in the **Filter Visible add-ons** field to quickly find the Office Connector add-on.
4. Open the details view of the add-on by clicking on the Office Connector add-on in the system add-ons list.
5. From the details view, you can:
 - Click **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.
 - View the **modules** that make up the add-on by expanding the modules list. You can enable or disable certain Office Connector modules.

On this page:

- [Enabling and Disabling the Office Connector and its Modules](#)
- [Configuring the Office Connector Options](#)


Related pages:

- [Office Connector Prerequisites](#)
- [Office Connector Limitations and Known Issues](#)
- [Working with the Office Connector](#)
- [Managing Add-ons and Macros](#)

 The information on this page *does not* apply to Confluence OnDemand.

To disable or enable a module:

1. Open the details view for the Office Connector add-on in the 'Manage Add-ons' page.
2. Expand the active modules link. The text of this link indicates the number of enabled modules out of the total modules in the add-on.

Screenshot: Enabling the Office Connector add-on and its modules

3. Hover over the module in the list to make the **Enable** or **Disable** button visible, and click the button to apply the action.

Only certain 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

Confluence administrators can configure settings that control the behaviour of the Office Connector on your Confluence site.

To set the configuration options for the Office Connector:

1. Select **Office Connector** under 'Configuration' in the left-hand panel of the 'Confluence Administration Console'. The 'Configure Office Connector plugin' screen appears.

Screenshot: Configuring the Office Connector options

Configure Office Connector plugin

Settings updated successfully

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.
☐ Cache in-memory:

Maximum file space for cache(MB):

Number of Conversion Queues: [Manage Queues](#)

Security

Authentication: ☐ Allow authentication tokens in the URL path

- Set the configuration options as described in the table [below](#).

The configuration options are 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 the button to appear in the page action icon or from the view page tab.

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 document. The warning will tell the user when they are about to overwrite existing content.
Advanced Formatting Options: Use the footnote macro for Word footnotes	Disabled	If this option is enabled, a Confluence page created from an imported Word document will use the {footnote} macro from Adaptavist to render any footnotes contained in the document. Note that you will need to install the Footnotes add-on onto your Confluence site. For more information about this add-on and macro, please refer to the Footnotes add-on .
Authentication: Allow authentication tokens in the URL path	Disabled	If this option is enabled, the Office Connector will use authentication tokens in the URL.
Temporary storage for viewfile macro	The Confluence Home directory.	The {viewfile} macro will cache data temporarily. This option allows you to set the location of the cache. Available settings are: • 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 location by editing the Office Connector's <code>directories.properties</code> file: 1. Go to the <code>bundled-plugins</code> directory in your Confluence Home directory. 2. Copy the Office Connector JAR file to a temporary location: <code>OfficeConnector-x.xx.jar</code>, where 'x.xx' is the version number.

3. Unzip the JAR file and find the `directories.properties` file in the `resources` directory. The content of the file looks like this:

```
#Complete
the
following
line to
set a
custom
cache
directory
.
#If
resetting
to blank,
don't
delete
anything
before or
including
the '='
com.benry
an.conflu
ence.word
.edit.cac
heDir=
```

		4. Edit the last line, adding the path to your required temporary location directly after the '=' character. For example: - On Windows: <pre>com.be nryan. conflu ence.w ord.ed it.cac heDir= c:\my\ path\</pre> - On Linux: <pre>com.be nryan. conflu ence.w ord.ed it.cac heDir= /home/ myuser name/m y/path</pre> 5. Save the file, recreate the JAR and put it in the bundled-plugins directory in your Confluence Home directory, overwriting the original JAR. Cache in-memory – The temporary file will be held in memory. We recommend this option if you are running in a clustered environment.
Maximum file space for cache (MB)	500	This is the maximum size of the cache used by the {viewfile} macro. (See above.)

Number of Conversion Queues	6	This is the maximum number of threads used to convert PowerPoint, Excel files or PDF slide shows. You can use this setting to manage Confluence performance, by limiting the number of threads so that the Office Connector does not consume too many resources. Click Manage Queues to view attachments that are still pending conversion.

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 [Customising Confluence](#).

We have 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:

- [Managing Add-ons and Macros](#)
- [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.

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](#).

Related pages:

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

Customising the Confluence Dashboard

If you are a Confluence Administrator, you can customise the site dashboard, affecting the way all users will see the dashboard. Some of the actions below require Confluence Administrator permissions, whereas others require System Administrator permissions.

Confluence users can customise their own view of the dashboard too. See the [user's guide](#).

Sending users to a space home page instead of the dashboard

See [Configuring the Site Home Page](#).

Editing the top left-hand section of the dashboard

See [Editing the Site Welcome Message](#).

Disabling the 'Popular' tab on the dashboard

In some environments, you may prefer not to display the 'Popular' tab on the dashboard. For example, if your wiki allows only a small group of people to log in and contribute content or comments, then the tab may not be relevant to you.

To prevent the tab from appearing, you can disable the relevant plugin module. You need System Administrator [permissions](#) to do this. Go to the **Dashboard Macros** plugin (See [Configuring a Plugin](#)), choose **Manage plugin modules** and disable the **Popular Tab** module.

On this page:

- [Sending users to a space home page instead of the dashboard](#)
- [Editing the top left-hand section of the dashboard](#)
- [Disabling the 'Popular' tab on the dashboard](#)
- [Advanced customisations](#)
 - [Editing the bottom left-hand section of the dashboard](#)
 - [Editing the top right-hand action bar](#)
 - [Modifying the global template or layout](#)

Related pages:

- [Customising your Personal Dashboard](#)
- [Changing the Look and Feel of Confluence](#)

 *The information on this page does not apply to Confluence OnDemand.*

Advanced customisations

These configurations require knowledge of plugin development and/or the Velocity template language. See our [guide to the Atlassian Plugin SDK](#) and our [introduction to Velocity](#).

Editing the bottom left-hand section of the dashboard

This section can be updated using Confluence web panels. You can add items to the dashboard by including a web panel with the key `atl.dashboard.left`:

```
<web-panel key="{key}" location="atl.dashboard.left">
    <resource name="view" type="velocity"
location="{location}"/>
</web-panel>
```

You can remove the existing entities panel by disabling the `global-entities-panel` plugin from the dashboard macros plugin.

Editing the top right-hand action bar

You can add more links to the top right navigation bar by adding web items to `system.dashboard.button`:

```
<web-item key="{key}" name="{name}" section="system.dashboard.button">
  <label key="{label}" />
  <link />
  <styleClass />
</web-item>
```

Modifying the global template or layout

You can also modify files to add content to the global dashboard.

To make modifications to the dashboard, modify the global template `/confluence/decorators/global.vm` or the layout at  > **General Configuration** > **Layouts** > **Global Layout**.

For example, search the global layout for these macros:

```
$helper.renderConfluenceMacro("{recently-updated-dashboard:dashboard|showProfilePic=true}")
```

To modify the bundled plugin macros used in the Confluence dashboard:

1. Modify the `atlassian-bundled-plugins.zip` file located at `<Confluence install>/confluence/WEB-INF/classes/com/atlassian/confluence/setup`.
2. Update the `confluence-dashboard-macros-x.x.jar` file, rezip it and then put it back to `<Confluence install>/confluence/WEB-INF/classes/com/atlassian/confluence/setup`. Refer to [How to Edit Files in Confluence JAR Files](#)
3. Delete the JAR from `<confluence-home>/bundled-plugins`.
4. Restart Confluence.

To customise the space list, you can work with `spaceList.vm`.


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.



On this page:

- [Changing the site logo](#)
- [Changing the site icon \(favicon\)](#)
- [Changing the default space logo](#)
- [Changing a specific space logo](#)

Related pages:

- [Changing the Look and Feel of Confluence](#)
- [Customising Colour Schemes](#)
- [Confluence Administrator's Guide](#)

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:

- [Working with 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
- **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		#205081	
Top Bar Text		#FFFFFF	
Header Button Background		#3B7FC4	
Header Button Text		#FFFFFF	
Top Bar Menu Selected Background		#3B73AF	
Top Bar Menu Selected Text		#FFFFFF	
Top Bar Menu Item Text		#333333	
Menu Item Selected Background		#3B73AF	
Menu Item Selected Text		#FFFFFF	
Page Menu Selected Background		#3B73AF	
Page Menu Item Text		#333333	
Heading Text		#333333	
Space Name Text		#999999	
Links		#3B73AF	
Borders and Dividers		#CCCCCC	
Tab Navigation Background		#3B73AF	
Tab Navigation Text		#FFFFFF	
Tab Navigation Background Highlight		#205081	
Tab Navigation Text Highlight		#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.

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.

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.

Related pages:

- [Applying a Theme to a Space](#)
- [Applying a Theme to a Site](#)
- [Configuring the Documentation Theme](#)
- [Confluence Administrator's Guide](#)
- [Creating a Theme](#)

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:

- [Applying a Theme to a Space](#)
- [Configuring 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...](#)

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](#)
- [Applying a Theme to a Space](#)
- [Configuring the Documentation Theme](#)
- [Confluence Administrator's Guide](#)

 The information on this page *does not apply to* Confluence OnDemand.

Customising Site and Space Layouts

You can modify Confluence's look and feel by editing the 'decorator' (layout) files. Modifying these files allows you to change the look and feel of:

- The Confluence site as a whole, which includes all spaces within the Confluence site.
- An individual space within the Confluence site.

This page tells you how to customise the layout files for your Confluence site as a whole. These customisations:

- Modify the default 'decorator' files of each space in your site.
- Are reflected in every space unless the space's own equivalent layout files have been customised.

You need [System Administrator](#) permissions to perform these customisations.

You can also customise the layout files for a given space. For more information, refer to [Customising Space Layouts](#). Space layout customisations override the equivalent site customisations.

Note: If you modify the look and feel of Confluence by following these instructions, you will need to update your customisations when upgrading Confluence. The more dramatic the customisations are, the harder it will be to reapply your changes when upgrading. Please take this into account before proceeding with your customisation. For more information on updating your customisations, please refer to [Upgrading Customised Site and Space 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](#)
- [Confluence Administrator's Guide](#)

Confluence is built on top of the open source [SiteMesh](#) library, a web-page layout system. Read more on the [SiteMesh website](#). 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 a simple programming language called 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 and the footer.
- **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. If you are using Confluence to generate a static website, for example, you will need to modify these layouts.

Editing a site decorator file

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Select **Layouts** under **Look and Feel** in the left-hand navigation panel.
 - Click **View Default** to view the .vmd file.
 - Click **Create Custom** to edit the default .vmd file. This will open the .vmd file in edit mode.

3. Make changes and click **Update**.

If something goes wrong: Click **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.



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.



*The information on this page **does not apply to** Confluence OnDemand.*

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** on 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** at the top of the screen.
Note: The 'Space Admin' option appears only if you are a space administrator for the space or you are a super user (a member 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.

RELATED TOPICS

[Customising Site and Space Layouts](#)

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.

 The information on this page *does not apply* to Confluence OnDemand.

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.

RELATED TOPICS

Content by label

There is no content with the specified labels

**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 on [the SiteMesh website](#).

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

Common Modifications

- Rename 'Dashboard' by searching for Dashboard. To change "Dashboard" to "My Portal", change dashboard.name=Dashboard to dashboard.name=My Portal

Common Modifications

Task	Search For	Notes
Rename 'Dashboard'	Dashboard	The dashboard.name parameter has the name. To change 'Dashboard' to 'My Portal', change dashboard.name=Dashboard to dashboard.name=My Portal and update any other occurrences of the word 'Dashboard' in the instance

Modify login page text	login.	The <code>login.instructions</code> parameter has the "Enter your account details below to login to Confluence" text

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-resources` plugin, which is a bundled plugin (add-on) that is installed automatically when you install Confluence.

Only administrators with access to the server where Confluence is running 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:

- In the Confluence web application folder, find the file `/confluence/WEB-INF/atlassian-bundled-plugins/confluence-email-resources-x.x.jar`
Note: This plugin is independently versioned, the version number will not necessarily match Confluence's version number.
- Copy this file to a working location and extract the jar file. Find out more about how to [edit files within .jar archives](#).
- Within the jar file, templates are stored in the `/templates/` folder. Edit the Soy templates to make your changes.
- 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).
- In your Confluence test instance go to  > **Add-ons** and upload your new jar file.
- 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](#)
- [Configuring the What's New Dialog](#)
- [Customising Default Space Content](#)
- [Customising the Getting Started Guide on the Dashboard](#)
- [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 [Working with Templates](#) and [Working with 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 [Adding 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 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.

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:

- [Creating Content](#)
- [Working with Templates](#)
- [Confluence Administrator's Guide](#)

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

Import

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:

- [Editing User Settings](#)
- [Recognised System Properties](#)
- [Configuring Indexing Language](#)
- [Installing a Language Pack](#)

 The information on this page *does not apply to* Confluence OnDemand.

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. The '**Language Configuration**' screen will appear. Select the language that you want to use as the default language for your Confluence site.

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 [site 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.

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**.

On this page:

- [Customising the Administrator Contact Message](#)
- [Disabling the Administrator Contact Form](#)
- [Configuring Spam Prevention](#)


Related pages:

- [Contacting Confluence Administrators](#)
- [Configuring Captcha for Spam Prevention](#)

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

You can configure Confluence to send people to any space home page when they log in or click the site logo, rather than to the dashboard.

The spaces available to set as the site home page will depend on the access permissions of the space and the site.

- The site home page must be accessible to the 'confluence-users' or 'users' group.
- If the site allows anonymous access, the site home page must also be accessible to anonymous users, that is, people who have not logged in to Confluence.

To configure the site-wide home page:

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 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.
5. Choose **Save**.

Related pages:

- [Editing the Site Welcome Message](#)
- [Changing the Site Title](#)
- [Customising Default Space Content](#)
- [Changing the Site Logo](#)
- [Confluence Administrator's Guide](#)


Accessing the dashboard with a site homepage set

If you choose to set a space homepage 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. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **Application Navigator**.
3. Enter the name for your link, for example, 'Dashboard'.
4. Enter the URL for your site dashboard, for example, `https://yoursite.com/wiki/dashboard.action`.
5. Choose **Add**.

A link to the dashboard will now appear in the Application Navigator.



Notes

- The [user's personal settings](#) will override the global setting.
- If you allow anonymous access to the dashboard, but not anonymous access to the site home page, then when logging on to the site, users will be redirected to the original dashboard instead of the site home page. To avoid this, either make the site home page accessible anonymously, or make the dashboard not accessible anonymously.

Configuring the What's New Dialog

The 'What's New' dialog pops up automatically when a user logs in for the first time after a major Confluence upgrade (such as an upgrade to Confluence 4.3). The dialog displays a summary of the new features for the release, sourced from the Atlassian website (by default).

Confluence administrators can configure the behaviour of the 'What's New' dialog, as follows:


- Change the URL that the 'What's New' dialog retrieves information from.
- Disable the dialog.

On this page:

- [Changing the 'What's New' Dialog URL](#)
- [Disabling the 'What's New' Dialog](#)

Related pages:

- [Disabling and enabling add-ons](#)



Screenshot above: An example of the 'What's New' dialog

Changing the 'What's New' Dialog URL

The 'What's New' dialog URL is stored in your Confluence `help-paths.properties` file. This URL is a concatenation of the `help.prefix` property with the `help.whats.new.iframe.link`.

Note: The `help.prefix` property also defines the base URL for Confluence help links, i.e. help links in the Confluence application.

To change the 'What's New' Dialog URL:

Follow the instructions in the 'Changing the Links for Individual Help Pages' section on [Local Confluence Documentation](#). You will need to update the `help.prefix` and `help.whats.new.iframe.link` properties, as desired.

For example, you may have installed your Confluence documentation behind a firewall at <http://www.example.com/> and created a page <http://www.example.com/whatsnew> that you use for change management. In this case, you would do the following:

- Set `help.prefix` to <http://www.example.com/>
- Set `help.whats.new.iframe.link` to `whatsnew`

There is an additional property `help.whats.new.full.link`. This is only used if the content pointed to by the updated URL isn't loaded in 10 seconds, in which case a 'timeout' screen is displayed with a link to the full 'What's New' content. For locally-hosted pages you can just set this property to the same value as `help.whats.new.iframe.link`.

Disabling the 'What's New' Dialog

The 'What's New' dialogue is enabled via a plugin. To disable the 'What's New' dialogue, you need to disable the **Confluence What's New** plugin in Confluence.

To disable the 'Confluence What's New' plugin:

Follow the instructions on [Disabling and enabling add-ons](#). Please note, the **Confluence What's New** plugin is a 'System Plugin'. Click **Show System Plugins** on the Manage Add-ons administration page to display the system plugins.

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.
2. Choose **Global Templates and Blueprints** in the left-hand panel.
3. 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.
4. 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.
5. Choose **Save**.

On this page:

- [Edit the default home page for a blank space](#)
- [Reset the original default content](#)

Related pages:

- [Working with Spaces](#)
- [Working with 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.

Customising the Getting Started Guide on the Dashboard

By default, the Confluence dashboard displays a quick-start guide for administrators under the site welcome message on the left. This section of the dashboard is visible to Confluence administrators and system administrators only. It is not configurable via the web interface, but you can update or remove it by editing the site layout as described below.

You need [System Administrator](#) permissions to perform this customisation.

Editing or removing the getting-started section

To customise the getting-started guide on the dashboard:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **Layouts** in the left-hand panel.
3. Choose **Create custom** (or **Edit**) next to **Global Layout**.
Note: If the global layout has already been customised, the 'Edit' option will be available. Otherwise, you will need to create the custom layout now, by choosing 'Create custom'.
4. Find the following code:

```
#if($permissionHelper.isConfluenceAdministrator($remoteUser))
    <div class="dashboard-item wiki-content">
        <h2>$i18n.getText("getstarted.heading")</h2>
        <ol id="dashboard-get-started">
            <li class="create-space">
                <h3><a
href="$req.contextPath/spaces/createspace-start.action">$i18n.getText("getstar
ted.add.space")</a></h3>

                <p>$i18n.getText("getstarted.add.space.desc")</p>
            </li>
            <li class="add-users">
                <h3><a
href="$req.contextPath/admin/users/browseusers.action">$i18n.getText("getstart
ed.add.users")</a></h3>

                <p>$i18n.getText("getstarted.add.users.desc")</p>
            </li>
            <li class="user-profile">
                <h3><a
href="$req.contextPath/users/editmyprofilepicture.action">$i18n.getText("getst
arted.choose.profile.picture")</a></h3>

                <p>$i18n.getText("getstarted.choose.profile.picture.desc")</p>
            </li>
        </ol>
    </div>
#end
```

5. Update the code as required:
 - To remove the 'get started' section, delete the entire block of text shown above.
 - Alternatively, edit the code to suit your requirements. See [Customising Site and Space Layouts](#) for guidelines.
6. Choose **Save**.

The default getting-started section

By default, the getting-started guide looks more or less like the screenshot below, starting with the heading 'Get started'.

To restore the default getting-started guide:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **Layouts** in the left-hand panel.
3. Choose **Reset Default** next to **Global Layout**.
Note: This will reset any other customisations applied to this layout too.

On this page:

- [Editing or removing the getting-started section](#)
- [The default getting-started section](#)
- [Notes](#)

Related pages:

- [Customising Site and Space Layouts](#)
- [Editing the Site Welcome Message](#)
- [Configuring the Site Home Page](#)
- [Changing the Site Title](#)
- [Changing the Site Logo](#)
- [Confluence Administrator's Guide](#)

 *The information on this page does not apply to Confluence OnDemand.*

Screenshot: The getting-started guide on the dashboard

Dashboard

Welcome to Confluence

Confluence is where your team collaborates and shares knowledge — create, share and discuss your files, ideas, minutes, specs, mockups, diagrams, and projects.

Get started

**Create a new space**
and start creating content.

**Invite your colleagues**
to join you in Confluence.

**Upload your picture**
and edit your profile.

Notes

If you modify the look and feel of Confluence by following these instructions, you will need to update your customisations when upgrading Confluence. The more dramatic the customisations are, the harder it will be to reapply your changes when upgrading. Please take this into account before proceeding with your customisation. For more information on updating your customisations, please refer to [Upgrading Customised Site and Space Layouts](#).

Editing the Site Welcome Message

The site welcome message appears at the top left of the Confluence dashboard. You can change the default message by editing the appropriate system template. For example, you may want the welcome message to display an introduction to your site or a message of the day.

To edit the site welcome message:

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 **Edit** next to **Default Welcome Message**.
4. Type your message into the template editor.
5. Choose **Save**.

The default site welcome message

By default, the site welcome message looks more or less like the screenshot below, starting with the heading 'Welcome to Confluence' and ending with '...diagrams, and projects'.

To restore the default site welcome message:

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 **Default Welcome Message**.


On this page:

- [The default site welcome message](#)
- [Using the template editor](#)
- [Including content from another page](#)

Related pages:

- [Configuring the Site Home Page](#)
- [Changing the Site Title](#)
- [Changing the Site Logo](#)
- [Customising Default Space Content](#)
- [Confluence Administrator's Guide](#)

Screenshot: Site welcome message at top left of the dashboard



Using the template editor

Enter text into the body of the template, and use the editor toolbar to apply styles, layout and formatting. You can add links and macros. In general, you can use the Confluence editor in the same way as on a page.

Notes:

- You cannot use template variables in the welcome message template.
- You cannot attach an image or other file to a template. Instead, attach the file to another page, and insert it into the body of the template.

For example:

- You can attach an image to a page and then choose **Insert > Image** to embed the image into the template.
- You can attach a PDF file to a page and then choose **Insert > Other Macros > PDF** to embed the PDF file into the template.

Including content from another page

It may be useful to write your welcome message on a normal Confluence page and include the page into the welcome message template. Using a normal page means that you can allow other people, who are not Confluence administrators, to change the welcome message.

To include content from another page:

1. Create a Confluence page as usual and add your welcome message as the page content. Remember to limit the size of the content, because it must fit nicely onto the dashboard. For this example, let's assume the title of your page is 'Dashboard Message'. You can put it in any space you like.
2. Add [page restrictions](#) or [space permissions](#) to the 'Dashboard Message' page or space, to suit your requirements. You may want to restrict the editing of the page to a group of people, or you may want to allow any employee to edit the page. This will determine who can update the welcome message on the dashboard.
3. Edit the welcome message template, and add the [Include Page](#) macro to display the content from your 'Dashboard Message' page.
4. Save the welcome message template. The dashboard will display the content of the template immediately, including the content of your 'Dashboard Message' page. Similarly, if you or anyone else edits the page, the welcome message on the dashboard will change 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 like [JIRA](#). Linking two applications allows you to share information and access one application's functions from within the other. For example, you could use the [JIRA Issues Macro](#) to display JIRA issues on a Confluence page.

Related Topics

- [Linking to Another Application](#)
- [Configuring Workbox Notifications](#)
- [Integrating JIRA and Confluence](#)
- [Registering External Gadgets](#)

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.

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 your JIRA server and view JIRA issues in 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. In the **Application URL** box, supply the URL of the application you want to link to and then select **Create new link**.
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 so the link can be set up. This information is required 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 application types. After the link has been set up, it will appear on the "Configure Application Links" page. You can use this page to change the configuration of application links to make them more secure or to change the link settings:

- To edit the settings of the application link (for example, to [change the authentication type](#) of the link), select **Edit**.
- If you've set up multiple links to the same type of application (for example, multiple JIRA servers), you can use the **Make Primary** link to specify which application is the default instance. See [Making a Primary Link for Links to the Same Application Type](#) for more information.
- After you've linked applications, you also connect the areas of those applications that contain information relating to your project or team (for example, you can connect a project's Confluence space with a JIRA

project). These types of links are called [project links](#).

Configuring Workbox Notifications

People can view and manage in-app notifications and tasks in their [Confluence workbox](#). This page tells you how to enable in-app notifications and configure some related settings.

In addition, people can receive notifications from JIRA and other Confluence servers in their 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 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](#)

Related pages:

- [Managing Notifications in Confluence](#)
- [Confluence Administrator's Guide](#)

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 JIRA, 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.

Save

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	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. 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.

Including notifications from JIRA

Confluence workbox can include notifications from your JIRA issue tracker. In Confluence OnDemand, you can do this if you have JIRA OnDemand too.

To include notifications from JIRA:

JIRA 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 server will appear in the list of linked applications below this option.
 - People will see JIRA notifications in their workbox, as described in [Managing Notifications in Confluence](#).

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 from sending notifications to Confluence

You may wish to configure Confluence to display its own notifications in its workbox, but prevent notifications from JIRA from appearing in the workbox, even when JIRA 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 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 *ConfluenceChatty*. (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
☒ 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

Please refer to the guide to [Installing Confluence and JIRA Together](#).

JIRA and Confluence are designed to complement each other. Collect your team's thoughts, plans and knowledge in Confluence, track your issues in JIRA, and let the two applications work together to help you get your job done.

Below are some ways you can get JIRA and Confluence working together.

Setting Up Trusted Communication between JIRA and Confluence

An administrator can configure JIRA (3.12.0 or later) and Confluence to communicate in a trusted way, so that Confluence can request information from JIRA on behalf of the currently logged-in user. JIRA will not ask the user to log in again or to supply a password.

Trusted communication is used when embedding information from one application (for example, a list of JIRA issues) into another application (for example, a Confluence page).

Read more about [trusted communication](#).

Inserting JIRA issues

You can insert issues from a JIRA site onto your Confluence page using the 'Insert JIRA Issue' dialogue box. You can also use this dialogue box to create a new issue on the JIRA site. See the [JIRA Issues Macro](#).

On this page:

- [Setting Up Trusted Communication between JIRA and Confluence](#)
- [Inserting JIRA issues](#)
- [Viewing Confluence Content in JIRA or JIRA Content in Confluence](#)
- [Integrating JIRA and Confluence User Management](#)
- [Useful Plugins](#)

Related pages:

- [Integrating Confluence with Other Applications](#)
- [Confluence Administrator's Guide](#)



The information on this page does not apply to Confluence OnDemand.

Viewing Confluence Content in JIRA or JIRA Content in Confluence

Using Gadgets

You can embed a Confluence activity stream or a Confluence page in JIRA's dashboard. Likewise, JIRA gadgets can be rendered on a Confluence page. See [Adding a Confluence Gadget to a JIRA Dashboard](#) and [Gadget Macro](#) for information on how to set up gadgets.

Using the JIRA Issues macro

For versions earlier than Confluence 3.1 and JIRA 4.0, use the `{jiraissues}` macros to embed JIRA reports and portlets into your Confluence site

Any JIRA search result can be embedded in a Confluence page using the [JIRA Issues macro](#) with your choice of included fields and field ordering, and any JIRA gadgets can be embedded in a Confluence page by [Registering External Gadgets](#).

Integrating JIRA and Confluence User Management

To save you having to enter users into both JIRA and Confluence, you may benefit from using [Atlassian Crowd](#) as the user repository for both applications. Alternatively you can configure Confluence to use JIRA's user database. See [Connecting to Crowd or JIRA for User Management](#).

Useful Plugins

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](#).

The [JIRA Linker plugin](#) provides a custom field that helps you find an URL, particularly a Confluence page, so you can add a page link into a JIRA issue.

Installing Confluence and JIRA Together

This page describes Atlassian's recommendation for installing JIRA and Confluence on the same server. Refer to [Here Be Dragons](#) for instructions on integrating all Atlassian applications.



Do not deploy multiple Atlassian applications in a single Tomcat container —

Deploying multiple Atlassian applications in a single Tomcat container is **not supported**. We do not test this configuration and upgrading any of the applications (even for point releases) is likely to break it. There are also a number of known issues with this configuration (see [this FAQ](#) for more information).

We also do not support deploying multiple Atlassian applications to a single Tomcat container for a number of practical reasons. Firstly, you must shut down Tomcat to upgrade any application and secondly, if one application crashes, the other applications running in that Tomcat container will be inaccessible.

Finally, we recommend not deploying *any other applications* to the same Tomcat container that runs Confluence,

especially if these other applications have large memory requirements or require additional libraries in Tomcat's `lib` subdirectory.

 The information on this page *does not apply to* Confluence OnDemand.

Recommended Setup - Separate Stand-Alone Installations

Atlassian recommends running JIRA and Confluence in separate stand-alone instances running behind an Apache Web Server. See the guides for:

- [Installing Confluence](#)
- [Running Confluence behind Apache](#)
- [Installing JIRA](#)
- [Integrating JIRA with Apache](#)

Advantages

- Each application can be restarted without affecting the other.
- If one webapp hangs for any reason (eg. running out of memory), it doesn't affect the other.
- Any problems can be debugged more easily. Logs are separate and product-specific, rather than everything going to `catalina.out`. Thread and heap dumps are smaller and more relevant.
- It reduces the likelihood of jar conflicts (eg. jars that must be installed in `common/lib` or `lib` for Confluence running off Apache Tomcat version 6 or above), particularly if you later want to install a third webapp not from Atlassian.
- Apache HTTP Web Server is well suited for running publicly available sites, with extensive modules for security and efficiency. It also allows for flexibility with URLs (ie <http://confluence.atlassian.com>, <http://confluence>, and so on).

 Apache Web Server is recommended and reliable. It is also a third-party product, and therefore not developed nor supported by Atlassian. See [Atlassian Support Offerings](#) for details.

Setting Up Trusted Communication between JIRA and Confluence

An administrator can configure [JIRA](#) and Confluence to communicate in a trusted way, so that Confluence can request information from JIRA on behalf of the currently logged-in user. JIRA will not ask the user to log in again or to supply a password.

Trusted communication is used when embedding information from one application (for example, a list of JIRA issues) into another application (for example, a Confluence page).

 **Potential security risk**
Do not configure a trusted application unless you trust all code in that application to behave itself at all times. Trusted communication uses public/private key cryptography to establish the identity of the trusted server, so you must also be sure that the trusted application will maintain the security of its private key. Read the details of the [security risks](#) below.

Prerequisites

The following setup is required:

- JIRA 4.2.0 or later.
- Confluence 3.5.0 or later.
- In order to authenticate successfully against JIRA, the Confluence user must also be registered as a JIRA user with the same username.

Note: It is highly recommended that your JIRA and Confluence instances share a **common user base**, rather than two separate user bases with duplicated usernames. You will receive an error if Confluence passes JIRA a username which JIRA cannot recognise. Also, with separate user bases you run the risk that the same

username may be used by two different people. The trusted application does not supply the user's password, so the trusting application will assume the username belongs to the user registered in the trusting application's own user base.

Tip: Try [Atlassian Crowd](#) for a tidy user management solution.

On this page:

- [Prerequisites](#)
- [Why do we need Trusted Communication?](#)
- [Overview](#)
- [Configuring JIRA to trust Confluence](#)
- [Adding the macro to a Confluence page](#)
- [Viewing the Confluence page](#)
- [Security Risks](#)
- [Troubleshooting](#)

Related pages:

- [JIRA Issues Macro](#)
- [Troubleshooting the JIRA Issues Macro and Trusted Applications](#)
- [Connecting to LDAP or JIRA or Other Services via SSL](#)
- [Confluence Administrator's Guide](#)

 *The information on this page does not apply to Confluence OnDemand.*

Why do we need Trusted Communication?

The [JIRA Issues](#) macro allows you to embed a list of JIRA issues into a Confluence page. Prior to Confluence 2.7, if you wanted to display JIRA issues that had restricted viewing, then you needed to store the JIRA user's credentials (username and password) in the macro code directly on the Confluence page. This was not very secure.

The reasons we require the user credentials are:

- Your JIRA instance might not be public, and you might not want to allow anonymous access to your issues.
- You might have security restrictions on some of your issues. You may not want to allow someone to leak data from your JIRA project by using the JIRA Issues Macro on a Confluence page.

Overview

Here is a summary of the integration points in a trusted communications relationship. Each of the following points is described in more detail in the sections below.

- A JIRA or Confluence system administrator [configures JIRA to trust Confluence](#).
- A Confluence user [adds one of the macros to a Confluence page](#).
- A Confluence user or anonymous user [views the Confluence page](#).

Configuring JIRA to trust Confluence

Trust only has to be established once between the two applications. Once trust has been established, it is entirely transparent to the Confluence users.

You can use Application Links to enable trust relationships between two applications. Linking two applications allows you to share information and access one application's functions from within the other.

You can configure an application link to use Trusted Applications as the authentication mechanism. For instructions, see [Linking to Another Application](#).

Adding the macro to a Confluence page

The Confluence user can add and edit the macros as described on the following page: [JIRA Issues macro](#).

The following options are available for determining the issues which will be retrieved from JIRA and displayed on the Confluence page:

What you want to do	Macro parameter	URL parameter	Comments
Display the JIRA issues which the logged-in user is authorised to see. And if the user is not logged in, display only issues which allow unrestricted viewing.			Do not specify any authentication parameters. In this case, the behaviour depends on the way your administrator has set up trusted communication between JIRA and Confluence. Here is a summary of the behaviour. If trusted communication is enabled , the authorisation will work seamlessly. When a logged-in user views your page, they will see only the JIRA issues they are allowed to see. And if they are not logged in, they will see only the issues which allow unrestricted viewing. If trusted communication is disabled , the Confluence page will show only the JIRA issues which allow unrestricted viewing.
Ensure that Confluence will display only the JIRA issues which allow unrestricted viewing.	anonymous		Regardless of who the user is (logged in or not), the Confluence page will show only anonymously-visible issues. Confluence will not attempt to set up a trusted communication link with JIRA in this case.
Use a pre-determined username and password to access the JIRA issues.		&os_username=MYNAME&os_password=MYPASSWORD	Not recommended. Prior to Confluence 2.7, this was the only way of displaying issues with restricted viewing. For Confluence 2.7 and later, this method will still work. Confluence will not attempt to set up a trusted communication link with JIRA in this case.

Viewing the Confluence page

When a user views a Confluence page which contains a JIRA Issues macro, this is what happens:

- If the macro markup contains an explicit username and password in the URL parameter, Confluence will not request trusted communication with JIRA. Confluence will retrieve the JIRA issues which the specified username is authorised to see. This behaviour is the same as Confluence versions prior to 2.7.
- If the macro markup contains the anonymous parameter, Confluence will retrieve only the JIRA issues which allow unrestricted viewing. Confluence will not attempt to set up a trusted communication link with JIRA in this case.
- If the user is anonymous (not logged in), Confluence will retrieve only the JIRA issues which allow unrestricted viewing. Confluence will not attempt to set up a trusted communication link with JIRA in this case.
- If the user is logged in, then Confluence attempts trusted communication with JIRA. Confluence sends the username to JIRA. JIRA returns a set of issues which that username is authorised to access, based on the JIRA user base and the JIRA groups and permissions. Confluence displays those issues on the page.
- If JIRA or Confluence encounters a problem during the trusted communication process, an error message may appear on the Confluence page above the macro output – see [.Setting Up Trusted Communication between JIRA and Confluence v5.3#troubleshooting](#) below.

Security Risks

Please take the following considerations into account when setting up trusted communication:

- When you configure JIRA to trust an application, you are allowing the application to access JIRA in the name of a particular user. The trusted application passes JIRA the user's login name, but no other authentication information. JIRA does not request the user's password. By doing this, you are **bypassing JIRA's authentication mechanism**.
- Do not configure a trusted application unless you **trust all code in that application** to behave itself at all times.
- Trusted communication uses public/private key cryptography to establish the identity of the trusted server. The trusted application needs to maintain the security of its private key. Confluence stores its private key in the database. **So you must be sure that the Confluence database is secure, and also any full backups of the database.**
- Ensure that you **specify an IP address** for your Confluence site when configuring trusted applications in JIRA. Do not use the wild card *. *. *. * as the IP address. Failure to configure IP address restrictions is a security vulnerability, allowing an unknown site to log into your JIRA site under a user's login ID.
- Be aware of the risks associated with using separate user bases, as explained [above](#). **We strongly recommend a common user base between the trusted and trusting applications.**
- When configuring an application to trust another application, you should use a trusted network or SSL to **protect the sensitive information passed between the applications during the configuration procedure**. This will help to prevent [man-in-the-middle attacks](#).

Troubleshooting

Below are the warning messages which may appear on your Confluence page, above the output of the JIRA Issues macro.

Warning Message	Cause	Solution	Warning Message Can be Turned Off?

javax.net.ssl.SSLHandshakeException: sun.security.validator.ValidatorException: PKIX path building failed: sun.security.provider.certpath.SunCertPathBuilderException: unable to find valid certification path to requested target	JIRA is running over SSL	Add JIRA's SSL Certificate to the Java Keystore	No
The JIRA server does not recognise your user name. Issues have been retrieved anonymously.	The logged-in Confluence user is not registered in the JIRA user base.	Add the username to your JIRA user base. It is highly recommended that at your JIRA and Confluence instances share a common user base.	No
The JIRA server does not trust this Confluence instance for user authentication. Issues have been retrieved anonymously. You can set the macro to always use an anonymous request by setting the 'anonymous' parameter to 'true'.	Your JIRA instance has not been configured to trust your Confluence instance.	One of the following solutions: • Configure JIRA to trust Confluence. • Disable trusted communications for the JIRA macros in Confluence. • Use the anonymous parameter in all your JIRA Issues macros.	Yes
The JIRA server does not support trust requests. Issues have been retrieved anonymously. You can set the macro to always use an anonymous request by setting the 'anonymous' parameter to 'true'.	Your JIRA instance is not able to handle trusted communications (i.e. the JIRA version is earlier than 3.12.0).	One of the following solutions: • Download the latest version of JIRA and then configure JIRA to trust Confluence. • Disable trusted communications for the JIRA macros in Confluence. • Use the anonymous parameter in all your JIRA Issues macros.	Yes

Failed to login trusted application: confluence:14159892 due to: com.atlassian.security.auth.trustedapps.CertificateTooOldException: OLD_CERT; Certificate too old.	There is a date/time difference between the JIRA server and Confluence server.	• Certificate Too Old KnowledgeBase Entry	-

Consult [Troubleshooting the JIRA Issues Macro and Trusted Applications](#) for further troubleshooting.

Registering External Gadgets

You can register gadgets from external web sites (such as [JIRA](#), [iGoogle](#) or [Gmail](#)) with your Confluence installation, so that the gadgets appear in the [macro browser](#) and people can add them to Confluence pages via a [gadget macro](#).

Choose one of the following ways to register the external gadgets on Confluence:

- **Subscribe to all of the external 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.
- **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](#).

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. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **External Gadgets** in the left-hand panel.
3. **Click** the **Gadget Feeds** tab.
4. Enter the base URL of the application you want to subscribe to, in the text box labelled **Gadget Feed URL**. For example, `http://example.com/jira` or `http://example.com/confluence`.
5. Choose **Add**. Confluence will convert the URL to a gadget feed and place it in the list of 'Added Gadget

Feeds'.

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](#)
- [Adding JIRA Gadgets to a Confluence Page](#)
- [Linking to Another Application](#)

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	Delete
https://pug.jira.com/rest/gadgets/1.0/g/feed/	Delete

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 obtain that gadget's URL and copy it to your clipboard.

Getting a gadget's URL from an Atlassian application

If your web application is another Atlassian application such as Confluence or JIRA:

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 obtain the gadget's 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. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **External Gadgets** in the left-hand panel.
3. Paste your gadget's URL into the **Gadget Specification URL** field in the 'Add a new Gadget' section.
4. 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.

Managing your Confluence License

The license on your Confluence site entitles you to run Confluence and to have Atlassian support for a specified period. It also defines the number of users who are entitled to log in to the Confluence site.

Read how to [find the details of your existing license](#), and [get a Confluence license](#) if you do not have one already.

Are too many people authorised to use your site, exceeding the number allowed by the license? Try [reducing the user count](#), or see the [licensing and pricing overview](#) on the Atlassian website if you want to upgrade to a higher user count.

You may also need to find the [support entitlement number \(SEN\)](#) when dealing with the Atlassian support team.

To quickly check the status of your licence you can also go to  > **General Configuration** > **Atlassian Support Tools** > **Health Check**.

Related pages:

- [Upgrading Beyond Current Licensed Period](#)
- [Confluence Installation and Upgrade Guide](#)
- [Confluence Administrator's Guide](#)

Viewing and Editing License Details

When you upgrade or renew your Confluence license, you will receive a new license key. You will need to update your Confluence installation with the new license key.

You can access your existing license key, or generate an evaluation license key, at <http://my.atlassian.com>.

Updating your license details in Confluence

To update your Confluence license:

1. If you do not already have a license key, get your existing license key, or generate an evaluation license key, at <http://my.atlassian.com>.
2. Log in to Confluence as a user with Confluence Administrator or System Administrator permissions.
3. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
4. Choose **License Details** in the left-hand panel.
5. Enter your new license details into the **License** field.
6. Choose **Save**.

If you are running a Confluence cluster, you will need to:

- Update each server's Confluence license separately.
- Ensure that the new license has enough nodes to cover all servers that are currently running in your cluster. To check the number of active servers in your cluster, see the [Cluster Administration](#) page.

On this page:

- [Updating your license details in Confluence](#)
- [Viewing your license details](#)
- [Understanding the user count for your Confluence license](#)
- [Exceeding your licensed user count](#)
- [Downgrading your Confluence license to pay for fewer users](#)

Related pages:


- [Reducing the User Count for your Confluence License](#)
- [Managing Confluence Users](#)
- [Confluence Administrator's Guide](#)

Screenshot: License details

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
Date Purchased	Jan 25, 2013
License Type	Confluence: Evaluation
Licensed Users	Unlimited
License Expiration	Your evaluation expires in 5 days, 7 hours .
Support Entitlement Number	SEN-L21804
Server ID	BQMH-JCP3 -JCP3 -JCP3 (Atlassian sales or support may ask you to provide this ID)
License	 Save

Viewing your license details

To view the details of your Confluence license:

1. Log in to Confluence as a user with Confluence Administrator or System Administrator permissions.
2. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
3. Choose **License Details** in the left-hand panel.

The 'License Details' screen tells you:

- What type of license you have (for example: Commercial, Academic, Community, or Evaluation).
- How many users your Confluence site is licensed to support, and how many are currently registered ('signed up currently'). See [below](#) for more about the user count.
Choose **Refresh** to make sure you see the latest count.
- How much time remains in your one-year support and upgrades period (for full licenses) or 30-day trial (for trial licenses).
- Your server ID, which:
 - is generated when you install Confluence for the first time
 - exists for the life of the Confluence installation

- survives an upgrade
- is held in the database
- is not bound to a specific license
- is the same for all servers in a cluster.

Understanding the user count for your Confluence license

The number of registered users allowed on your Confluence site may be limited, depending on your license type. See the [licensing and pricing overview](#) on the Atlassian website. If you have an 'unlimited' license, then the number of registered users is not significant.

The number of registered users is also called the 'user count' or the number of users 'signed up currently'. It is determined as follows:

- It includes only those users who have the 'can use' global permission for the Confluence site. (See [Global Permissions Overview](#) for more about the 'can use' permission.)
- It does not include anonymous users, who may access your Confluence site if you have allowed anonymous access. (See [Setting Up Public Access](#) for more about allowing anonymous access.)
- It does not include [deactivated users](#).

Exceeding your licensed user count

If you exceed the number of users included in your licence, 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 - see [Reducing the User Count for your Confluence License](#).

Downgrading your Confluence license to pay for fewer users

If you want to downgrade your Confluence license to one which allows fewer users, please make sure first that your new license covers your current user count.

- View your license details as described [above](#).
- Check whether the number of users 'signed up currently' is lower than the number allowed by the new license.
- If you currently have more users signed up than the new license allows, please follow these [instructions on reducing the user count](#).
- When the number of users 'signed up currently' is lower than the number allowed by your new license, you can add the new license key to Confluence as described above.

Getting a Confluence License

Need a Confluence license or license key?

- If you do not yet have a license, you can get a [free multi-user evaluation license](#) or a [10-user starter license](#) immediately.
- If you already have a Confluence license, you can retrieve your key or generate a new key from the [license viewer](#).
- For enterprise, non-profit, open source and educational licenses, see [Confluence licensing and pricing](#).
- If you cannot find your key or are having problems, contact sales@atlassian.com.

Related pages:

- [Viewing and Editing License Details](#)
- [Reducing the User Count for your Confluence License](#)
- [Confluence Administrator's Guide](#)

Reducing the User Count for your Confluence License

This page tells you how to reduce the number of users that count towards your Confluence license. You may want to reduce your user count in Confluence if you have exceeded your license limit, or if you want to change to a lower-tier license to reduce costs.

Understanding the user count for your Confluence license

The number of registered users allowed on your Confluence site may be limited, depending on your license type.

See the [licensing and pricing overview](#) on the Atlassian website. If you have an 'unlimited' license, then the number of registered users is not significant.

The number of registered users is also called the 'user count' or the number of users 'signed up currently'. It is determined as follows:

- It includes only those users who have the 'can use' global permission for the Confluence site. (See [Global Permissions Overview](#) for more about the 'can use' permission.)
- It does not include anonymous users, who may access your Confluence site if you have allowed anonymous access. (See [Setting Up Public Access](#) for more about allowing anonymous access.)
- It does not include [deactivated users](#).

On this page:

- [Understanding the user count for your Confluence license](#)
- [Reducing the user count](#)

Related pages:

- [Viewing and Editing License Details](#)
- [Managing Confluence Users](#)
- [Confluence Administrator's Guide](#)

Reducing the user count

The recommended method for reducing your user count is to remove or deactivate the users. You can remove users who do not require access to Confluence and have never created content in Confluence. You can deactivate users who have created content but no longer require access to Confluence. See [Deleting or Deactivating Users](#).

Alternatively, if you have [connected Confluence to an LDAP directory](#), you may want configure Confluence to synchronise a subset of users from LDAP rather than all users. This is described in the following knowledge base article: [Changing the Number of Users Synchronized from LDAP to Confluence](#). This can be a complicated procedure and we recommend that you do not use this method unless necessary.

Finding Your Confluence Support Entitlement Number (SEN)

There are three ways to find your Support Entitlement Number (SEN):

▼ [Method 1: Check in the Confluence Administration Interface](#)

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **License Details** in the left-hand panel. The SEN is shown:

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
Date Purchased	Jan 25, 2013
License Type	Confluence: Evaluation
Licensed Users	Unlimited
License Expiration	Your evaluation expires in 5 days, 7 hours .
Support Entitlement Number	SEN-L21804
Server ID	BQMH-JCP3 -JCP3 -JCP3 (Atlassian sales or support may ask you to provide this ID)
License	Save

▼ [Method 2: Log into my.atlassian.com as the Account Holder or Technical Contact](#)

Your Support Entitlement Number is available in <http://my.atlassian.com>:

Support Entitlement Number

Welcome

Hey [redacted] you've been an Atlassian customer for 3 years now. Thank you for your support! You can manage your 8 licenses below. For questions relating to managing your account, licensing or purchasing, please see the [purchasing and licensing FAQ](#) or [contact a customer service representative](#). To request technical support, please visit our [Support Portal](#).

Licenses [New Evaluation License](#) | [Expand All](#) | [Collapse All](#)

Product	Name	Support Expires	Support
Confluence 2 Users: Personal Server	[redacted]	02 Jun 2009	
Confluence 500 Users: Commercial License	Atlassian	10 Apr 2011	Request Support

Confluence 500 Users: Commercial License

Name: Atlassian (edit)

SEN: SEN- [redacted]

Server ID: [redacted] (edit)

Technical Contact: [Jeremy Largman \(jlargman@atlassian.com\)](#)

Email:

[Add](#)

Billing Contact: [Jeremy Largman \(jlargman@atlassian.com\)](#)

Email:

[Add](#)

Licensed To: Atlassian

License: [redacted]

This license key is compatible with Confluence 3 or above. If you are using an older version of Confluence, you can [downgrade your license key](#).

Actions: [Renew](#) | [Upgrade](#) | [Download Source](#) | [View Developer License](#)

History:

10 Apr 2008	Initial Purchase
10 Apr 2009	Maintenance Renewal
01 Jun 2009	Converted to Confluence 500 Users: Commercial License
10 Apr 2010	Maintenance Renewal

Method 3: Atlassian Invoice

Your Support Entitlement Number (SEN) appears on the third page of your Atlassian invoice.

See [Finding Your Support Entitlement Number](#) in the support space for more general information about how Atlassian Support uses this number.

! The information on this page *does not apply to* Confluence OnDemand.

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](#)

Related pages:

- [Managing Add-ons and Macros](#)
- [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.

The embedded HSQLDB database for evaluation purposes

The Confluence installation includes an embedded [HSQLDB](#) database, supplied for the purpose of **evaluating Confluence**.

If you are using the embedded database, the database files are stored in the \database directory under your [Confluence Home Directory](#). See also [Important Directories and Files](#).

Note: The embedded HSQLDB database is not suitable for production Confluence sites.

Production sites should use an external database. See our [guide to database configuration](#). When using the default HSQLDB database, you run the risk of irrecoverable data loss because HSQLDB is not [transaction safe](#).

- Corruption is occasionally encountered after sudden power loss. It can usually be corrected using the data recovery procedure [documented in our knowledge base](#).
- HSQLDB is suitable for evaluation purposes, but the risk can only be eliminated by switching databases. This is essential when you move from an evaluation to a production site. External databases may also provide superior speed and scalability.

To find out if you are still using the embedded database, go to  > **General Configuration** > **Atlassian Support Tools** > **Health Check**.

On this page:

- [The embedded HSQLDB database for evaluation purposes](#)
- [Selecting an external database](#)
- [Database setup](#)
- [Optimising database performance](#)
- [Database troubleshooting](#)
- [Notes](#)

Related pages:

- [Database JDBC Drivers](#)
- [Supported Platforms](#)
- [Embedded HSQLDB Database](#)
- [Managing Confluence Data](#)
- [Confluence Administrator's Guide](#)

Selecting 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 will 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](#)
- [List Of Supported Databases](#)
- [Database Troubleshooting](#)

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](#)

Optimising database performance

To improve database responsiveness:

- [Improving Database Performance](#)

- [Database Troubleshooting](#)

Database troubleshooting

For solving database-related problems:

- [Troubleshooting External Database Connections](#)
- [Troubleshooting the Embedded HSQLDB Database](#)
- [Interpreting 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 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](#)
- [Confluence Administrator's Guide](#)

Database	JDBC driver bundled with Confluence?	JDBC drivers	More information
PostgreSQL		8.4-701.jdbc3 The JDBC 3 driver will work under the 1.6 JVM. If you want to use the JDBC 4 driver, you can download it from the PostgreSQL website . However, we recommend that you use the bundled JDBC 3 driver.	Database Setup for PostgreSQL

Microsoft SQL Server		jTDS 1.2.2 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		5.1.11 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 • For Oracle 11.1, use the 10.2.0.4 or 11.1.0.7.0 driver (Java 6 ojdbc6.jar). • For Oracle 11.2, use the 11.2.0.1.0 driver (Java 6 ojdbc6.jar). We recommend using the thin drivers only. Tip: Search for the JAR file name on the download site. See the Oracle JDBC driver FAQ.	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=HTTP
  )(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=HTTP.
- 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 available on this page: [Database JDBC Drivers](#).
2. Copy the driver JAR file into the lib folder of your Tomcat installation: <TOMCAT-INSTALLATION>/lib.

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](#)
- [Important Directories and Files](#)
- [Confluence Installation and Upgrade Guide](#)

Step 3. Configure Tomcat

1. 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.
2. 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>
```

3. 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 -->

    <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"
      maxActive="25"
      maxIdle="5"
      maxWait="10000"
    />

    <Manager pathname="" />
  </Context>
</Host>
```

4. Change the username and password to match your Oracle login.
5. 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
```

6. If required, choose different `maxActive` 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](#).

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](#)

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.

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**). 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](#).
- 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> a

nd 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.

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](#)
- [Important Directories and Files](#)
- [Confluence Installation and Upgrade Guide](#)

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`:

```
<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"
  maxActive="20"
  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.
 - `maxActive` — 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.
 - `maxWait` — 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.SQLServerIntJdialect</proper
ty>
<property
name="hibernate.connection.datasource">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 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](#).
- 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 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](#).

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 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:

1. Run the MySQL installation wizard:
 - a. If you are connecting Confluence to your existing MySQL server, choose **Reconfigure Instance**.
 - b. Choose **Advanced Configuration**.
 - c. 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.
 - d. 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.*
 - e. Set the **InnoDB Tablespace** settings to your requirements. (The default settings are acceptable.)
 - f. 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.
 - g. 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.
 - h. 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 (often named `my.ini` on Windows operating systems) in your MySQL server. 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 32M:

```
[mysqld]
...
max_allowed_packet=32M
...
```

- 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
```

(Refer to [MySQL Option Files](#) for detailed instructions on editing `my.cnf` and `my.ini`.)

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**.

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`).

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 `<CONFLUENCE_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 Tomcat installation: `<TOMCAT - INSTALLATION>/lib`.

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](#)
- [Important Directories and Files](#)
- [Confluence Installation and Upgrade Guide](#)

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`:

```
<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"
  maxActive="15"
  maxIdle="7"
  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.
- If you plan to use non-Latin characters, add `"&useUnicode=true&characterEncoding=utf8"` on the end of the above URL.

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.
 - `maxActive` – 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.
 - `maxWait` – 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.
 - `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[net.sf.hibernate.dialect.MySQLDialect]]></pr
operty>
<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`.

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](#)
- [Important Directories and Files](#)
- [Confluence Installation and Upgrade Guide](#)

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:

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
  username="postgres"
  password="postgres"
  driverClassName="org.postgresql.Driver"
  url="jdbc:postgresql://localhost:5432/yourDatabaseName"
  maxActive="20"
  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.
 - maxActive — 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.
 - maxWait — 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 HSQLDB Database

The Confluence installation includes an embedded [HSQLDB](#) database, supplied for the purpose of **evaluating Confluence**.

If you are using the embedded database, the database files are stored in the \database directory under your [Confluence Home Directory](#). See also [Important Directories and Files](#).

Note: The embedded HSQLDB database is not suitable for production Confluence sites.

Production sites should use an external database. See our [guide to database configuration](#). When using the default HSQLDB database, you run the risk of irrecoverable data loss because HSQLDB is not [transaction safe](#).

- Corruption is occasionally encountered after sudden power loss. It can usually be corrected using the data recovery procedure [documented in our knowledge base](#).
- HSQLDB is suitable for evaluation purposes, but the risk can only be eliminated by switching databases. This is essential when you move from an evaluation to a production site. External databases may also provide superior speed and scalability.

To find out if you are still using the embedded database, go to  > **General Configuration** > **Atlassian Support Tools** > **Health Check**.

Related pages:

- [Important Directories and Files](#)
- [Database Configuration](#)

Connecting to HSQLDB using DBVisualizer

The purpose of this guide is to walk you through connecting to Confluence's embedded [Hypersonic SQL Database](#) using the Database Administration tool [DBVisualizer](#).

Below are step by step instructions on how to Configure DBVisualizer and connect it to HSQLDB.

Prerequisites

1. Download and install the latest copy of [DBVisualizer](#).


2. You will also need to download a copy (preferably the latest version) of [HSQLDB](#)
3. Extract the contents of the HSQLDB archive
4. Ensure that Confluence is not running.

Connection Procedure

Please ensure that you read and follow the instructions below carefully.


 **Remember to backup your <confluence-home>/database folder before attempting any modifications**

1. Enter Connection Name



1. Click on the icon highlighted in **Red**
2. Enter an identifiable name for the connection. e.g. conf2.5.4-std

2. Select JDBC Driver




1. From the drop down list select **HSQLDB Embedded**

2. Click on Load Driver Files
3. Browse to directory where the HSQLDB .jar file is located. Confluence bundles this and it can be found at <confluence-installation>/confluence/WEB-INF/lib/hsqldb-*.jar.

3. Select Database Path

1. Browse to your <Confluence-Home> directory
2. Open the **Database** folder
3. Select the conflencedb.properties file


4. Enter Connection Details



1. Remove the ".properties" from the end of conflencedb
2. Type in sa for the username
3. Leave the password field blank

refer to the example screenshot above if you are unsure

5. Connect to embedded Database



1. Click on Test Connection to verify that the details are correct.
2. Click on **"Finish"** to complete the setup
3. Select the connection from the list on the left hand side.
4. You can now click on **"Connect"** to connect to the embedded database.

HSQL database manager

Alternatively, you can use [HSQLDB's database manager](#). Just copy the value of `hibernate.connection.url` in `confluence.cfg.xml` as the URL and you're good to go.

Related Topics

[Universal SQL client Squirrel](#)

[HSQL](#)

[Enable Hibernate Logging](#)


[Database Tables Reference](#)

[Confluence data model](#)

[Database Tables Reference](#)

Below is a diagram of the Table References in Confluence (2.5.4).

This may be useful for Database Administrators that need to manually create the Database tables.



Right Click and Select *Save Link As* [here](#) to download this image.

Troubleshooting the Embedded HSQLDB Database

Note: HSQLDB should not be used as a production database. It is included for evaluation purposes only. For more information, see [Embedded HSQLDB Database](#).

Resolving the error: "User not found: SA"

Please refer to our [knowledge base article](#).

Hibernate logging

You may find it useful to [enable detailed Hibernate logging](#) when debugging problems with HSQLDB.

Connecting to HSQLDB

You may need to connect to the database to retrieve information, or for troubleshooting purposes. Please follow the instructions on [Connecting to HSQLDB using DBVisualizer](#).

Related pages:

- [Database Configuration](#)
- [Confluence Administrator's Guide](#)

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:

- You should use this page when moving from the embedded database to an external database.
- Provided your dataset is not large, you may use this method to move from one type of external database to another. For example, from Oracle to PostgreSQL.
- Similarly, if the dataset is not large, you can use this method to upgrade your database to a new version of the same database. Note: You do not need to migrate your data if you are upgrading the database in place.

Note: If you are simply moving your database from one server to another you can just 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).

Limitations of database migration

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.

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
- Troubleshooting

Related pages:

- Database Configuration
- Confluence Installation Guide
- About Add-ons
- Confluence Home Directory
- Confluence Administrator's Guide

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. See [Adding JIRA Gadgets to a Confluence Page](#).
- Review any other resources that other systems are consuming from Confluence.

Method two – for installations with a large volume of attachments

Before you start

Before proceeding with these instructions please check the following.

- Your existing installation must be Confluence 2.2 or later.
- Your attachments must be stored in the file system, *not in your database*. (To migrate between attachment storage systems, see [Attachment Storage Configuration](#).)

The instructions below will only work if both of the above are true.

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. See [Adding JIRA Gadgets to a Confluence Page](#).
- 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

If you are unable to restore your XML backup, consult our [troubleshooting guide](#).

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.

On this page:

- [JDBC connection settings](#)
- [Creating a UTF-8 database](#)
- [Updating existing database to UTF-8](#)

 *The information on this page **does not apply** to Confluence OnDemand.*

JDBC connection settings

MySQL

Append "useUnicode=true to your JDBC URL:

```
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/databas  
e?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;
```

- You will also need to set the Server Characterset 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`

- 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
```

- 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 a 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.

RELATED TOPICS:

[Configuring Character Encoding](#)

[Known Issues for MySQL](#)

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="transactionManager"
class="org.springframework.orm.hibernate.HibernateTransactionManager">
  <property name="sessionFactory">
    <ref bean="sessionFactory"/>
  </property>
  <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, migration 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.

Improving Database Performance

Diagnosis

Use native database tools to assess the impact of your database. If you'd like to check what Confluence is doing from it's side, you can [enable SQL logging](#). If you analyze [thread dumps](#), as this is done in general [Troubleshooting Confluence Hanging or Crashing](#) guide, you may find the kinds of threads like this:

```
"http-8080-Processor150" daemon prio=1 tid=0x08543368 nid=0x11aa in
Object.wait() [0x665a4000..0x665a51b0]
at java.lang.Object.wait(Native Method)
- waiting on <0x83140488> (a
com.mchange.v2.resourcepool.BasicResourcePool)
at
com.mchange.v2.resourcepool.BasicResourcePool.awaitAcquire(BasicResource
Pool.java:968)
at
com.mchange.v2.resourcepool.BasicResourcePool.checkoutResource(BasicResourcePool.java:208)
- locked <0x83140488> (a com.mchange.v2.resourcepool.BasicResourcePool)
```

These threads are waiting for a database connection. It could be that the database is not performing optimally, or it may just need tuning for allowing more connection threads. Both are discussed below.

Upgrade your Database and Drivers

SQL Server 2000, Oracle 9i, and MySQL with 3.1 drivers are among some of the issues with database performance. Ensure you are using updated versions of databases and their drivers.

Upgrade your hardware

Atlassian does not offer specific recommendations on hardware for database performance. Use good judgment and native OS and database tools for your assessment.

Ensure you have the Latest Database Indices

Confluence has improved database performance over time. You'll want to make sure you have all the latest, if you're getting hung threads waiting for db connections.

Confluence 2.10 or Manual .ddl Indices

With 2.10 and later, Confluence includes database indices bundled. Confluence 2.10 [automatically creates the necessary database indexes](#) when you upgrade. If you are not on 2.10, you may have run the ddl manually during the upgrade process. To check, you can look against [these](#).

Additional Indices not Included in 2.10

- One import db index is the **lower case page title index**. Prior to Confluence 3.0, querying for a page by title and space key can take a long time due to table scans necessary on a lowercase *where* clause. On most databases it is possible to add a lowercase index on these columns that helps with performance. See [Creating a Lowercase Page Title Index](#) for instructions on how to do this. Prior to 2.10, apply lowercase title indexes (all Confluence versions).
- The compound database index for the ATTACHMENTDATA table is described in [CONF-13819](#).
- A composite index on some of the columns in SpacePermissions table is described in [CONF-14488](#).

Tuning the Database Connection Pool

This is described in the knowledge base article [Confluence Slows and Times out During Periods of High Load due to DB Connection Pool](#).

Configure a Database Query Timeout

If a database is getting overloaded, you can prevent it from crashing Confluence by [Configuring a Database Query Timeout](#).

Related Articles

[Troubleshooting Database Issues](#).

Creating a Lowercase Page Title Index

Diagnosis

Confluence sometimes has performance problems retrieving pages by title because the query uses the lower() function. For example, the query looks something like this:

```
select * from CONTENT where lower(TITLE) = :title and SPACEID = :spaceid
```

Database profiling might show a query like the following taking a long time to execute (emphasis added):

```
select ... from CONTENT page0_, SPACES space1_
where page0_.CONTENTTYPE='PAGE'
and ((lower(space1_.SPACEKEY)= @P0 and page0_.SPACEID=space1_.SPACEID)
and(lower(page0_.TITLE)= @P1 )
and(page0_.PREVVER is null )and(page0_.CONTENT_STATUS='current' ))
```

Typically, databases don't use indexes when you use a function in a where clause; they do a table scan instead. This makes the performance of this query not ideal ([CONF-11577](#)).

Generic solution

On many databases (e.g. Oracle, PostgreSQL, DB2 for z/OS), it is possible to create the index using the normal "create index" syntax, just using the function instead of the column name.

```
create index CONFITLE_LOWER on CONTENT(lower(TITLE));
```


Sources:

- <http://www.postgresql.org/docs/current/static/sql-createindex.html>
- <http://asktom.oracle.com/tkyte/article1/>

SQL Server

On SQL Server, you can add a computed column to the database table and then add an index on this column.

```
alter table CONTENT add TITLE_LOWER as lower(TITLE);
create index CONFTITLE_LOWER on CONTENT(TITLE_LOWER);
```

Sources:

- [http://msdn.microsoft.com/en-us/library/aa258260\(SQL.80\).aspx](http://msdn.microsoft.com/en-us/library/aa258260(SQL.80).aspx)
- <http://blogs.msdn.com/psssql/archive/2009/03/09/how-to-use-computed-columns-to-improve-query-performance.aspx>

MySQL

It is not currently possible to create a lowercase index on MySQL. Confluence 3.0 includes some caching improvements which should alleviate this performance problem on this database.

Source:

- <http://dev.mysql.com/doc/refman/5.1/en/create-index.html>

Workaround for MySQL databases, using a **case-insensitive** collation:

Please check whether your MySQL database has been set to use case-sensitive or case-insensitive collation. The queries to check whether your database is set to case-insensitive collation are:

```
show full columns from content where field = 'title';
show full columns from spaces where field = 'spacekey';
```

If the **collation_name** is returned as **<encoding>_ci**, the **ci** indicates case-insensitive collation.

If the database has been set to use case-insensitive collation, you can try removing **lower** from the following queries, in your `ContentEntityObject.hbm.xml` file residing in your `<Confluence-Install>/confluence/WEB-INF/lib/confluence-2.x.x.jar/com/atlassian/confluence/core/`:

```
<query name="confluence.page_findLatestBySpaceKeyTitle"><![CDATA[
  from Page page
  where lower(page.space.key) = :spaceKey and
         lower(page.title) = :pageTitle and
         page.originalVersion is null and
         page.contentStatus = 'current'
]]></query>

<query
name="confluence.page_findLatestBySpaceKeyTitleOptimisedForComments"><![CDATA[
  from Page page
  left join fetch page.comments as theComments
  left join fetch theComments.children
  where lower(page.space.key) = :spaceKey and
         lower(page.title) = :pageTitle and
         page.originalVersion is null and
         page.contentStatus = 'current'
]]></query>
```

DB2 for Linux or Windows

DB2 supports [indexes on generated columns](#) which are used for queries with a matching predicate. You can implement it like this:

```
ALTER TABLE CONTENT ADD COLUMN TITLE_LOWER GENERATED ALWAYS AS (LOWER(TITLE));  
CREATE INDEX CONFTITLE_LOWER ON CONTENT(TITLE_LOWER)
```

Related pages

- [Improving Database Performance](#)
- [CONF-10030: Queries that use 'lower' do not use index because of case sensitivity](#)

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 restarting (or for Confluence WAR distributions, the application server running Confluence would need restarting).

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	select 1
Microsoft SQL Server	select 1
Oracle	select 1 from dual
PostgreSQL	select 1

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 Directory](#)
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"
    maxActive="20"
    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 OnDemand.*

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 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](#).

On this page:

- [Establishing a production system backup solution](#)
- [Which files need to be backed up?](#)
- [How do I restore?](#)
- [Other processes](#)

Related pages:

- [Site Backup and Restore](#)

 *The information on this page **does not apply** to Confluence OnDemand.*

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:

- **confluence.cfg.xml**

The rest of the directories will be auto-populated on start up. You may also like to backup these directories:

- **config** – if you have modified your ehcache.xml file.
- **index** – if your site is large or reindexing takes a long time – this will avoid the need for a full reindex when restoring.

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:

- thumbnails
- viewfile.

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.

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.

On this page:

- [Configuring Confluence Backups](#)
- [Enabling Backup Path Configuration](#)
- [Notes](#)

Related pages:

- [Confluence Administrator's Guide](#)

⚠ The information on this page *does not apply* to Confluence OnDemand.

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.

✓ 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-
Backup File Date Pattern	yyyy_MM_dd
Backup Path	
Custom backup paths are not enabled. More about custom backup paths	
☒ Backup Attachments	
Submit Cancel	

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 Directory](#).
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](#). Feel free to submit new scripts or post updates by logging in and adding them to the page as a comment.

 *The information on this page **does not apply to** Confluence OnDemand.*

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 **Administration Console**.

You need to have [System Administrator](#) permissions in order to perform this function.

Note:

- Atlassian recommends that you follow the [Production 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.)

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** folder.

kups.

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.

Related pages:

- [Restoring a Site](#)
- [Configuring Backups](#)
- [Production Backup Strategy](#)
- [Confluence Administrator's Guide](#)

 *The information on this page **does not apply** to Confluence OnDemand.*

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 Directory](#).
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 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.

Related pages:

- [Production Backup Strategy](#)
- [Manually Backing Up the Site](#)
- [Confluence Administrator's Guide](#)

 The information on this page *does not apply to* Confluence OnDemand.

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.
- **Backward version compatibility.** Confluence supports backward compatibility for site backups (but **not** for [space backups](#)). You can successfully restore backups of a site from an older version of Confluence to a newer version of Confluence. You cannot restore backups from a newer version to an older version or across major versions. For example, if you create a site backup in Confluence 2.4.3, it cannot be restored into a Confluence 2.2.2 site. It can however, be restored into 2.4.5 or 2.5.x, because 2.4.5 and 2.5.x are newer versions of Confluence. Importing an old XML backup file to a new major version (for example, Confluence 3.5 to Confluence 4.0) is not supported.
- **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

This page tells you how to import the contents of a Confluence space into another Confluence site, via an XML backup file.

You can [export the content of a space](#), including pages, comments and attachments. The process involves converting the data in the space into XML format. The end product is a zip file that contains XML file(s) and optionally, all the attachments in the space. To transfer this data to another Confluence site, restore this zip file as described below.

You need to have [System Administrator](#) permissions in order to restore a space from an XML zip file.

Before you start

Before you start, please check the following important notes.

Note 1: You cannot restore to a previous version

Let's assume you are importing a space from site A to site B. In this case, site B must be running the same version of Confluence as site A, or a later version than site A. The space import is not backwards compatible.

Note 2: You cannot restore to a different major Confluence version

Confluence only supports compatibility for space import and export when executed within the **same major version** of Confluence. (This issue is logged as [CONF-26111](#).)

Clarifying our terminology: By **major version**, we mean the version defined in the first two sections of the release number. For example, Confluence 2.2 and Confluence 2.3 are different major versions. Confluence 2.2.1 and Confluence 2.2.6 are the same major version.

Restoration data must share the same major version number. This means that a space export created in one major version of Confluence cannot be imported into a different major version of Confluence. For example, if you create a space export in Confluence 2.3.5, it cannot be imported into a Confluence 2.4.1 site. It can be however imported into 2.3.7. Similarly, a space export created in 5.0 can not be imported into 5.1.2. However, it can be restored into a Confluence 5.0.2 site.

If you need to import a space from a different major version, see the workaround described [below](#).

On this page:

- [Before you start](#)
 - [Note 1: You cannot restore to a previous version](#)
 - [Note 2: You cannot restore to a different major Confluence version](#)
 - [Note 3: Make sure the space does not already exist](#)
- [Restoring a space from an XML backup](#)
- [Workaround for restoring spaces between major releases](#)

Related pages:

- [Restoring a Site](#)
- [Confluence Administrator's Guide](#)

 *The information on this page does not apply to Confluence OnDemand.*

If you try to restore a space from a different major version of Confluence, you will see an error message similar to the one below and the import action will stop.

Screenshot: Clash of major versions on space restore

The following error(s) occurred:

- ◆ Restore denied. You can only restore space backups exported from the same major version (e.g. 2.2.x or 2.3.x).

Note 3: Make sure the space does not already exist

Confluence will only allow you to restore a space if there is not already a space with that key on the site. If you already have a space with the same key, you will need to delete or rename the existing space before restoring 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 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**.

Workaround for restoring spaces between major releases

If you need to import a space from an earlier major version, 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 same version as the site where you want to import the space. See [Upgrading Confluence](#).
5. Export the space from your temporary Confluence site. It will now have the correct 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 OnDemand.

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).

 The information on this page *does not apply* to Confluence OnDemand.

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.

RELATED TOPICS

[Important Directories and Files](#)
[Migrating to a Different Database](#)

 [Administrators Guide Home](#)  [Confluence Documentation Home](#)

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 OnDemand.

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.

Seeing an error when creating or importing a backup?

Problem	Solution
Exception while creating backup	Follow instructions below
Exception while importing backup	Follow Troubleshooting XML backups that fail on restore instead

 The information on this page *does not apply to* Confluence OnDemand.

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.hibernate.
ObjectNotFoundException:
08 No row with the given identifier exists: 2535, of class:
com.atlassian.confluence.core.ContentEntityObject
09 at
net.sf.hibernate.ObjectNotFoundException.throwIfNull(ObjectNotFound
Exception.java:24)
10 at
net.sf.hibernate.impl.SessionImpl.immediateLoad(SessionImpl.java:19
46)
11 at
net.sf.hibernate.proxy.LazyInitializer.initialize(LazyInitializer.j
ava:53)
12 at
net.sf.hibernate.proxy.LazyInitializer.initializeWrapExceptions(Laz
yInitializer.java:60)
13 at
net.sf.hibernate.proxy.LazyInitializer.getImplementation(LazyInitia
lizer.java:164)
14 at
net.sf.hibernate.proxy.CGLIBLazyInitializer.intercept(CGLIBLazyInit
ializer.java:108)
15 at
com.atlassian.confluence.core.ContentEntityObject$$EnhancerByCGLIB$
$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.toGenericX
ML(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.propertyset.B
ucketPropertysetItem@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.

RELATED TOPICS

[Enabling detailed SQL logging](#)

[Administrators Guide Home](#) [Confluence Documentation Home](#)

[Migrating from HSQLDB to MySQL](#)



If you've gone through [Migrating to Another Database](#) and cannot migrate because of a [failed xml backup](#), this page might help.

Disclaimer

MySQL Migration Toolkit is released by the makers of MySQL and as such, problems with the software should be directed to them. Atlassian Support does not offer support for the Migration Toolkit, nor do we provide support for this migration path. These instructions are offered for strictly informational purposes, and your mileage may vary.

**Backup Reminder**

Please backup your database and your home folder before attempting this.



*The information on this page **does not** apply to Confluence OnDemand.*

Resources needed


- Empty MySQL DB with appropriate credentials to allow creation, deletion, and insertion of tables and rows.
- A Windows machine that can both communicate to the Confluence server and the destination DB.
- [MySQL Migration Toolkit](#)
- [HSQL Database Engine](#)

Preparation for migrating to MySQL from HSQLDB

1. Shutdown Confluence
2. Make a copy of the confluence home folder for backup purposes
3. Install the Migration Toolkit
4. Unzip the hsqldb package.
5. Copy the hsqldb.jar from hsqldb/lib into C:\Program Files\MySQL\MySQL Tools for 5.0\java\lib
6. Start the MySQL Migration Toolkit

Running the Migration Toolkit

You should be presented with the following screen.

**Choose Direct Migration**

Metadata

Configuration Type

Choose the type of configuration you have set up.

Configuration type

Before you can start the migration process you have to specify your system configuration.

☒ **Direct Migration**
MySQL Migration Tool is installed on source or target machine

Use this configuration if you have installed the MySQL Migration Service on either the source or target machine.

Please note that if the MySQL Migration Tool is not located on either the source or target machine there will be a huge overhead of network traffic and a major performance loss.

In that case please use the Three Way Configuration by installing the MySQL Migration Agent on the source or target machine.

☐ **Agent Based Migration**
Use MySQL Migration Agent installed on source machine

To setup the Three Way Configuration you have to install the MySQL Migration Agent on either the source machine or the target machine.

Source Database

Source Database

Select the source database you want to migrate from.

Source Database Connection

Database System: Select a RDBMS from the list of supported systems

Driver: Choose from the list of available drivers for this RDBMS

Connection Parameters

Source Connection Parameter
Please enter the connection parameters to connect to the database.

Stored Connection:

Class Name: Classname of the driver to use.

Connection String: Jdbc Connection String

Username: Explicit username if not submitted in the connection string

Password: Explicit password

Database System:	Generic JDBC
Connection String:	jdbc:hsqldb: file:PATH TODATABASEFOLDER\conflencedb \\
Username:	sa
Password:	No password. Leave this field blank

Destination Database

❗ Please make sure that the computer that is running MySQL Toolkit is able to access the MySQL server and that the user listed has the ability to create, drop, insert, and update tables.

❗ If your MySQL user has a \$ character in the password (such as 'pa\$sword'), please change the password or create a temporary account with full permissions. If you do not, the toolkit will throw an "Illegal group reference" error and you will not be able to proceed with the migration.

Target Database

Select the destination database.

Target Database Connection

Database System:

Select a RDBMS from the list of supported systems

Driver:

Choose from the list of available drivers for this RDBMS

Connection Parameters



Target Connection Parameter
Please enter the connection parameters to connect to the database.

Stored Connection:

Hostname:

Port:

Name or IP address of the server machine - TCP/IP port

Username:

Name of the user to connect with.

Password:

The user's password.

Advanced Settings

Connection String:

Jdbc Connection String

Connecting to Servers

Connecting to Servers

Establishing database connections.

Connection Progress**Tasks to execute**

The following tasks will now be executed. Please monitor the execution progress. Press [Advanced >>] to see the log.

- ☒ Connecting to source database system
- ☒ Retrieve schema information from source database system
- ☒ Test connection to target database system

Execution completed successfully.

You should see the toolkit trying to connect. If you have problems, please click on the advanced options and sql will show you debugging information. Click Advanced to see the log. If you see "Java Heap Space: Out of Memory", you can start the *MySQL Migration Toolkit* with a `-Xmx` flag to allocate more memory to the JVM.

After this screen you should come to reverse engineering. Click next.

Source Schemata Selection**Source Schemata Selection**

Choose the schemata you want to migrate.

Schemata

Select all schemata that have to be migrated.

Show as List ☐

INFORMATIO...



PUBLIC

1 schema selected.

You should see 2 databases, **INFORMATION_SCHEMA** and **PUBLIC**. Choose **PUBLIC**

Object Type Selection

Migration

In this step the selected object will be migrated.

Migration of Meta Data**Tasks to execute**

The following tasks will now be executed. Please monitor the execution progress. Press [Advanced >>] to see the log.

- ☒ Execute Migration Process
- ☒ Generate SQL Create Statements

Execution completed successfully.

Click Next.

Object Type Mapping**Object Creation Options**

Please define how the object creation should be performed.

Object Creation Options**Database Object Creation Parameters**

Select the desired options for the object creation. Click Next > to start the creation process.

☒ Create Objects Online

Select this option to create the objects on the target database. If there is a problem during the creation process you will be informed and can fix the used statement by pressing the [Details >>] button.

☐ Create Script File for Create Statements

If you want to store the object creation in a script file enable this option. You can use this option in parallel to creating the objects online option if you want to have a backup of the SQL commands.

Filename: ...

Click **Show Details** on both sections. For **Migration Method for Type Schema**, choose **Multilanguage**. For **Migration Method for Type Table**, choose **Data Consistency/Multilanguage**

Click **Advanced**. Check **Enabled Detailed Mappings in Next Step**

Detailed Object Mapping

Click to rename the **destination database** to be the one set aside to migrate to.

From this point on, you should be able to click next all the way through to finish the migration.

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.

However, if migrating from HSQLDB to MySQL, you might have a better experience using the [MySQL Migration Toolkit](#).

Seeing an error when creating or importing a site or space backup?

Problem	Solution

Exception while creating backup	Follow Troubleshooting failed XML site backups instead
Exception while importing backup	Follow 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](#)
[Confluence Administrator's Guide](#)

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 at `lassian-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 disable 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.propertyset.B
ucketPropertysetItem@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

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

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](#)
- [Migrating between Attachment Storage Systems](#)

Related pages:

- [Working with Confluence Logs](#)
- [Working with Confluence LogsConfluence Administrator's Guide](#)

Attachment Storage Options

Local File System

By default, Confluence stores attachments in the `attachments` directory within the configured Confluence home folder. If you are looking to run [Confluence Clustered](#), attachments must be stored in the database.

Database

Confluence gives administrators the option to store attachments in the database that Confluence is configured to use.

Here are some reasons why, as an administrator, you may want to choose this storage system:

- Ease of backup.
- Avoiding issues with certain characters in attachment file names.



While storing attachments in the database can offer some advantages, 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 perform a migration, 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. Confluence administrators should keep in mind that the amount of disk space used by Confluence is directly proportional to the number and size of attachments put into the system.

To configure the maximum size allowed for an attachment:

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 maximum size next to **Attachment Maximum Size**.
The default is 10mb.
5. 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.

Example

To specify 250 kb you would use the following JVM parameter:
`-Datlassian.indexing.contentbody.maxsize=256000`

Related pages:

- [Recognised System Properties](#)
- [Working with Attachments](#)
- [Confluence Administrator's Guide](#)

Outcomes of Limiting Attachment Indexing Size

Limiting the size of attachment indexing has the following effects:

- Decreases the size of the index when large attachments are present.
- Decreases the memory used in indexing large attachments.
- Prevent excerpts of large attachments being displayed in search results.

For more details, please refer to the following issue in our issue tracker: [CONF-10512](#).

Hierarchical File System Attachment Storage

For Confluence version 3.0, the structure of attachments stored on the filesystem was changed. In versions of Confluence prior to 3.0, attachments were stored in directories corresponding to the id of the content to which they belong. The more content in Confluence with attachments, the more directories you would have immediately beneath your configured attachments directory. This directory structure has been changed in Confluence 3.0 and since the default configuration of Confluence is to [store attachments in the filesystem](#), this change is likely to have relevance to administrators of most existing Confluence installations.

If you are installing Confluence for the first time, there will be no consequences as a result of this change. If you are upgrading from a previous version of Confluence, the migration to this new filesystem structure should happen automatically during the upgrade.

The reason for introducing this change was to address the issue [CONF-13004](#). Certain file systems have a limit on the number of files that can be stored in a directory and large Confluence installations were reaching this limit. In addition, storing too many files at a single directory level can cause performance degradation in some circumstances. This new attachment storage strategy ensures this will no longer be the case.

 The information on this page *does not apply* to Confluence OnDemand.



Backup Confluence Home

Before upgrading to Confluence 3.0, as with any upgrade you must ensure you have a backup of your Confluence home directory before you proceed.

The New Directory Layout

The attachment storage layout was chosen to fulfil the following main requirements:



1. Limit the number of entries at any single level in a directory structure.
2. Partition attachments per space making it possible for a system admin to selectively back up attachments from particular spaces (see the [JIRA issue](#) for more details).

An attachment in Confluence can be thought of as having a number of identifying attributes: *id*, *space id* and *content id*. That is to say, the attachment logically belongs to a piece of content which logically belongs in a space (not all content belongs to a space). For attachments 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> , modulo 250
6	The next 3 least significant digits of the <i>content id</i> , modulo 250
7	The full <i>content id</i>
8	The full <i>attachment id</i>

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:

Attachments:**Directory Structure:**

To find the directory where attachments for a particular space are stored, you can use the JSP `findspaceattachments.jsp` at the location `<confluence url>/admin/findspaceattachments.jsp`. This JSP requires a space key and returns the directory on the file system where attachments for that space are stored.

Attachment D in the above diagram is stored in a slightly different structure. Attachments that are not conceptually within a space replace the level 2 - 4 directories with a single directory called 'nonspaced'. Examples of such attachments are the global site logo and also attachments on draft content.

Upgrading to the new attachment storage structure

As mentioned previously, this upgrade is only necessary if you have Confluence configured to store attachments on the [file system](#).

If migration is not necessary due to a different storage configuration (for example, because attachments are stored in the database), then no migration will occur during upgrade and the [Confluence log](#) will simply show the following messages -

```
INFO [main] [AbstractUpgradeManager] upgradeStarted Starting automatic
upgrade of Confluence
INFO [main] [UpgradeTask] isUpgradeNeeded The configured
attachmentDataDao does not store
    attachment data on the file system so the
HierarchicalFileSystemAttachmentUpgradeTask is not necessary.
INFO [main] [AbstractUpgradeManager] upgradeFinished Upgrade completed
successfully
```

Should migration be required, it will occur automatically during upgrade and the [log](#) will show output similar to this -

```
INFO [main] [UpgradeTask] doUpgrade Beginning
HierarchicalFileSystemAttachmentUpgradeTask. Depending on the size of
the
    attachment data this may take some time.
INFO [main] [UpgradeTask] run 4023 pages may have attachments to be
moved to a new hierarchical structure.
INFO [main] [UpgradeTask] run 0 of 4023 pages have had their attachments
moved to the new structure
INFO [main] [UpgradeTask] run 500 of 4023 pages have had their
attachments moved to the new structure
INFO [main] [UpgradeTask] run 1000 of 4023 pages have had their
attachments moved to the new structure
INFO [main] [UpgradeTask] run 1500 of 4023 pages have had their
attachments moved to the new structure
INFO [main] [UpgradeTask] run 2000 of 4023 pages have had their
attachments moved to the new structure
INFO [main] [UpgradeTask] run 2500 of 4023 pages have had their
attachments moved to the new structure
INFO [main] [UpgradeTask] run 3000 of 4023 pages have had their
attachments moved to the new structure
INFO [main] [UpgradeTask] run 3500 of 4023 pages have had their
attachments moved to the new structure
INFO [main] [UpgradeTask] run 4000 of 4023 pages have had their
attachments moved to the new structure
INFO [main] [UpgradeTask] run Successfully moved the attachments for all
4023 pages to the new hierarchical structure.
INFO [main] [UpgradeTask] doUpgrade Completed
HierarchicalFileSystemAttachmentUpgradeTask.
INFO [main] [AbstractUpgradeManager] upgradeFinished Upgrade completed
successfully
```



It should be noted that for most implementations of Java, the migration to the new data structure involves moving the files (not copying them). Hence, there should not be a need to have additional disk space available. It also means that the migration should be relatively fast.

Have you previously applied the CONF-8298 patch?

The patch or workaround on the [CONF-8298](#) issue changed the structure of attachment storage but not to the most efficient possible structure. So during the Confluence 3.0 upgrade process this intermediate ([CONF-8298](#)) structure will be detected and automatically upgraded.

Troubleshooting the upgrade

 It should be noted that in the event of a failure, your attachment directory may be in an inconsistent state and your first step in troubleshooting should be to restore the backup of your home directory.

There are a number of reasons the migration could fail. This will be shown in the log with a message similar to "Failed to move the attachments for all pages to the new hierarchical structure."

Immediately preceding this message in the log will be entries for each page whose attachments could not be moved. The following table shows examples of these messages and offers some possible explanations.

Example Message	Description
The configured attachment directory <directory name> could not be found or was not a directory.	The configured Confluence attachment directory is not accessible. Check confluence home for the attachment directory and ensure the permissions are correct to allow reading and writing for this directory.
It is not possible to migrate the attachments to the new structure since files already exist which the attachment process may need to create.	Your attachments directory contains files or directories which the upgrade task wants to create. That is, a top level directory called ver003 containing directories or files with names containing up to 3 digits (e.g. 1, 213). This could be due to a previous failed attempt to migrate the attachments. You should restore a previous good copy of your attachments directory and remove any files or directories with this naming pattern before retrying.
Couldn't find current Confluence content for the id <content Id>. The attachment is a non-spaced attachment (e.g. global logo, draft attachment, etc) and will be migrated to the nonspaced directory.	This is a normal message indicating that the attachment being migrated does not belong to a space e.g. global logo, global description, personal information (on profile pages) and attachments on draft content.
Problem while accessing the database for content id content Id so its attachments will not be migrated.	It was not possible to access the database at this point during the migration. You will need restore your Confluence attachment directory from the backup and attempt the upgrade again, once the database is accessible again.
Could not create the new attachment directory directory.	The upgrade task could not create the new directory to contain the attachment being moved. Does the server user have sufficient permission to perform this operation in the indicated directory? Is there sufficient disk space?
Failed to move the current attachment directory <some path> to the new location of <some other path>.	The upgrade task could not move the directory. Does the server user have sufficient permission to perform this operation in the indicated directory?

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.

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.

On this page:

- [Database diagrams](#)
- [Database tables and references](#)
- [Authentication](#)
- [Content](#)
- [Clustering](#)
- [System information](#)
- [Spaces](#)
- [Appearance](#)
- [Miscellaneous](#)

Related pages:

- [Managing Confluence Data](#)
- [Connecting to HSQLDB using DBVisualizer](#)
- [Confluence Administrator's Guide](#)

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 key name	Primary key name
AO_9412A1_A OUSER	ID	AO_9412A1_U SER_APP_LIN K	USER_ID	fk_ao_9412a1_ user_app_link_ user_id	AO_9412A1_A OUSER_pkey
attachments	attachmentid	attachmentdata	attachmentid	fk9dc3e34d34a 4917e	attachments_p key
attachments	attachmentid	attachments	prevver	fk54475f9017d 4a070	attachments_p key
attachments	attachmentid	content_label	attachmentid	fkf0e7436e34a 4917e	attachments_p key
attachments	attachmentid	imagedetails	attachmentid	fka768048734a 4917e	attachments_p key

content	contentid	attachments	pageid	fk54475f908c38fbea	content_pkey
content	contentid	bodycontent	contentid	fka898d4778dd41734	content_pkey
content	contentid	confancestors	ancestorid	fk9494e23c37e35a2e	content_pkey
content	contentid	confancestors	descendentid	fk9494e23cc45e94dc	content_pkey
content	contentid	content	prevver	fk6382c05917d4a070	content_pkey
content	contentid	content	parentid	fk6382c05974b18345	content_pkey
content	contentid	content	parentcommentid	fk6382c059b97e9230	content_pkey
content	contentid	content	pageid	fk6382c0598c38fbea	content_pkey
content	contentid	content_label	contentid	fkf0e7436e8dd41734	content_pkey
content	contentid	content_perm_set	content_id	fkbf45a7992caf22c1	content_pkey
content	contentid	contentproperties	contentid	fk984c5e4c8dd41734	content_pkey
content	contentid	extrnlnks	contentid	fk97c10fe78dd41734	content_pkey
content	contentid	likes	contentid	fk4514b9c8dd41734	content_pkey
content	contentid	links	contentid	fk45157998dd41734	content_pkey
content	contentid	notifications	pageid	fk594acc88c38fbea	content_pkey
content	contentid	spaces	homepage	fk9228242d11b7bfee	content_pkey
content	contentid	spaces	spacedescid	fk9228242d2c72d3d2	content_pkey
content	contentid	trackbacklinks	contentid	fkf6977a478dd41734	content_pkey
content_perm_set	id	content_perm	cps_id	fkbd74b31676e33274	content_perm_set_pkey
cwd_app_dir_mapping	id	cwd_app_dir_group_mapping	app_dir_mapping_id	fk_app_dir_group_mapping	cwd_app_dir_mapping_pkey
cwd_app_dir_mapping	id	cwd_app_dir_operation	app_dir_mapping_id	fk_app_dir_mapping	cwd_app_dir_mapping_pkey
cwd_application	id	cwd_app_dir_group_mapping	application_id	fk_app_dir_group_app	cwd_application_pkey
cwd_application	id	cwd_app_dir_mapping	application_id	fk52050e2fb347aa6a	cwd_application_pkey

cwd_application	id	cwd_application_address	application_id	fk_application_address	cwd_application_pkey
cwd_application	id	cwd_application_attribute	application_id	fk_application_attribute	cwd_application_pkey
cwd_directory	id	cwd_app_dir_group_mapping	directory_id	fk_app_dir_group_dir	cwd_directory_pkey
cwd_directory	id	cwd_app_dir_mapping	directory_id	fk_app_dir_dir	cwd_directory_pkey
cwd_directory	id	cwd_directory_attribute	directory_id	fk_directory_attribute	cwd_directory_pkey
cwd_directory	id	cwd_directory_operation	directory_id	fk_directory_operation	cwd_directory_pkey
cwd_directory	id	cwd_group	directory_id	fk_directory_id	cwd_directory_pkey
cwd_directory	id	cwd_group_attribute	directory_id	fk_group_attribute_dir_id	cwd_directory_pkey
cwd_directory	id	cwd_user	directory_id	fk_user_dir_id	cwd_directory_pkey
cwd_directory	id	cwd_user_attribute	directory_id	fk_user_attribute_dir_id	cwd_directory_pkey
cwd_group	id	cwd_group_attribute	group_id	fk_group_attribute_id_group_id	cwd_group_pkey
cwd_group	id	cwd_membership	parent_id	fk_parent_grp	cwd_group_pkey
cwd_group	id	cwd_membership	child_group_id	fk_child_grp	cwd_group_pkey
cwd_user	id	cwd_membership	child_user_id	fk_child_user	cwd_user_pkey
cwd_user	id	cwd_user_attribute	user_id	fk_user_attribute_id_user_id	cwd_user_pkey
cwd_user	id	cwd_user_credential_record	user_id	fk76f874f73aee0f	cwd_user_pkey
external_entities	id	external_members	extentityid	fk d8c8d8a5f25e5d5f	external_entities_pkey
groups	id	external_members	groupid	fk d8c8d8a5117d5fda	groups_pkey
groups	id	local_members	groupid	fk6b8fb445117d5fda	groups_pkey
keystore	keyid	trustedapp	public_key_id	fkddb119ca9c85adb1	keystore_pkey
label	labelid	content_label	labelid	fkf0e7436e27072aef	label_pkey
os_group	id	os_user_group	group_id	fk932472461e2e76db	os_group_pkey
os_user	id	os_user_group	user_id	fk93247246f73aee0f	os_user_pkey

pagetemplates	templateid	content_label	pagetemplateid	fkf0e7436ed32042e4	pagetemplates_pkey
pagetemplates	templateid	pagetemplates	prevver	fkbc7ce96a17d4a070	pagetemplates_pkey
spacegroups	spacegroupid	spacegrouppermissions	spacegroupid	fk487d958b16994414	spacegroups_pkey
spacegroups	spacegroupid	spaces	spacegroupid	fk9228242d16994414	spacegroups_pkey
spaces	spaceid	content	spaceid	fk6382c059b2dc6081	spaces_pkey
spaces	spaceid	notifications	spaceid	fk594acc8b2dc6081	spaces_pkey
spaces	spaceid	pagetemplates	spaceid	fkbc7ce96ab2dc6081	spaces_pkey
spaces	spaceid	spacepermissions	spaceid	fk d33f23beb2dc6081	spaces_pkey
trustedapp	trustedappid	trustedapprestriction	trustedappid	fke8496ba235d1d865	trustedapp_pkey
users	id	local_members	userid	fk6b8fb445ce2b3226	users_pkey

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 content table.
content	A persistence table for the ContentEntityObject class of objects. The subclass is indicated by the contenttype column.
content_label	Arbitrary text labels for content.
label	The other half of the content_label 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 safetynumber 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.
extrnltns	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. The simple way is to go directly to the database. We recommend [DbVisualizer](#), and have basic instructions for [connecting it to HSQLDB](#).

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 spacenames, and the last date and time at which any content was added or changed.

 The information on this page *does not apply* to Confluence OnDemand.

Alternatively, this one simply identifies spaces whose 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.

It's also possible to present the information in a wiki page, using the [SQL plugin](#), which can be installed using the [Plugin Exchange](#). You'll also need to define a database resource in `conf/server.xml` and `confluence/WEB-INF/web.xml`, as described [here](#). Having done so, you can use wiki markup code like the following, replacing `confluenceDS` with the name of your own local datasource:

```
h3. Space activity
{sql:datasource=confluenceDS|output=wiki}
SELECT spaces.spacename AS Space, MAX(content.lastmoddate) AS LastModified
FROM content, spaces
WHERE content.spaceid = spaces.spaceid
GROUP BY Space;
{sql}
```

The result will be something like this:

Space activity:

space	lastmodified
Private Space	2007-10-11 11:34:04.914
Another space	2007-10-11 11:39:39.716
More space	2007-10-11 11:40:11.688

You can try the [Chart plugin](#) in combination with the SQL plugin to give more visually attractive results.

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 [Importing Content Into Confluence](#).

You can also export Confluence content to various formats. See [Exporting Confluence Pages and Spaces to Other Formats](#).

Related pages:

- [Managing Confluence Data](#)
- [Confluence Administrator's Guide](#)

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.

For a list of important files on the server side, see [Important Directories and Files](#).

This section contains the following guidelines:


- [Important Directories and Files](#)
- [Application Server Configuration](#)
- [Web Server Configuration](#)
- [Starting Confluence Automatically on System Startup](#)

Related pages:

- [Getting Started as Confluence Administrator](#)
- [Supported Platforms](#)
- [Confluence Administrator's Guide](#)

 *The information on this page does not apply to Confluence OnDemand.*

Diagram: A Confluence installation



Important Directories and Files

The Installation Directory

The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.

Important Files and Directories

- `confluence/WEB-INF/classes/confluence-init.properties` : This file tells Confluence where to find the [Confluence Home Directory](#). This file is modified by the administrator when [installing Confluence](#).
- `confluence/WEB-INF/lib/` : This directory is used when deploying plugins, especially those plugins that cannot automatically be loaded through the Administration Console.
- `confluence/WEB-INF/classes/log4j.properties` : Confluence's logging configuration file. See [Working with Confluence Logs](#).
- `conf/server.xml` : SSL configuration. See [Running Confluence Over SSL or HTTPS](#).

Memory Settings

The file used to edit JAVA_OPTS memory settings will depend on the method used to install Confluence, as well as the operating system used for your installation.

- Windows Users
 - **Confluence** — bin/setenv.bat
 - **Confluence Installer** — wrapperwin32.conf
- Mac/Linux Users
 - **Confluence** — bin/setenv.sh
 - **Confluence Installer** — wrapperosx.conf

The Temp Directory

The temp directory is configured in the Java runtime and some Confluence components write temporary files or lockfiles into this directory.

For EAR/WAR installations typically, this directory is /tmp on Linux systems, or C:\Temp on Windows.

For Standalone installations 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.
```

The Confluence Home Directory

The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you are using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.

✓ Tip: Another term for 'Home directory' would be 'data directory'.

Administrators can expect the Confluence Home Directory to grow quite large in a busy site.

The location of this directory is configured by the system administrator during installation (see `confluence-init.properties` above). 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.

Important Files and Directories

- `confluence.cfg.xml` : Confluence's core configuration file; includes the configuration for connecting to its database.
- `attachments/` : All file attachments in the Confluence site are stored under this directory. This is the only place Confluence keeps attachment files.
- `backups/` : If Confluence is configured to produce daily backups, these are kept in this directory. Administrators should occasionally delete old or unwanted backups from this directory to prevent it from growing too large.
- `config/` : Miscellaneous global and per-space configuration files are kept in this directory.
- `database/` : If Confluence is being run from the [embedded HSQL database](#), the database files will be kept in this directory.
- `index/` : The full-text search index is kept in this directory. Removing or modifying files in this directory may cause search to no longer function. Rebuilding the search index from Confluence's global administration screen will completely regenerate the contents of this directory.
- `temp/` : Confluence stores temporary files in this directory, especially during backups and exports. A daily job within Confluence deletes files that are no longer needed.
- `thumbnails/` : Stores temporary files for image thumbnails. The contents of this directory can be safely deleted, as Confluence will regenerate thumbnails as required.

Database

All other data, including page contents, links, archived mail and so on, is kept in the database. If you have configured Confluence to use the [embedded HSQL database](#), 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.

**Tip**

All of Confluence's persistent data is stored either in the Confluence Home Directory, or the database. If you have backup copies of both of these, taken at the same time, you will be able to restore Confluence from them (see [Restoring Data from other Backups](#)).

RELATED TOPICS

[Confluence Home Directory](#)
[Confluence Installation Directory](#)
[Embedded HSQLDB Database](#)
[Database Configuration](#)

Confluence Data Directory Configuration

Here is a link listing [important Confluence files](#).

The home directory defines the location of the directory where Confluence will store its data, including attachments, indexes and backups. Administrators can set this location by defining a value for the file `<MY-INST ALL>/confluence/WEB-INF/classes/confluence-init.properties`. To find what your home directory is currently set to, open this file and check the `confluence.home` property. It is unset on new installations.



*The information on this page **does not apply to** Confluence OnDemand.*

Windows Configuration

On Windows, this path:

```
C:\confluence\data
```

will be written like so:

```
confluence.home=C:/confluence/data
```

Note that all backslashes (`\`) are written as forward slashes (`/`).

Linux/Solaris Configuration

On any Linux-based system, the property is defined using the normal directory syntax:

```
confluence.home=/var/confluence/
```

Symbolic links

If your `confluence.home` directory contains a symbolic link, you must define the absolute path.



Please note that there can be no symbolic links within the `confluence.home` directory. 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.

Confluence Home Directory

Often in the documentation, you'll see a reference to the 'Confluence Home directory'.

What is the Confluence Home Directory?

The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you are using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.

✓ Tip: Another term for 'Home directory' would be 'data directory'.

⚠ The information on this page *does not apply to* Confluence OnDemand.

Finding the Confluence Home Directory

The location of the Confluence Home directory is defined when you [install Confluence](#). This location 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.

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.

When Confluence first starts up, it reads the `confluence-init.properties` file to determine where to look for the Home directory. The property that determines this is `confluence.home`, for example:

```
confluence.home=/var/atlassian/application-data/confluence
```

Once Confluence is running you can find the Confluence Home directory via the Administration console, under Administration > System Information > Confluence Information - Confluence Home.

Content of the Confluence Home Directory

The Confluence home directory contains some of the configuration data used by Confluence. Other data is stored in the database. This section outlines the purpose of the files and directories in the Confluence home directory.

`confluence.cfg.xml`

This file contains all of the information necessary for Confluence to start up, such as:

- Product license
- Context path
- Database details, such as location and connection pool settings
- Paths to important directories

`attachments`

This directory contains every version of each attachment stored in Confluence. This directory is not used when Confluence is configured to store attachments in the database. Attachments are always stored in the database in clustered instances of Confluence.

Since Confluence 3.0, the directory structure has been defined by the [Hierarchical File System Attachment](#)

Storage method.

For versions before Confluence 3.0, paths within this directory had the following structure:

```
/attachments/PAGE_ID/ATTACHMENT_ID/VERSION
```

You can specify an alternative directory for attachment storage by setting the `attachments.dir` property in `confluence.cfg.xml`.

backups

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:

```
daily-backup-YYYY_MM_DD.zip
```

You can specify an alternative directory for backups by setting the `daily.backup.dir` property in `confluence.cfg.xml`.

bundled-plugins

Confluence ships with a set of *bundled* plugins. These are plugins written by the Atlassian and the Confluence community that we think provide useful and broadly applicable functionality in Confluence. The `bundled-plugins` 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. It will simply be replaced the next time Confluence starts up.

database

This is where Confluence stores its database when configured to run with the HSQL embedded database. 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.

index

This is where Confluence stores its indexes for rapid retrieval of often used data. The Confluence index is used heavily by the application for content searching and recently updated lists and as such is critical for a running Confluence instance. It is important to note however that should the data in this directory be lost or corrupted, it can be restored by running a full reindex from within Confluence. This 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 `lucene.index.dir` property in `confluence.cfg.xml`. As this is the most heavily accessed directory in the Confluence home directory you might want to consider hosting it on the fastest disk available. It would also be useful if the disk holding the Confluence index was not heavily used by any other application to reduce access contention.

plugin-cache

All Confluence plugins are stored in the Confluence database. To allow for quicker access to classes contained within the plugin JARs, Confluence will cache these plugins in the `plugin-cache` 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.

resources

The `resources` directory stores any space logos used in your Confluence instance. For each space with a space logo, there is a directory within `resources` named after the space's key. That directory contains the space's logo.

temp

The temp directory is used for various runtime functions such as exporting, importing, file upload and indexing. As the name suggests, and file in this directory is of temporary importance and is only used during runtime. This directory can be safely emptied when Confluence is offline.

An alternative directory may be specified for temporary data by setting the `webwork.multipart.saveDir` property in `confluence.cfg.xml`.

thumbnails

When Confluence generates a thumbnail of an image (for example when the `gallery` macro is used), the resulting thumbnail is stored in this directory for quicker retrieval on subsequent accesses. This directory is essentially a thumbnail cache, and deleting files from this directory simply means the thumbnail will have to be regenerated on the next access.

RELATED TOPICS

[Confluence Installation Directory](#)
[Important Directories and Files](#)
[Embedded HSQLDB Database](#)

Confluence Installation Directory

The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.

 The information on this page *does not apply to* Confluence OnDemand.

RELATED TOPICS

[Confluence Home Directory](#)
[Important Directories and Files](#)

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

```
JkOptions +ForwardURICompatUnparsed
```

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.

Please note, you cannot simply copy the WAR file or expanded WAR directory from an old Confluence EAR/WAR version in the old application server to Tomcat. **This will not work.**

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 or Confluence EAR-WAR.
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 Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you are using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.
 -  **Tip:** Another term for 'Home directory' would be 'data 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** (if you are using Confluence) **or your Confluence webapp** (if you are using Confluence EAR-WAR edition). The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.

3. Switching Application Servers

1. Install Confluence on your new application server. We recommend that you [install Confluence](#) (from the zip file) as it is preconfigured with Tomcat. If you want more control over the installation process, you can [install Confluence EAR-WAR on Tomcat](#) however this requires more manual configuration. Regardless of which method you choose, as part of the installation process:
 - If you are connecting to your database via a standard JDBC connection, enter the URL, username and password for your existing database.
 - If you are connecting to your database via datasource, use the settings for your existing database when you configure the JDBC datasource in your new server. Refer to the appropriate guide below:
 - [Configuring a MySQL Datasource in Apache Tomcat](#)
 - [Configuring a SQL Server Datasource in Apache Tomcat](#)
 - [Configuring a PostgreSQL Datasource in Apache Tomcat](#)
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. Make sure you shutdown the old server before you startup 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-Directory>\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-Directory>\bin\setenv.sh or the <Installation-Directory>\bin\setenv.bat file, you may want to make the modifications in your new installation. The parameters are specified in the JAVA_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 "defineCGLIBCClassInJavaPackage";  
};
```

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 mod_proxy and mod_jk
mod_proxy (also known as reverse proxy)	• recommended connection method • simple HTTP proxy to application server • works with all application servers • if application paths are consistent, there is minimal load on the web server	• application paths must be consistent to avoid complex and slow URL rewriting • works with name-based virtual hosting, both on web server and app server • web server keeps a pool of connections to application server
mod_jk (also known as AJP)	• uses the AJP binary protocol • provides failover (and load balancing, which Confluence supports only with a clustered license) • only works with some application servers (typically Tomcat) • if application paths are consistent, there is some load on the web server to translate requests to AJP	

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; the [JIRA docs](#) describe the process for Ubuntu/Debian layout):

```
# 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>
```

▼ 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:

Externally accessible (Apache) URL	http://confluence.example.com/
Application server URL	http://app-server.internal.example.com:8090/confluence/

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.

```
# 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>
```

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 JIRA 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 and Confluence on the same application server. There is a [sample configuration for Tomcat](#) you can use after configuring Apache.

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 `/`.

JIRA external URL	http://jira.example.com/
JIRA application server URL	http://jira-app-server.internal.example.com:8080/
Confluence external URL	http://confluence.example.com/
Confluence application server URL	http://confluence-app-server.internal.example.com:8090/

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:

Externally accessible (web server) URL	http://www.example.com/confluence/
Application server URL (HTTP)	http://app-server.internal.example.com:8090/confluence/

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.

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-3.0.1-std.tar.gz
ln -s confluence-3.0.1-std/ 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-6-sun

case "$1" in
    # Start command
    start)
        echo "Starting $APP"
        /bin/su -m $USER -c "$CATALINA_HOME/bin/startup.sh &> /dev/null"
        ;;
    # Stop command
    stop)
        echo "Stopping $APP"
        /bin/su -m $USER -c "$CATALINA_HOME/bin/shutdown.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-1.5.0-sun
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-6-sun-1.6.0.16/
export JDK_HOME=/usr/lib/jvm/java-6-sun-1.6.0.16/
cd /usr/local/confluence/current/bin
./startup.sh
```

Example stopscript:

```
#!/bin/bash
export JAVA_HOME=/usr/lib/jvm/java-6-sun-1.6.0.16/
export JDK_HOME=/usr/lib/jvm/java-6-sun-1.6.0.16/
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:

```
tomcat6 //US//Confluence --Startup auto
```

6. If you have a less than a 512 megabytes of memory, skip this step. For users with large Confluence installations, you can increase the maximum memory Confluence can use. (The default is 256MB). For example, you can set the maximum memory to 512 megs using:

```
tomcat6 //US//Confluence --JvMx 512
```

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:

```
tomcat6 //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'](#)

-  [Unable to Install Service on Windows Vista](#)
-  [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](#)

- 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](#).

- [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](#)

Related pages:

- [Customising your Confluence Site](#)
- [Confluence Administrator's Guide](#)

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](#)
- [Viewing Site Statistics](#)
- [Viewing and Editing License Details](#)
- [About Add-ons](#)
- [Live Monitoring Using the JMX Interface](#)
- [Tracking Customisations Made to your Confluence Installation](#)
- [Confluence Administrator's Guide](#)

 *The information on this page does not apply to Confluence OnDemand.*

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.

 *The information on this page does not apply to Confluence OnDemand.*

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 [the Windows Service](#) - 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.

Related Topics

- [Viewing System Information](#)
- [Cache Statistics](#)
- [Viewing and Editing License Details](#)
- [Viewing and Managing Installed Plugins](#)

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.

 The information on this page *does not apply* to Confluence OnDemand.

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.vm, template/includes/macros.vm, decorators/mail.vmd, decorators/space.vmd, template/includes/personal-sidebar.vm
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](#).

RELATED TOPICS

 [Administrators Guide Home](#)  [Confluence Documentation Home](#)

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.

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](#)).

Related pages:

- [How Do I Get More Statistics from Confluence?](#)
- [Cache Statistics](#)
- [Viewing Space Activity](#)
- [Live Monitoring Using the JMX Interface](#)
- [Installing and Configuring Plugins](#)
- [Confluence Administrator's Guide](#)

 *The information on this page **does not apply** to Confluence OnDemand.*

Screenshot: Global Activity

Global Activity

Activity for 15 February 2013


Period: [month](#) | [week](#) | [day](#)

Day: [<< Previous](#) | [Now](#) | [Next >>](#)

Viewing



Editing



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**.

 The information on this page *does not apply to* Confluence OnDemand.

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.

 The information on this page *does not apply to* Confluence OnDemand.

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=
```

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`.

RELATED TOPICS

Content by label

There is no content with the specified labels



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:

- [Searching Confluence](#)
- [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.

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**.

Related pages:

- [Choosing a Default Language](#)
- [Installing a Language Pack](#)
- [Content Index Administration](#)
- [Creating a Lowercase Page Title Index](#)
- [Rebuild the Content Indices from Scratch](#)
- [Confluence Administrator's Guide](#)

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.

Related pages:

- [Searching Confluence](#)
- [Confluence Administrator's Guide](#)

 *The information on this page **does not apply** to Confluence OnDemand.*

To configure the quick navigation feature:

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**.

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 [Glify 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 minute (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.

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
- Searching Confluence
- Configuring the Confluence Search and Index
- Confluence Administrator's Guide

 *The information on this page does not apply to Confluence OnDemand.*

Screenshot: Index summary

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%

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.**

NOT BUILT

0%

Build

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.

It may help to increase the heap memory allocation of Confluence by following the instructions in the [JIRA documentation](#).

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:

- [Searching Confluence](#)
- [Confluence Administrator's Guide](#)

Rebuilding the Ancestor Table



Due to a known issue

 **CONF-32174** - Rebuild Ancestor Table is Inaccessible RESOLVED, the page to rebuild the ancestor table is not accessible in Confluence 5.4.1 and later. You will need to apply the workaround described in the bug report to be able to access this functionality.

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
```

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

Page Level Permissions

[Rebuild Ancestor Table](#)

Related pages:

- [Searching Confluence](#)
- [Confluence Administrator's Guide](#)



*The information on this page **does not apply** to Confluence OnDemand.*

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](#)

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.

Related pages:

- [Setting Up an External Search Tool to Index Confluence](#)
- [Configuring the Confluence Search and Index](#)
- [Confluence Administrator's Guide](#)

 The information on this page *does not apply to* Confluence OnDemand.

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](#)
- [Confluence Administrator's Guide](#)

 The information on this page *does not apply to* Confluence OnDemand.

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 angled 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. Manually enter your **Host Address**, **User Name** and **Password** details (recommended)

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 [Subscribing to Email Notifications of Updates to Confluence Content](#).
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 **Tools > Watch**. See [Watching 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

Set up a mail session for the Confluence distribution to use Gmail as follows:

1. Stop Confluence.
2. Move (don't copy) activation-1.0.2.jar and mail-1.4.1.jar from <confluence-install>/confluence/WEB-INF/lib to <confluence-install>/lib.
Note: The version numbers on these jar files may vary, but that should not matter. As of Confluence 5.2.3, activation-1.0.2.jar no longer exists, and does not need to be moved or downloaded
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. Submit, 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 [Subscribing to Email Notifications of Updates to Confluence Content](#).

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:

- [Subscribing to Email Notifications of Updates to Confluence Content](#)
- [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 OnDemand.*

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:

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.

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](#)

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 `char`s 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.

 The information on this page *does not apply to* Confluence OnDemand.

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.

Iñtërnâtiônàlizætiøn 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

Iñtërnâtiônàlizætiøn

Test 2: Form submission

This is the test string pasted by you into the web form and submitted back to Confluence

Iñtërnâtiônàlizætiøn

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

iñtërnâtiônàlizætiøn

Expected result (converting Java string to lowercase)

iñtërnâtiônàlizætiøn

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

IÑTËRNÂTIÔNÀLIZÆTIØN

Expected result (converting Java string to uppercase)

IÑTËRNÂTIÔNÀLIZÆTIØN

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.



The information on this page *does not apply* to Confluence OnDemand.

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:

 The information on this page *does not apply to* Confluence OnDemand.

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

INtERNAtIOnAlIZæTIOŃ 

Expected result (converting Java string to uppercase)

INtERNAtIOnAlIZæTIOŃ

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](#)
- [Enabling Trackback](#)
- [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.

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 be 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.

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.

On this page:

- [Introduction to Confluence's WebDAV Client Integration](#)
- [Restricting WebDAV Client Write Access to Confluence](#)
- [Disabling Strict Path Checking](#)
- [Virtual Files and Folders](#)
- [Using a WebDAV Client to Work with Pages](#)
- [Known Issues](#)

Related pages:

- [Global Permissions Overview](#)
- [Disabling and enabling add-ons](#)
- [Disabling or Enabling Confluence Add-ons](#)
- [Attachment Storage Configuration](#)

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.
 4.  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
☒	☒	☒

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.

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 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.

Setting Up a WebDAV Client 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 [Troubleshooting WebDAV](#) 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 8080, which is the default port value used by the popular application server Apache Tomcat that runs many Confluence EAR / WAR installations, or 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' webdavs (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.

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 the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Update the **Long Number Format** and **Decimal Number Format** to suit your requirements.
5. 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 [[searchterms@google](#)] instead.

Here is a screenshot showing the shortcuts currently defined on <http://confluence.atlassian.com>:

Key	Expanded Value	Default Alias	Operations
cache	http://www.google.com/search?q=cache:		Remove
imdb	http://us.imdb.com/Title?		Remove
jira	http://jira.atlassian.com/secure/QuickSearch.jspx?searchString=	JIRA Issue %s	Remove
googlegroups	http://groups.google.com/groups?q=		Remove
google	http://www.google.com/search?q=		Remove
dictionary	http://www.dict.org/bin/Dict?Database=*&Form=Dict1&Strategy=*&Query=		Remove

Shortcut links are added and maintained by Confluence administrators from the **Administration Console**.

On this page:

- [Creating shortcut links](#)
- [Using shortcut links](#)
- [Deleting shortcut links](#)

Related pages:

- [Working with Links](#)
- [Confluence Administrator's Guide](#)

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. (Optional. Available in Confluence version 2.3 and later.) 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 [Working with 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	Demonstration
a JIRA issue	CONF-1000@JIRA	http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=CONF-1000	CONF-1000
a Google search	Atlassian Confluence@Google	http://www.google.com/search?q=Atlassian+Confluence	Atlassian Confluence@Google

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

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**.

Related pages:

- [Choosing a Default Language](#)
- [Installing a Language Pack](#)
- [Confluence Administrator's Guide](#)

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.

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**.

Related pages:

[Confluence Remote API Reference](#)

 *The information on this page does not apply to Confluence OnDemand.*

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.

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 the **Threaded Comments** checkbox to enable threaded mode. Deselect the check box to disable threaded mode and display all comments in flat mode.
5. Choose **Save**.

Related pages:

- [Commenting on pages and blog posts](#)
- [Confluence Administrator's Guide](#)

Enabling 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 sent by other sites.

To enable trackback:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Select **Further Configuration** in the left panel.
3. Choose **Edit**.
4. Select the **Trackback** checkbox then **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](#)

 *The information on this page **does not apply to** Confluence OnDemand.*

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](#).

 *The information on this page **does not apply** to Confluence OnDemand.*

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.

 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.

RELATED TOPICS

[How to Edit Files in Confluence JAR Files](#)

[Administrators Guide Home](#) [Confluence Documentation Home](#)

Configuring System Properties

This page describes how to set Java properties and options on startup for Confluence Stand-alone and EAR/WAR versions.



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 **JAVA_OPTS=**
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 JAVA_OPTS=%JAVA_OPTS%**
3. Refer to the list of parameters [below](#).

i Add all parameters in a space-separated list. Make sure to keep the string %JAVA_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:

```
tomcat6w //ES//<SERVICENAME>
```

- i** In the above example, it would be `tomcat6w //ES//Confluence121213135538`
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

Property	Since	Default Value	Module...	Effect
atlassian.forceSchemaUpdate	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.
confluence.home	1.0	Any filesystem path	Confluence and atlassian-config	If this system property is set, Confluence will ignore the contents of the confluence-init.properties file, and use this property as the setting for the Confluence Home directory.

<code>confluence.dev mode</code>	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.
<code>confluence.disable.mailpolling</code>	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
<code>confluence.i18n.reloadbundles</code>	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> .

<code>confluence.ignore.debug.logging</code>	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.
<code>confluence.jmx.disabled</code>	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>
<code>confluence.optimize.index.module</code>	2.2	20	Confluence	Number of index queue flushes before the index is optimised.
<code>confluence.plugins.bundled.disable</code>	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.
<code>atlassian.mail.fetchdisabled</code>	3.5	false	Confluence	Disables mail fetching services for IMAP and POP
<code>atlassian.mail.senddisabled</code>	3.5	false	Confluence and atlassian-mail	Disables sending of mail

<code>atlassian.disable.caches</code>	2.4	true	atlassian-plugins, atlassian-cache-sevlet	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.
<code>confluence.html.encode.automatic</code>	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.
<code>org.osgi.framework.bootdelegation</code>	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.,com.yourkit." .
<code>confluence.diff.pool.size</code>	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).

<code>confluence.diff.timeout</code>	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.
<code>confluence.html.diff.timeout</code>	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.
<code>atlassian.user.experimentalMapping</code>	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 CONF-12319 , fixed in Confluence 3.1.1.

<code>confluence.import.use-experimental-importer</code>	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.
<code>atlassian.webrresource.disable.minification</code>	3.3	false	atlassian-plugins	Disables automatic minification of JavaScript and CSS resources served by Confluence.

<code>index.queue.thread.count</code>	3.3	See "Effect"	Confluence	Sets the number of threads to be used for the reindex job. The value has to be in the range of 1 to 10 (inclusive), i.e. at least one thread but no more than 10 threads will be used. There is no default value, i.e. • 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 10 threads will be used). • 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) • If you set <code>index.queue.thread.count=200</code>, then ten threads (the maximum allowed) will be used to reindex the content.

<code>index.queue.batch.size</code>	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.
<code>password.confirmation.disabled</code>	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: administrative actions , change of email address and Captcha for failed logins . Disabling password confirmations is useful if you are using a custom authenticator.

confluence.browser.language.enabled	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 choose a language preference .
upm.pac.disable	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 Atlassian Plugin Exchange . This is useful for application servers that do not have access to the Internet. See the UPM documentation .
confluence.reindex.documents.to.pop	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
confluence.reindex.attachments.to.pop	3.5.9	10	Confluence	Indicates how many attachments each indexing thread should process at a time during a full re-index.

<code>confluence.upgrade.active.directory</code>	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 - Sizerlimit Exceeded exceptions in your logs. For more details, see Unable to Log In with Confluence 3.5 or Later Due to 'LDAP error code 4 - Sizerlimit Exceeded'
<code>confluence.context.batching.disable</code>	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.

<code>com.atlassian.logout.disable.session.invalidation</code>	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.
<code>officeconnector.spreadsheet.xlsxmaxsize</code>	4.0.5	2 ²¹	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 CONF-21043 for more details.
<code>com.atlassian.confluence.extra.calendar3.display.events.calendar.maxpercalendar</code>		200	Team Calendars	Specifies the maximum number of events per calendar. This property is effective only if the Team Calendars plugin is installed on your Confluence site.
<code>com.atlassian.confluence.allow.downgrade</code>	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 After Downgrading, Confluence Will No Longer Run for details.

confluence.mbox.directory	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.
confluence.upgrade.recovery.file.enabled	5.5	true	Confluence	By default, Confluence will create 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.

<code>confluence.junit.report.directory</code>	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.

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](#).

Property	Since	Default Value	Module...	Effect
<code>atlassian.forceSchemaUpdate</code>	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.
<code>confluence.home</code>	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.

<code>confluence.dev mode</code>	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.
<code>confluence.disable.mailpolling</code>	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
<code>confluence.i18n.reloadbundles</code>	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> .
<code>confluence.ignore.debug.logging</code>	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.

<code>confluence.jmx.disabled</code>	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>
<code>confluence.optimize.index.module</code>	2.2	20	Confluence	Number of index queue flushes before the index is optimised.
<code>confluence.plugins.bundled.disable</code>	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.
<code>atlassian.mail.fetchdisabled</code>	3.5	false	Confluence	Disables mail fetching services for IMAP and POP
<code>atlassian.mail.senddisabled</code>	3.5	false	Confluence and atlassian-mail	Disables sending of mail
<code>atlassian.disable.caches</code>	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.

<code>confluence.html.encode.automatic</code>	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.
<code>org.osgi.framework.bootdelegation</code>	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., com.yourkit. ".
<code>confluence.diff.pool.size</code>	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).
<code>confluence.diff.timeout</code>	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.
<code>confluence.html.diff.timeout</code>	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.

<code>atlassian.user.experimentalMapping</code>	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 CONF-12319 , fixed in Confluence 3.1.1.
<code>confluence.import.use-experimental-importer</code>	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.
<code>atlassian.webrresource.disable.minification</code>	3.3	false	atlassian-plugins	Disables automatic minification of JavaScript and CSS resources served by Confluence.

<code>index.queue.thread.count</code>	3.3	See "Effect"	Confluence	Sets the number of threads to be used for the reindex job. The value has to be in the range of 1 to 10 (inclusive), i.e. at least one thread but no more than 10 threads will be used. There is no default value, i.e. • 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 10 threads will be used). • 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) • If you set <code>index.queue.thread.count=200</code>, then ten threads (the maximum allowed) will be used to reindex the content.

<code>index.queue.batch.size</code>	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.
<code>password.confirmation.disabled</code>	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: administrative actions , change of email address and Captcha for failed logins . Disabling password confirmations is useful if you are using a custom authenticator.

<code>confluence.browser.language.enabled</code>	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 choose a language preference .
<code>upm.pac.disable</code>	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 Atlassian Plugin Exchange . This is useful for application servers that do not have access to the Internet. See the UPM documentation .
<code>confluence.reindex.documents.to.pop</code>	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
<code>confluence.reindex.attachments.to.pop</code>	3.5.9	10	Confluence	Indicates how many attachments each indexing thread should process at a time during a full re-index.

<code>confluence.upgrade.active.directory</code>	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 - Sizerlimit Exceeded exceptions in your logs. For more details, see Unable to Log In with Confluence 3.5 or Later Due to 'LDAP error code 4 - Sizerlimit Exceeded'
<code>confluence.context.batching.disable</code>	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.
<code>com.atlassian.logout.disable.session.invalidation</code>	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.

officeconnector.spreadsheet.xlsxmaxsize	4.0.5	2 ²¹	Office Connector	Indicates the maximum size in bytes of an Excel file that can be viewed using the viewxls macro. If empty, the maximum size defaults to 2Mb. See CONF-21043 for more details.
com.atlassian.confluence.extra.calendar3.display.events.calendar.maxpercalendar		200	Team Calendars	Specifies the maximum number of events per calendar. This property is effective only if the Team Calendars plugin is installed on your Confluence site.
com.atlassian.confluence.allow.downgrade	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 After Downgrading, Confluence Will No Longer Run for details.

<code>confluence.mbox.directory</code>	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.
<code>confluence.upgrade.recovery.file.enabled</code>	5.5	true	Confluence	By default, Confluence will create 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.

<code>confluence.junit.report.directory</code>	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.

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. You can access the logs from the Confluence Administration Console, via the [support tool](#). If you cannot access the Confluence Administration Console, 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.

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](#)

 *The information on this page does not apply to Confluence OnDemand.*

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/atlas`

sian-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`.

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.

RELATED TOPICS

[Important Directories and Files](#)
[Enabling Detailed SQL Logging](#)
[Enabling user access logging](#)
[Generating a Thread Dump](#)
[Enabling Page Request Profiling](#)
[Troubleshooting Problems and Requesting Technical Support](#)

 [Administrators Guide Home](#)
 [Confluence Documentation Home](#)

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`.

 *The information on this page **does not apply** to Confluence OnDemand.*

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.
 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	
	INFO ▼	Add entry

Existing Levels

Class/Package Name	Current Level	New Level	
atlassian.plugin	INFO	INFO ▼	Remove
com.atlassian.confluence.admin.actions	INFO	INFO ▼	Remove

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.

 The information on this page *does not apply* to Confluence OnDemand.

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.

Reference : [log4j manual](#)

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](#)
- [Managing 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](#)
- [Running Confluence Over SSL or HTTPS](#)
- [Connecting to LDAP or JIRA or Other Services via SSL](#)
- [Configuring RSS Feeds](#)
- [Preventing and Cleaning Up Spam](#)

Related pages:

- [Giving People Access to Content](#)
- [Configuring a Confluence Environment](#)
- [Confluence Administrator's Guide](#)

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, passwords are hashed through SHA1 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.

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.

On this page:

- [Application Security Overview](#)
- [Finding and Reporting a Security Vulnerability](#)
- [Publication of Confluence Security Advisories](#)
- [Severity Levels](#)
- [Our Patch Policy](#)
- [Published Security Advisories](#)

Related pages:

- [Security Patch Policy](#)
- [Severity Levels for Security Issues](#)
- [How to Report a Security Issue](#)
- [Configuring Confluence Security](#)
- [Confluence Administrator's Guide](#)



The information on this page does not apply to Confluence OnDemand.

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 Patch Policy

Atlassian's approach to releasing patches for security issues is detailed in [Security Patch 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 `clearExpiredRememberMeTokenSTrigger`.

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
doc-sidebar	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.
confluence.list.pages.cookie	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.
confluence.browse.space.cookie	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.
confluence-language	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.
AJS.conglomerate.cookie	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 autocomplete, 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.

 The information on this page *does not apply* to Confluence OnDemand.

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 OnDemand.

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

 The information on this page *does not apply* to Confluence OnDemand.

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:

 The information on this page *does not apply* to Confluence OnDemand.

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>
```

Managing 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.

Page Information

Title:	The new logo
Author:	Sophie Staunton Feb 15, 2013
Last Changed by:	Sophie Staunton Feb 15, 2013
Tiny Link: (useful for email)	http://localhost:8090/x/HwA8
Export As:	Word · PDF
Operations:	Copy

Labels

There are no labels assigned to this page.

Hot Referrers

(2)	https://mail.google.com/mail/u/1/	[ban]
(1)	https://atlassian.hipchat.com/chat	[ban]

To manage your external referrers:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose '**Manage Referrers**'.

The following actions will be available:

- **Record or ignore all external referrers:** By default, Confluence records the number of hits made to a page from the link on the external site. If you turn this option off, Confluence will not record the hits.
- **Show or hide all external referrers:** By default, Confluence lists the external referrers as '**Hot Referrers**' on the '**Info**' screen of a page, as shown below. If you turn this option off, external referrers will not be listed on the page.
- **Specify which external referrers to exclude:** You can decide which referrers you want to exclude from being displayed on your site.

Screenshot: Manage external referrers

Manage Referrers

Record the external URLs that link to Confluence pages.

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.

Edit

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.

Add

Excluded Referrer URL Prefixes

URL Prefix

Operations

There are currently no Referrers being excluded.

Excluding 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.

You 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>

There are two instances where you may want to do this:

1. If you are 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. If 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.

i You need to be a Confluence administrator and to know the URL of the site to add it to the excluded referrers list.

To add a URL to the excluded referrers list:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.

2. Choose **Manage Referrers**
3. Enter the URL in the **Referrer URL Prefix** field (you must include <http://>)
4. 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.

Add

Excluded Referrer URL Prefixes

URL Prefix	Operations
	Purge All
http://evilspamsite.blogspot.com	Delete · Purge

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:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose '**Manage Referrers**'.
3. Deselect '**Show Referrers in Page Info**'.

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

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:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose '**Manage Referrers**' in the left-hand panel.

3. 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.

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.

 The information on this page *does not apply* to Confluence OnDemand.

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
```

 *The information on this page **does not apply** to Confluence OnDemand.*

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.

Related Topics

Content by label

There is no content with the specified labels





[Administrators Guide Home](#) [Confluence Documentation Home](#)

Configuring Captcha for Spam Prevention

You need to be a Confluence administrator to configure Captcha for spam prevention in Confluence.

If your Confluence site is open to the public 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.
- Edit a page.
- Send a request to the Confluence administrators.

Captcha is the technical term for a test that can distinguish a human being from an automated agent such as a web spider or robot. You can read more about Captcha on [Wikipedia](#).

When Captcha is switched on, 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



You can configure Confluence to enforce Captcha for certain types of users. You can exempt logged-in users (they will have completed a Captcha when they signed up) or members of particular groups.

By default, Captcha for spam prevention is disabled. If you enable it, the default is that Captcha for spam prevention will apply to anonymous users only. Only anonymous users will have to perform the Captcha test when creating comments or editing pages. Captcha images will not be shown to logged-in users.

Related pages:

- [Configuring Confluence Security](#)
- [Confluence Administrator's Guide](#)

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.

 Shortcut links (e.g. CONF-2622@JIRA) and internal links to other pages within Confluence are not tagged.

 *The information on this page **does not apply** to Atlassian OnDemand sites with multiple apps. If you are using Confluence-only OnDemand, the information does apply.*

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.

Related Topics

Content by label

There is no content with the specified labels



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 the technical term for a test that can distinguish a human being from an automated agent such as a web spider or robot. You can read more about Captcha on [Wikipedia](#).

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](#)

Related pages:

- [Confluence Administrator's Guide](#)

 *The information on this page **does not** apply to Confluence OnDemand.*

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 \(Cross Site Request Forgery\)](#) 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



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.

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 or Confluence WAR distribution using Tomcat.** 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 7.** JDK 1.7 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. If you are using JDK 1.6 or older, please refer to the [Java SE documentation](#) to see the differences in the `keytool` utility from your JDK to JDK 1.7.

i The default connector port for Confluence is 8090, while a plain Tomcat installation (used for EAR / WAR distribution) will default to 8080.

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](#)



The information on this page *does not apply* to Confluence OnDemand.

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 -alias tomcat -keyalg RSA
```

- On OS X or UNIX-based systems, run the following command at the command prompt:

```
$JAVA_HOME/bin/keytool -genkeypair -alias tomcat -keyalg RSA
```

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'`.
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 `.keystore` file 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.

i 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_PKC12_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_PKC12_KEYSTORE_NAME>
-srcstoretype PKCS12 -srcstorepass <MY_PKC12_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: <http://tomcat.apache.org/tomcat-5.5-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](#).
- If any of your users will access Confluence from **Internet Explorer 7 on Vista**, please note the following additional points when using Java's keytool utility:
 - Make sure that you specify the -keyalg RSA option, as shown in the example of the keytool command above. The default is the SHA1 algorithm, which results in an error 'Internet Explorer cannot display the webpage' on IE7 on Vista.
 - You may also need to specify the -sigalg MD5withRSA option. Otherwise, SHA1 will be used even if you specify the -keyalg RSA option. See [this Atlassian blogpost](#) for more information.
- **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_INSTALLATION>/conf/server.xml and add the following line within the <Context ... />element:

```
<Valve className="org.apache.catalina.authenticator.NonLoginAuthenticator"
    disableProxyCaching="true" securePagesWithPragma="false" />
```

Related Topics

- [SSL Configuration HOW-TO](#) in the Apache Tomcat 6.0 documentation
- [SSL Configuration HOW-TO](#) in the Apache Tomcat 5.5 documentation
- [keytool - Key and Certificate Management Tool](#) in the Java SE documentation
- [Connecting to LDAP or JIRA or Other Services via SSL](#)
- [Supported Platforms](#)

Connecting to LDAP or JIRA or Other Services via SSL

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 JIRA 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](#).

 *The information on this page **does not** apply to Confluence OnDemand.*

 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](#)

Related Topics

[Configuring an SSL Connection to Active Directory](#)
[Configuring Web Proxy Support for Confluence](#)
[Running Confluence Over SSL or HTTPS](#)

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**.

On this page:

- [Notes](#)

Related pages:

- [Using the RSS Feed Builder](#)

 *The information on this page does not apply to Confluence OnDemand.*

Screenshot: Configuring RSS feeds

Maximum RSS Items	200
Limit the maximum number of items an RSS Feed can request.	
RSS timeout	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 you have a public-facing Confluence site, your site 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, it is easier to delete them via SQL, as described below.

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 Confluence 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 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.

Notes

- See [CONF-1469](#). Your comments that issue are very much appreciated.

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 need to have [System Administrator](#) permissions in order to configure and execute jobs.

Accessing Confluence's Scheduled Jobs Configuration

To access Confluence's Scheduled Jobs configuration page:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Choose '**Scheduled Jobs**' to open the 'Scheduled Jobs' page. For each job listed down this page, the following information is shown:
 - **Job** — the name of a job.
 - **Status** — the job's status, which is either 'Scheduled' (it is currently enabled) or 'Disabled'. See [below](#) for details on disabling or re-enabling a job.
 - **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's last execution.
 - **Actions** — allows you to configure the job, execute it manually, view a history of previous executions or disable the job.

On this page:

- [Accessing Confluence's Scheduled Jobs Configuration](#)
- [Executing a Job Manually](#)
- [Configuring 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](#)

 The information on this page *does not apply to* Confluence OnDemand.

Screenshot: Scheduled Jobs

Scheduled Jobs

Job	Status	Last Execution	Next Execution	Avg. Duration	Actions
Back Up Confluence	Scheduled	14-Feb-2013 08:07:55	20-Feb-2013 02:00:00	9079	History · Run · Edit · Disable
Check Cluster Safety	Scheduled	19-Feb-2013 15:01:00	19-Feb-2013 15:01:30	3	History · Run
Clean Index Queue	Scheduled	14-Feb-2013 08:07:55	20-Feb-2013 02:00:00	3869	History · Run · Edit · Disable
Clean Temporary Directory	Scheduled	14-Feb-2013 08:07:55	20-Feb-2013 04:00:00	3759	History · Run · Edit
Clear Expired Mail Errors	Scheduled	14-Feb-2013 08:07:59	20-Feb-2013 03:00:00	0	History · Run · Edit
Clear Expired Remember Me Tokens	Scheduled		20-Feb-2013 00:00:00	0	Run · Edit
Email Daily Reports	Scheduled	14-Feb-2013 08:07:55	20-Feb-2013 00:00:00	5304	History · Run · Edit · Disable
Flush Did You Mean Index	Scheduled	19-Feb-2013 14:00:00	19-Feb-2013 16:00:00	15	History · Run · Edit · Disable
Flush Index Queue	Scheduled	19-Feb-2013 15:01:26	19-Feb-2013 15:01:27	0	History · Run
Flush Local Task Queue	Scheduled	19-Feb-2013 15:01:00	19-Feb-2013 15:02:00	0	History
Flush Mail Queue	Scheduled	19-Feb-2013 15:01:00	19-Feb-2013 15:02:00	0	History · Run · Edit · Disable
Flush Task Queue	Scheduled	19-Feb-2013 15:01:00	19-Feb-2013 15:02:00	0	History · Run · Disable
Optimise Indexing	Scheduled	19-Feb-2013 15:00:00	20-Feb-2013 03:00:00	361	History · Run · Edit
scheduledjob.desc.refreshSearcherJob	Scheduled	19-Feb-2013 15:01:00	19-Feb-2013 15:02:00	0	History · Run · Edit · Disable
scheduledjob.desc.summaryEmail	Scheduled	19-Feb-2013 15:00:00	19-Feb-2013 16:00:00	506	History · Run · Disable

Executing a Job Manually

1. Access the 'Scheduled Jobs' configuration page ([above](#)).
 2. Locate the job you wish to execute manually and click its '**Run**' link in the 'Actions' column. The job will be run immediately.
 - ✔ Refer to 'Types of Jobs' ([below](#)) for detailed descriptions about each job.
-  Not all jobs can be run manually.

Configuring a Job's Schedule

1. Access the 'Scheduled Jobs' configuration page ([above](#)).
 2. Locate the job whose schedule you wish to configure and click its '**Edit**' link in the 'Actions' column. The job's 'Edit Schedule for job' dialog box opens.
 - ✔ Refer to 'Types of Jobs' ([below](#)) for detailed descriptions about each job.
 3. Enter an appropriate cron expression to define the frequency with which the job is executed.
 - ✔ Refer to 'Cron Expressions' ([below](#)) for more details about their syntax. To revert the job's schedule back to its default settings, click the '**Default**' button.
 4. Click '**Save**' to record your job's new schedule.
-  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.

1. Access the 'Scheduled Jobs' configuration page ([above](#)).
2. Locate the job you wish to disable/re-enable.
 - ✓ Refer to 'Types of Jobs' ([below](#)) for detailed descriptions about each job.
 - If a job is enabled, click its **Disable** link in the 'Actions' column to disable the job.
 - If a job is disabled, click its **Enable** link in the 'Actions' column to enable the job.

i Not all jobs in Confluence can be disabled.

Viewing a Job's Execution History

1. Access the 'Scheduled Jobs' configuration page ([above](#)).
2. Locate the job whose execution history you wish to view and click the **History** link.
 - i** If a job has not completed at least one execution, its 'History' link will not be available.
 - ✓ Refer to 'Types of Jobs' ([below](#)) for detailed descriptions about each job.

The 'History for job' dialog box opens, showing a list of previous executions of the job in reverse chronological order, including the:

 - Start date and time
 - End date and time
 - The length of time (in milliseconds) that it took to complete the job

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 backup 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 Index Queue	Triggers a periodical clean of the index queue to ensure that its size does NOT grow indefinitely.	Per cluster	At 2am every day
Clean Temporary Directory	Cleans up temporary files generated in the 'temp' subdirectory of the Confluence home directory. This temp directory may be created by exports etc.  This does not 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 mail error queue . 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 to all subscribers.  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 Did You Mean Index	Note: The 'Did You Mean' feature is no longer available in Confluence. This job is therefore redundant, and will be removed at some time in the future. Flushes changes to the 'Did You Mean' index, which keeps the 'Did You Mean' feature up to date. Confluence records each content update in the 'Did You Mean' index.	Per node	Every 2 hours from 12 am
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 minute
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 mail queue .	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
Optimise Indexing	Compacts the confluence indexes to maintain searching performance. ✔ This task is demanding on system resources and does not need to be performed too regularly. If you see Confluence performance deteriorate around 3pm, try scheduling this job for 3am only and check if search performance remains reasonable.	Per node	At 3am and 3pm every day
Poll Mail	Polls POP accounts on all spaces that have them configured.	Per cluster	Every minute

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 do need to manually add them.

Note: The HTML Include macro is disabled by default in Confluence. See [enabling HTML macros](#).

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/th ispage</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. Unless your instance is running in an environment without internet access, we do not recommend disabling the whitelist.

Operating Large or Mission-Critical Confluence Installations

Request for interest

Hi there,

We are currently working on a completely new version of our clustered offering that provides greater scalability as well as High Availability.

If you are interested in joining our closed beta program for the new offering, please, send an email to kelvin@atlassian.com with the subject: **Confluence Clustered** and I will get back to you ASAP.

Cheers,

Kelvin Yap
Enterprise Advocate

This page gives guidelines for operational management teams who are responsible for a large Confluence installation, or for a Confluence installation which is crucial to the business of their organisation.

Introduction to this Page

Motivation for Presenting these Guidelines

Most Confluence installations start off small. Ten people in an early-adoption department use it for a couple of weeks. Everything works well and the good news starts spreading. Adoption increases throughout the organisation. More and more people use the wiki, and more and more rely on Confluence being up and running. After a while even the CEO starts blogging. And then a system outage occurs.

Now what?

Wikis like Confluence often grow into mission-critical applications within just a few months. Often adoption is so fast that IT departments haven't had the time to scale up their support.

We have assembled some requirements to help you make sure that your installation of Confluence can be mission critical. There are no surprises to be found here — all of the requirements would apply to any other piece of software that is mission critical within your organisation.

Who should Read these Guidelines?

The guidelines **do not apply to you** if you are using Confluence with just a few dozen users, and no one really minds if Confluence is down for a couple of hours because your database has crashed.

But if any one of the following applies to you, then these guidelines are **a must read for you!**

- The wiki has become your organisation's documentation base.
- Your users can't work properly when Confluence is down.
- Your boss or customer threatens to terminate your contract if you don't meet a strict service level agreement (SLA), such as 99.9% availability.

On this page:

- [Motivation for Presenting these Guidelines](#)
- [Who should Read these Guidelines?](#)
- [Dedicated Hardware for Confluence](#)
- [Dedicated Qualified Staff](#)
- [Constant Monitoring of Production Systems](#)
- [Adherence to Strict Upgrade Procedures](#)
- [Testing of Upgrades before Production Implementation](#)
- [Enforcing Security Guidelines](#)
- [Load-Testing Environments](#)
- [Tuning](#)
- [Related Topics](#)



The information on this page does not apply to Confluence OnDemand.

Requirements of Large or Mission-Critical Confluence Installations

Dedicated Hardware for Confluence

In a small work group with a few dozen or even hundreds of users, your Confluence installation can happily share the CPUs, memory and disks with other low-profile applications and a database.

But with thousands or even tens of thousands of users, you need dedicated hardware that runs Confluence and nothing else, and it needs to be fast hardware with plenty of RAM. While you *can* run Confluence in a virtualised environment such as VMware, we suggest you don't do it for mission-critical or high-load installations unless you are a real expert in virtualisation. Otherwise your other VMs might have performance problems which propagate to Confluence.

If you experience database-related problems, you should consider moving the Confluence database to a dedicated machine. Confluence itself can run queries that impact the performance of other applications, and other application problems or scheduled tasks can have an adverse affect on the usability of Confluence.

Dedicated Qualified Staff

If your Confluence installation is mission critical and your service level agreements require 24/7 up time, you need to be able to pinpoint problems quickly. You need qualified staff, dedicated to looking after Confluence, who are available during business hours and possibly beyond.

If you require assistance from the [Atlassian Support](#) team, you may need to answer some pretty technical questions to help us diagnose what is going on in your systems. Also keep in mind that Atlassian support assists you in finding problems in Confluence, but we can't help you administer your systems.

In particular, we recommend that you have dedicated staff in the roles listed below.

Operations Team with General Administrators

If your organisation relies on Confluence being up and running around the clock with very little downtime, you need people who can set up, maintain, tune and improve your Confluence installation. This requires at least one person, but ideally you will have a team of operational engineers.

If your wiki is mission critical, chances are that other IT systems within your organisation have already made it necessary to have such an operations team. So you will probably not need to hire someone specifically to administrate Confluence. But it is vital that supporting and maintaining Confluence is added to the list of responsibilities of that operations teams, and that you can get them to troubleshoot and analyse Confluence at short notice.

If problems arise and you need to contact [Atlassian Support](#), these engineers will be our first point of contact. We may ask them to provide details of log files, application-server settings, monitoring systems, and so on.

Network Staff

If Confluence is mission critical for large numbers of users, it is vital that you have dedicated network staff available to track down problems when they arise.

A mission-critical installation will usually be used by hundreds or even thousands of users, and you don't want to keep them waiting because a network card breaks, or because someone has made an undocumented change to the network and you don't have an expert around who can figure it out.

Again, this only applies to mission-critical systems. If you use Confluence for less critical collaboration and knowledge sharing, and a broken network cable causing a day's downtime is no major catastrophe, then you will not need dedicated networking staff.

Database Staff

If Confluence is mission critical for a large number of users, you need an experienced database administrator (DBA) available to troubleshoot database performance issues and other potential problems. It is dangerous not to have an experienced full-time DBA at hand at short notice when running a mission critical application. While

small installations of Confluence basically work 'out of the box', any system that involves high load or high-availability requirements needs continual monitoring, optimising and fine tuning of the Confluence database. Database monitoring is no trivial task — it's not something that anyone can learn quickly.

Developers

You may have decided to customise Confluence by changing its source-code, or by writing your own plugins. If your server is mission-critical, you must nominate staff who will be responsible for that code, and they must be up for the task. Otherwise you might end up in a situation in which your server experiences downtimes because of custom code is broken, or does not work with a newer version of Confluence anymore, but you can't fix the problem because no one knows how the customized code works, and you can't uninstall it either because it has become critical for your Confluence usage pattern. Keep good track of changes, and have someone available to jump into action if there is a problem. Don't let the summer intern write mission-critical plugins, unless you have more senior staff to maintain that code as long as it is in use.

Constant Monitoring of Production Systems

You will need to monitor your production systems constantly.

When the wiki is the lifeblood of your organisation, you need know exactly what is going on inside, so that you can plan for future needs and analyse potential bottlenecks.

Monitoring involves a number of essential tasks, including those listed below:

- Monitoring log files.
- Checking for HTTP-availability and performance (e.g. by getting the same page every five minutes and displaying the time on a graph).
- Looking at many different parameters such as load, connections, IO, database-trends, and so on.
- Charting long-term trends.
- Keeping an access log of requests to the web server. This is vital, especially when requesting [performance-related support](#) from Atlassian.


Monitoring a web application like Confluence implies also monitoring the subsystems it uses. Many outages and downtimes are caused by broken mail servers, databases running out of space, file systems filling up and so on. It is often possible to detect these trends way before the actual web application breaks down. Keep an eye on the file system, and if you see it is getting closer to 90% utilisation, you can mend the situation without Confluence breaking down. Or even if the worst case happens (e.g. the database breaks down and Confluence is affected straight away) then having the proper monitoring for the database server makes troubleshooting a lot easier.



Tools for Monitoring Confluence


At Atlassian we use [Hyperic](#). But the list of monitoring systems is long and we can't recommend a specific product over the other. If your organisation has a monitoring system already, make sure you hook up Confluence to it. If you don't have a monitoring system yet, you need to install one as soon as you feel Confluence is mission critical.

As an example of what our monitoring UI looks like, have a look at this screenshot:



The following screenshot shows one of our sensors looking at the HTTP response times of our documentation

wiki over the last 8 days. You can clearly see an incident four days ago. Having the graph (and regularly looking at it) allowed us to pinpoint the problem. We analysed the access logs and found that webpage-profiling had been enabled but not disabled again, which caused performance problems.



This page would get too long if we described all our monitoring sensors - but just to give you an impression, this is what we monitor on the JVM level alone.

JVM basics

- Current Loaded Classes
- Daemon Thread Count
- Heap Memory Committed
- Heap Memory Max
- Heap Memory Used
- Loaded Classes
- Loaded Classes per Minute
- Object Pending Finalization Count
- Peak Thread Count
- Thread Count
- Unloaded Classes
- Unloaded Classes per Minute

JVM garbage collection

- Collection Count
- Collection Count per Minute
- Collection Time
- Collection Time per Minute

JVM memory: (Metrics for Eden space, Old Gen, Survivor space, Perm Gen)

- Committed Memory
- Used Memory

We get the same level of detail for our database, for the file system, for the CPU, for the network, and so on. Not all of this is needed all the time. But if your company depends on an application, then the more information you have at your fingertips the better. Fortunately these metrics can be extracted quite easily once you have a monitoring system in place.

Adherence to Strict Upgrade Procedures

Your organisation will have its own upgrading procedure. Here are a few recommendations that you should add to your list:

- Our main recommendation: Never change more than one component at a time. Sometimes it may be tempting to upgrade the server hardware when you upgrade Confluence, but we recommend you don't do that. It makes pinpointing errors much more difficult. So, for example, don't upgrade hard disks in conjunction with a Confluence version upgrade, don't change the Confluence configuration at the same time as you upgrade your Apache software, and don't upgrade a major third-party plugin the day you move your database system to a new machine. The list is endless, these were just a few examples to get you thinking.
- After each upgrade step, run Confluence for a couple of days to check that everything is still fine.
- Keep track diligently of what you change, and when. It will be nearly impossible for us to help you if you can't tell us what exactly you changed at what time.
- Keep a copy of all log files produced during the upgrade, together with notes about what changed between successive restarts.

Always take careful note of the upgrade notes published with the [Confluence Release Notes](#) of each Confluence

version, as well as the [Confluence Upgrade Guide](#).

Example

Here you can see an extract of our change log for <http://confluence.atlassian.com> — the server that hosts this very page.

Sydney time	Server time	Event	Reason/Purpose (including JIRA issues)
	2008-03-25 22:18	Started upgrade to 2.8-m9-r3 (build #1314)	
	2008-03-25 22:25	App server brought down due to failed database upgrade	
	2008-03-26 00:51	Server brought back up after database restored from backup. Running 2.8-m9-r3.	
	2008-03-28 04:18	GC algorithm changed from concurrent to parallel collector. Max heap increased from 1.4 GB to 2.0 GB	
	2008-04-24	Hyperic agent started with connection to Resin.	
	2008-05-08 20:30 - 22:30	Manual updates to menu.css, comments.js and comments.css in webapp	Temporary fix for @JIRA , @JIRA which was impacting performance
	2008-05-12	Updated cache sizes for five caches, bounced server.	Cache efficiency was low on these caches.
2008-05-13 18:00-18:20	2008-05-13 03:00-03:20	Upgrade from Resin 3.0 to Tomcat 5.5	
2008-05-14 16:30-17:00		Upgrade from Confluence 2.8.1-rc2 to 2.8.1-rc3	
	2008-05-14 20:30	Install new cronjob as j2ee for automating access log analysis	@JIRA

Testing of Upgrades before Production Implementation

You should test upgrades in a staging environment.

Before rolling out a new version of Confluence (or of the software or hardware that it uses, e.g. database systems, application servers, data storage), make sure that you test the upgrade with real data (e.g. a database dump) on a completely independent machine.

Here's an example of what such a test would pick up: The new release of Confluence may not be compatible with a custom third party plugin you have previously installed, thus breaking the plugin's functionality. You may not even know that anyone installed that plugin — but maybe many people are already using it. You'll want to find out about this before you actually roll out the new version of Confluence.

Here is an outline for a simple upgrade test:

1. Create a clone of your production environment, using a database dump to obtain a copy of the Confluence data. We'll call this your 'staging environment'.
2. Upgrade the staging environment to the new version of Confluence.
3. Ask a few selected users from different departments to check the pages they commonly access, but have them do it in the staging environment.

✔ Hint: In addition to finding weirdnesses with plugins, this may also show whether training for new functionality is needed in some of the departments. The IT department staff may be able to handle the upgrade to a new version of Confluence without training, but perhaps the sales representatives who use the wiki less often will need some training.



Getting a license for your staging environment

Confluence Knowledge Base: [How to get a Confluence developer license](#)

Enforcing Security Guidelines

Security is one of the most important issues for Confluence. We are constantly spending large amounts of effort to keep up with security threats and to Confluence's security model. We treat security breaches with utmost priority, and the recent releases have been improved to fend off advanced attack vectors like cross-site scripting (XSS), cross-site request forgery (XSRF) and header injection flaws. Altogether we believe that Confluence is a very secure product. But of course as with any software there are occasional bugs, and we are fixing security issues whenever they come up. We regularly release minor software releases that contain security fixes. This means you should upgrade your system frequently. Obviously this can affect your system's uptime. You should also make sure your whole infrastructure around Confluence is made robust as well (consider operating systems, web servers, application servers, networks, social engineering aspects, etc).

As with any other distributed system, you need to decide on a case by case basis if classified documents can be stored in it. It is common practice to store the most secure documents on computers that are not even connected to the physical intranet. Please contact your company's security officer to learn more about your enterprise's security procedures.

Make sure to have qualified staff around, so you can deal with security issues quickly. Once a security patch becomes available or a security incident happens, speed is essential.

Please refer to our dedicated [Configuring Confluence Security](#) page for more technical details.

Load-Testing Environments

Many customers ask us,

So, how many users and spaces can I put into Confluence, and what is the best hardware do to so?

The answer is, 'It depends'.

It depends a lot on your use case. Confluence is so successful because it can cover a huge range of use cases. If most of your users only access Confluence infrequently, it is no problem to have 70 000 to 100 000 users. But if each user is a power-user who uses the system the whole day, there's a substantial decrease in number Confluence can take without tuning. If your pages are short, simple, and don't contain a lot of macros, then the situation will be vastly different from a system that relies heavily on macros, background-tasks, or other features.

If your system is large (for example serving more than 10 000 users or storing more than 1000 spaces) or mission-critical (which it could be with as few as 1000 users who use it all the time) you need one or more load-testing environments.

Even if your system is working nicely for 20 000 users right now, it might take just another 2000 users to push it over the edge.

We recommend the following basic procedure:

- Set up an environment that closely resembles your production environment.
- Gather statistics from your production system.
- Regularly apply a similar kind of load (and slightly higher) to the load-testing environment.
- Analyse how well Confluence scales for your usage patterns.

Tuning

You may need to be able to tune your installation in the ways mentioned below.

Optimising your System

If you have large numbers of users, then downloading all the static content (CSS, default images, JavaScript-files) may result in a high additional load on the application server that can be offloaded to a caching web server.

Please refer to the following additional information:

- Our general [Performance Tuning](#) page.
- Information on [configuring a large Confluence installation](#).

Limiting Third-Party Plugins

You may have to restrict the number of third-party plugins installed on your Confluence instance.

Most third-party plugins are not specifically written for high-load environments. What works fine in low-load environments could have unexpected and adverse effects when thousands of users are competing for your application server's CPU time or for database IO.

A common source of problems is access to database connections. If you have fewer users than database connections, it does not matter if an operation holds on to a database connection for two seconds while it downloads some data from the internet. With hundreds of concurrent users, this could quickly become a bottleneck.

Confluence itself is tested and optimised to handle high loads and avoids these kinds of problems. But if you install a number of plugins that have not been tested against high load, your system may become unstable.

We recommend that you load test the common use cases of each unofficial third-party plugin if your Confluence installation is mission critical. Only activate plugins that are vital to your business, and never allow experimental plugins onto your production system until they have been tested in a staging environment.

Selecting and Tuning your JVM

You should select your JVM carefully and you may need to be able to tune it.

The selection of the JVM for your large Confluence instance can have a huge impact on the performance perceived by the users. Between versions 1.4 and 6 of the Sun Java JVM there have been some impressive improvements in performance, especially under high concurrent load.

Here are some essential guidelines:

- Always run the most recent point release of your selected JVM.
- Where ever possible run the most recent major release from your selected JVM manufacturer. The Sun JVM version 6 is much faster than 1.4, especially under high loads.
- Tune your garbage collection algorithms. Experiment with different algorithms and settings to get the response times you desire in your environment. Here are some specific guidelines for Sun JVM in the Sun documentation:
 - [Java 6](#)
 - [Java 5](#)
 - [Java 1.4](#)

Customising Confluence to Optimise Performance

You may need to customise Confluence for performance reasons. Depending on your usage scenario, there may be ways to enhance Confluence performance that become necessary when you reach a certain level of usage.

Here are some things you might decide to do:

- Remove the display of the space list on the Dashboard. See [Customising the Confluence Dashboard](#).
- Configure any search appliances or other crawlers which are configured to index the Confluence site:
 - These should be suitably rate limited.
 - Configure them to crawl only pages in the `/display/` URL path, and only current versions of

pages.

Please refer to our general [Performance Tuning page](#) for more details.

Related Topics

[Performance Tuning](#)
[Configuring a Large Confluence Installation](#)
[Confluence Clustering Overview](#)
[Requesting Performance Support](#)
[Confluence Administrator's Guide](#)
[Configuring Confluence](#)
[Server Hardware Requirements Guide](#)
[How to Fix Out of Memory Errors by Increasing Available Memory](#)

Configuring a Large Confluence Installation

Deploying *any* application to several thousand users requires care and planning, especially if those users are going to be relying on the application to get their work done.

General Advice

Staged Rollout

Do not try to deploy Confluence immediately to your whole organisation. Instead, roll it out department by department, or project by project.

How Confluence will scale given a particular software and hardware configuration depends very much on how Confluence is likely to be used in your organisation. Launching Confluence to everybody at once may seem like a neat idea, but it also means that any problems you might experience scaling the system up to your entire organisation will hit you *all at once*, annoy everyone and possibly hurt adoption.

Rolling Confluence out gradually will give you the chance to tune it as you go, resulting in a much more painless experience. There will also be organisational advantages: you can identify those teams or projects who are most likely to be successful 'early adopters', and those teams can experiment with how best a wiki might suit your organisation, and pass on their 'best wiki practices' as usage of Confluence expands.

Plugin Governance

Confluence plugins can add tremendous value. Before adding one, visit the [plugin's page](#) and explore its issues (available from the issue management link). Try the plugin in a test environment and make sure to note any adverse effects after adding it to a production environment. Test plugins independently when upgrading.

Backup strategy

Disable the XML backup and use the [Production Backup Strategy](#).

New Spaces Governance

For both performance and good practice, put some modest governance in place around the creation of new spaces, such as a simple request that includes a check for duplicates and some strategy around how to best use a space. Duplicates and unused spaces should be purged by a wiki gardener. Try to keep it to [one space per group](#).

Performance Tuning

Check our guides for [Performance Tuning](#).

Choosing User Management and Single Sign-On

We recommend that you choose and configure your user management solution as soon as possible, rather than adding it to your Confluence installation at a later date.

It is possible to [integrate with an LDAP repository](#), such as Microsoft Active Directory, or add a single sign-on solution later (especially with the addition of [Crowd](#)). But if possible it is best to configure your user management system up front. You can configure access for only a specific group or set of groups, thereby keeping the

gradual rollout.

Please refer to our detailed guide to [Configuring User Directories](#) and examine the [User Management Limitations and Recommendations](#).

Configuring your Application Server, Web Server and Database

Because Confluence can be deployed in so many server combinations, we do not currently have guides on the best tuning parameters for each individual server. We will be happy to provide support, however. If you have any tuning parameters that you find particularly useful for Confluence instances, feel free to share them with other Confluence users in the [Confluence Community space](#).

Best Practices

Troubleshooting possible memory leaks

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. Many of the issues reported there are exacerbated with a large installation.

Memory Usage

The Java virtual machine is configured with a "maximum heap size" that limits the amount of memory it will consume. If Confluence fills up this maximum heap size it will run out of memory, and start behaving unpredictably. You can keep track of Confluence's memory usage from the System Information screen of the administration console:

Java VM Memory Statistics	
Total Memory	313 MB
Free Memory	140 MB
Used Memory	173 MB

Memory Graph  **45 % Free**

This example shows that, at the time of writing, confluence.atlassian.com is using 173MB of an allocated 313MB of heap. (The JVM was configured with a maximum heap size of 450MB, but this information is not available in the graph. The 313MB figure shows that the full 450MB of heap has not yet been needed)

Database Connection Pool

Confluence will need a database connection for each simultaneous user connection to the server. It is also a good idea to have 5-10 connections spare for Confluence internal processes such as backups, re-indexing or daily notification jobs.

Running out of pooled connections will cause the server to slow down as more users are waiting for a connection to be freed before starting their own request, and will eventually cause visible system errors as Confluence times out waiting for a database connection.

If you are using Confluence's internal connection pool, you can increase the number of available connections by modifying the `hibernate.c3p0.max_size` property in `{confluence_home}/confluence-cfg.xml`, and restarting Confluence. **Make sure** you have also configured your database to be able to support that many simultaneous connections.

Cache Sizes

The [Performance Tuning](#) page includes some useful rules of thumb for configuring the sizes of Confluence's internal 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](#).

RELATED TOPICS

[Operating Large or Mission-Critical Confluence Installations](#)
[Performance Tuning](#)
[Confluence Clustering Overview](#)
[Requesting Performance Support](#)
[Managing Confluence Users](#)
[Confluence Administrator's Guide](#)
[Configuring Confluence](#)

Confluence Clustering Overview

Request for interest

Hi there,

Great to see your interest in Confluence Clustered!

We are currently working on a completely new version of our clustered offering that provides greater scalability as well as High Availability.

If you are interested in joining our closed beta program for the new offering, please, send an email to kelvin@atlassian.com with the subject: **Confluence Clustered** and I will get back to you ASAP.

Cheers,

Kelvin Yap
Enterprise Advocate

It is possible to run Confluence in a clustered environment instead of on a single server. This means that you can run multiple copies of Confluence in a cluster, so that clients (such as a browser) can connect to any copy and see the same information.



Consider your options carefully before deciding on a clustered installation

While we have tried to make clustering Confluence as easy and administrator-friendly as possible, it is a major architectural change and requires extra planning for deployment and upgrades. Please consider the information on the [Cluster Checklist](#) and then consult [Atlassian support](#) before making your final decision.

This page gives an overview and links to further pages with information on installing, configuring and administering a Confluence cluster.



*The information on this page **does not apply** to Confluence OnDemand.*

Before Deciding to Run a Confluence Cluster

1. Read and consider the details on the [Cluster Checklist](#).
2. Consider the difference between [clustering for scalability](#) and [clustering for high availability \(HA\)](#).
3. Contact [Atlassian support](#) for further information and advice.

Technical Overview

Read a [technical overview](#) of clustering in Confluence.

Server and Network Requirements

- [Server hardware requirements](#)
- [Technical overview of Confluence clustering](#)
- [Diagram of recommended network topology](#)

Installation and Upgrading

There are two methods of installing Confluence in a cluster, depending on whether you have existing data:

- [Fresh installation](#)
- [Existing data](#)

If you are upgrading an existing Confluence cluster to a new version of Confluence, refer to the [cluster upgrade guide](#).

Configuration and Administration

- [Cluster Administration page](#) in the **Administration Console**
- [Changing datasources in clusters](#)

Troubleshooting

- [Cluster troubleshooting](#)

RELATED TOPICS

[Operating Large or Mission-Critical Confluence Installations](#)
[Performance Tuning](#)
[Requesting Performance Support](#)
[Confluence Administrator's Guide](#)
[Configuring Confluence](#)

Technical Overview of Clustering in Confluence



Overview of clustering documentation

Refer to the [overview of Confluence clustering](#) in the *Administrators' Guide*.



*The information on this page **does not apply** to Confluence OnDemand.*

Introduction

From version 2.3, Confluence has had the ability to configure and run multiple copies of itself in a cluster, so that clients can connect to any copy and see the same information. In effect, a Confluence cluster behaves as a single, powerful Confluence installation. While we have tried to make clustering Confluence as easy and administrator-friendly as possible, it is a major architectural change from earlier versions (or non-clustered installations) and consequently, requires extra planning for deployment and upgrades.

This document will give a technical overview of clustering in Confluence, primarily for those users and developers who will be installing and configuring Confluence in a cluster. A [separate overview](#) is available for Confluence plugin developers.

Cluster topology

A simple description of the cluster topology for Confluence would be **multiple applications, shared data source**. A cluster of Confluence consists of:

- multiple homogeneous installations of Confluence (called *nodes* below)
 - a Confluence home directory for each installation.
- a distributed Oracle Coherence cache (formerly known as Tangosol Coherence), which all nodes use via a multicast group - see [networking summary](#) below
- a single database, which all nodes connect to

The user is responsible for configuring an appropriate HTTP load balancer in front of the clustered installations.

Typically this means using `mod_jk` or another application server load-balancing technology. The load balancer must be configured to support **session affinity**.

Communication between clustered nodes is minimised by using a distributed cache which propagates updates to all other nodes automatically. Where necessary, Coherence provides a locking mechanism for synchronising jobs and a RMI interface for more complex communication.

LAN Clustering Only

Atlassian **only supports clustering over a local area network**. While it is theoretically possible to configure Confluence to cluster across a WAN, the latency involved is likely to kill performance of the cluster. If you do want to go down that path, you will need to configure Coherence yourself. Atlassian Support won't be able to support that kind of a configuration, but you can always enlist an [Atlassian Expert](#) to help.

Homogeneous Confluence installations

All the Confluence installations must be running exactly the same application, down to the lowest level. Items that must be the same include:

- Confluence version
- Application server version
- JDK version
- Libraries and plugins in the Confluence classpath, WEB-INF/lib
- Libraries in the application server classpath

The [installation section](#) has more information how to ensure homogeneous node installations.

Creating a Confluence cluster

When installing Confluence in a clustered setup, you will be responsible for configuring your web server and load balancer to distribute traffic between each node. No additional software is required as Coherence is bundled with Confluence.

Here is an overview of the process:

1. Obtain a [clustered licence key](#) from Atlassian for each node
2. Upgrade a single node to the clustered licence
3. Start the cluster from that node's administration menu, specifying a name and optionally a preferred network interface
4. Restart the single node and test it
5. Copy the Confluence application and Confluence home directory to the second node
6. Bring up the second node and it will automatically join the cluster.

Copying the Confluence application and home directory helps ensure that the installations are homogeneous.


An alternative to this method is to copy the Confluence web application, but not the Confluence home directory. In this case, the installation wizard will require your cluster name to connect to the other nodes, and it will automatically configure itself. You will need to rebuild the index manually after this installation, however.

There is now full documentation for a [Confluence Cluster Installation](#).

Upgrade process

Another consequence of the homogeneous requirement is that upgrades must be done by following a strict process.

1. All cluster nodes are brought down
2. Upgrade a single node to the latest Confluence version



Confluence cluster topology (simplified)

3. Start the single node so it can upgrade the database
4. Upgrade subsequent nodes and start them one-by-one.

This is the only safe method of upgrading a Confluence cluster.

Single database

The Confluence database in a cluster is shared by all nodes. This means that the database must be able to scale to service *all* the Confluence nodes, which will probably mean implementing some kind of database cluster and JDBC-level load balancing. We can not offer support with scaling or tuning your database, you will need to talk to your DBA or database vendor.

For obvious reasons, you must have an external database to run Massive - you can not cluster Confluence when using the embedded HSQL database.

The most important requirement for the cluster database is that it have sufficient connections available to support the expected number of application nodes. For example, if each Confluence instance has a connection pool of 20 connections and you expect to run a cluster with four nodes, your database server must allow at least 80 connections to the Confluence database. In practice, you may require more than the minimum for debugging or administrative purposes.

In a cluster, attachments must be stored in the database. Configuring a cluster in an existing installation will automatically migrate your attachments to the database. Non-clustered installations still have the option of using the Confluence home directory for storing attachments.

While attachments are stored in the database, they are temporarily written to the cluster node's local filesystem, designated `<confluence-home>/temp` folder, when being streamed to users (so Confluence doesn't have to hold open database connections unnecessarily). For this reason, Confluence will still need enough temporary disk space to hold any attachments currently in transit.

Distributed cache

In a normal configuration, Confluence uses many caches to reduce the number of database queries required for common operations. Viewing a page might require dozens of permissions checks, and it would be very slow if Confluence queried the database for this information with every page view. However, caches must be carefully maintained so they are consistent with the application data. If the page permissions change, the old invalid data needs to be removed from the cache so it can be replaced with a fresh correct copy.

To preserve consistent caches across a cluster, Confluence uses a distributed cache called [Oracle Coherence](#), which manages replicating cache updates transparently across all nodes. The network requirements of the distributed cache are quite simple, but must be preserved if the cluster is to work properly.

To discover other nodes in the cluster, Confluence broadcasts a join request on a multicast network address. Confluence must be able to open a UDP port on this multicast address, or it will not be able to find the other cluster nodes.

Once the nodes are discovered, each responds with a unicast (normal) IP address and port where it can be contacted for cache updates. Confluence must be able to open a UDP port for regular communication with the other nodes.

Because the Coherence network requirements are different to those required by the Confluence database connection, the situation can arise where Confluence can use the database but not talk to the other nodes in the cluster via Coherence. When Confluence detects this, it will shut itself down in a [cluster panic](#).

For more details on the network configuration of the distributed cache, see the [networking summary](#)

Home directory

Confluence's home directory has a much-reduced role in a cluster. Because the application data must be shared between all nodes for consistency, the only information stored in the Confluence home directory is either node-specific, or needed to start Confluence. This includes information related to:

- database connection
- license
- cluster connection

The only application data stored in the Confluence home directory is the **Lucene search index**. Confluence

synchronises this data itself by keeping track of indexing tasks in the database.

This is also why we recommend copying the Confluence home directory from the first node when setting up subsequent nodes. If you did not copy the Confluence home directory, you would need to rebuild the search index from scratch on the subsequent nodes after installation.

Event handling

Broadcasting events to all nodes in a cluster is supported in Confluence, but not recommended. The cluster topology uses a shared data store so that application state does not need to be synchronised by events.

The event broadcasting is done only for certain events, like installing a plugin. When a plugin is installed in one node, Confluence puts the plugin data in the database, and notifies the other nodes that they need to load the plugin into memory.

Indexing

Confluence maintains a copy of its Lucene search index on each node of the cluster. This index is used for many things beside full-text searches, including RSS feeds and lists of recently updated content. Indexing in a cluster works like this:

1. Node 1 gets a request to save some page update
2. After saving the page in the database, Node 1 adds a "page-updated" index entry to the queue, which is in the database
3. Periodically, each node picks up the "latest entries" from the queue, where what is latest is determined from a timestamp on a file in the Confluence home directory which indicates when the queue was last inspected. This process is called "flushing the index queue".
4. Each node independently updates its local Lucene index. The "page-updated" index entry is internally changed into a delete-document task and an add-document task to apply the changes to Lucene.
5. Each node updates the timestamp on its index-queue-timestamp file to reflect the most recent processing or "flushing" of the index queue.

Because of step #3, if the timing of the nodes is not synchronised or changes sporadically (due to a virtualisation environment, typically), index changes will not be correctly synchronised in the cluster. This is the most common cause of index sync problems in clusters.

If a node is disconnected from the cluster for a short amount of time (less than three hours), it will be able to bring its copy of the index up-to-date when it rejoins the cluster. If a node is down for a long amount of time and its lucene index has become stale as a result, you may want to avoid the expensive operation of rebuilding the index. To do that, you must copy a "live" version of the Lucene index from an active node. Simply replace the contents of the `Confluence Home]/index` directory with those from an active node before bringing the stale node back up.

Job synchronisation

For tasks such as sending the daily report emails, it is important that only one node in the cluster does this. Otherwise you would get multiple emails from Confluence every day.

Confluence uses locks in the Coherence distributed cache to ensure only one node can be running certain jobs at a time. This ensures email notifications will only be sent once.

Activity tracking

[Activity tracking](#) does not work in a cluster, and will be disabled for clustered deployments. We're working on making the activity tracker clusterable in a future release. You can follow [this issue](#). You can try some [other options](#) for tracking usage.

Cluster panic

In some situations, there can be a network issue or firewall that prevents the distributed cache from communicating but still allows Confluence to update the database. This is a dangerous situation because when the caches on the detached nodes become inconsistent, users on different nodes will see different information and updates can be lost.

Confluence can detect this problem by checking a database value against a cached value, and if they differ, all the clustered nodes will be shut down with a 'Cluster panic' message. This is considered a fatal error because

the consequences can cause damage to your data. For those administrators that like to live on the edge, there is a system property to prevent cluster panic and allow data corruption. For more information, see [Cluster safety mechanism](#).

If a cluster panic does occur, you need to ensure proper network connectivity between the clustered nodes. Most likely multicast traffic is being blocked or not routed correctly. See the [networking summary](#) below.

Summary of network requirements

In addition to normal connectivity with its database, all clustered Confluence instances require access to a multicast group and the ability to open a UDP unicast port.

By default, the multicast address is automatically generated from the cluster name you provide when starting the cluster and the multicast port is fixed. During cluster setup, Confluence will prompt for the unicast IP address to use if the server has multiple network interfaces, and by default the unicast port is fixed. The cluster multicast group will be joined on the same network interface as the bound unicast IP address.

For any settings which are not configurable through the Confluence web interface, they can be configured via an XML file in the Confluence home directory for more exotic networking requirements.

Scaling Confluence On A Single Server

Since the maximum addressable memory on a 32 bit JVM is 4GB, some large servers may scale Java applications by running JVM instances concurrently. This would be implemented as separate, clustered Confluence nodes running on a single server and communicating internally. Because each JVM replicates the cache entirely, it may be useful to test a single, massive instance running a 64 bit JVM as an alternative. This configuration may result in superior performance than an internal cluster.

Geographically Distributed Clusters

Collocating nodes is strongly recommended as high latency will almost certainly degrade performance due to the overhead of cache replication. Cluster nodes will provide the best performance if servers are physically adjacent. However, as long as all nodes share a LAN, users may wish to test alternative configurations to see how performance is affected.

RELATED TOPICS

[Server Hardware Requirements Guide](#)
[Overview of Confluence Clusters](#)
[Developers' Guide to Clustering](#)
[Cluster safety mechanism](#)

Introduction

A mechanism was added in Confluence 2.3 and above to ensure database consistency when running multiple cluster nodes against the same database. This is called the *cluster safety mechanism*, and is designed to ensure that your wiki cannot become inconsistent because updates by one user are not visible to another. A failure of this mechanism is a fatal error in Confluence and is called *cluster panic*.

Because the cluster safety mechanism helps prevents data inconsistency whenever any two copies of Confluence running against the same database, it is enabled in *all* instances of Confluence, not just clusters.

 The information on this page *does not apply to* Confluence OnDemand.

How cluster safety works

A scheduled task, ClusterSafetyJob, runs every 30 seconds in Confluence. In a cluster, this job is run only on one of the nodes. The scheduled task operates on a *safety number* – a randomly generated number that is stored both in the database and in the distributed cache used across a cluster. It does the following:

1. **Generate** a new random number
2. **Compare the existing safety numbers**, if there is already a safety number in both the database and the

cache.

3. **If the numbers differ, publish a ClusterPanicEvent.** Currently in Confluence, this causes the following to happen:
 - disable all access to the application
 - disable all scheduled tasks
 - update the database safety number to a new value, which will cause all nodes accessing the database to fail.
4. If the numbers are the same or aren't set yet, **update the safety numbers:**
 - set the safety number in the database to the new random number
 - set the safety number in the cache to the new random number.

How to fix it

See [Confluence will not start due to fatal error in Confluence cluster](#).

Technical details

The cluster safety number in the database is stored in the CLUSTERSAFETY table. This table has just one row: the current safety number.

Changing Datasources Manually in a Cluster



The recommended way of changing database connections is to shut down the whole cluster, install Confluence into new and empty directories and use the [Setup Wizard](#) to configure all new database connection settings.

However, if you wish to manually change your settings, you may proceed as described below.



It is **strongly recommended** that you test all of the following in a staging or test instance of Confluence before performing these steps in your production environment.



*The information on this page **does not apply** to Confluence OnDemand.*

Step 1: Prepare

- Locate the `confluence-cfg.xml` file in the Confluence home directory.
- Make a backup copy of that file.
- Prepare the necessary changes to that file.

Step 2: Shut Down Confluence

You need to shut down all the nodes in the cluster, not just one.

Step 3: Apply your Changes

Apply your configuration changes to the required node.

Step 4: Restart the Changed Node

It is crucial that you bring up the node on which you applied the changes **first**. Otherwise you will get an error message, and have to shut down all instances again.

Step 5: Restart all Other Nodes

Done.

RELATED PAGES

[Overview of Confluence Clusters](#)

Cluster Troubleshooting

This page covers troubleshooting for a clustered installation of Confluence.

- For information about clustering in general, refer to the [overview of Confluence clustering](#).
- 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](#)
- [Confluence cluster debugging tools](#)
- [Didn't find a solution?](#)
- [Content by label](#)
- [Related](#)

 *The information on this page **does not** apply to Confluence OnDemand.*

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	'Database is being updated by an instance which is not part of the current cluster' Error Message
Database is being updated by an instance which is not part of the current cluster errors on a cluster	Add multicast route , Check firewall , Cluster Panic due to Multiple Deployments
Cannot assign requested address on startup, featuring an IPv6 address	Prefer IPv4
Error in log: The interface is not suitable for multicast communication	Change multicast interface , Add multicast route
Multicast being sent, but not received (detectable with Multicast Test)	Check firewall , Check intermediate routers , Increase multicast TTL
Any issue not covered here	Contact support

Confluence cluster debugging tools

There is an umbrella issue opened for all cluster debugging tools [here](#)

It includes the tools listed below.

Multicast

- Which multicast address?

The multicast address and port used by Confluence can be found on the [Cluster Administration 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.
Multicast Test	Coherence tool for testing multicast traffic from one node to another.
<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 the [Multicast Test](#) tool shows that 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's a known issue with IPv6, especially on Linux.

The fix is to add `-Djava.net.preferIPv4Stack=true` to `JAVA_OPTS`. This tells the JVM to try binding an IPv4 address first, and resort to IPv6 only if that fails.

Note: A more radical approach is to add `NETWORKING_IPV6=no` to `/etc/sysconfig/network`, yet probably should be left for a later consideration on a production machine.

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 `confluence.cfg.xml` in the Confluence home directory and add a property (or change the existing one) to select your desired network interface. For example to tell Confluence to use `eth1`:

```
<property name="confluence.cluster.interface">eth1</property>
```

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.

Create a file in the Confluence home directory called `tangosol-coherence-override.xml`. Add the following to it, setting the TTL value appropriately (1 is the default):

```
<?xml version='1.0'?>
<coherence>
<cluster-config>
  <multicast-listener>
    <time-to-live system-property='tangosol.coherence.ttl'>1</time-to-live>
  </multicast-listener>
</cluster-config>
</coherence>
```

Alternatively, simply start Confluence with the system property: `-Dtangosol.coherence.ttl=1`. Again, 1 is the default value, and you should change it to something appropriate to your network topology.

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.

For best results, use the simplest network topology possible for the cluster traffic between the nodes. For two nodes, that means a single network cable. For larger numbers, try using a single high-quality switch.

Advanced Tangosol configuration

If the solution to your problem involves changes to the Tangosol configuration, these changes should **not** be made to the Confluence configuration in `confluence/WEB-INF/classes/`. Instead, to ensure your configuration survives upgrades, make your changes via:

- Tangosol system properties
- creating a `tangosol-coherence-override.xml` file in the Confluence home directory.

Examples of making these changes are shown in the [increasing the TTL](#) section.

Didn't find a solution?

Check Related Articles from the Confluence Knowledge Base

Content by label

There is no content with the specified labels



Open JIRA Features and Bug Reports

T	Key	Summary	Assignee	Reporter	P	Status	Resolution	Created	Updated
	CONF-27248	WorkBox (notifications and tasks) support for Confluence clustered	Unassigned	Chris Hubing		OPEN	Unresolved	Nov 14, 2012	Mar 25, 2014
	CONF-12689	Support Confluence cluster upgrades without an outage	Unassigned	Igor Minar		OPEN	Unresolved	Aug 06, 2008	Jun 24, 2012
	CONF-9297	Confluence should be able to automatically recover from cluster panics	Unassigned	Gary Weaver		OPEN	Unresolved	Aug 27, 2007	Apr 29, 2014
	CONF-32697	lucene-query-mapper plugin module cannot be found on plugin install in a cluster	Unassigned	Paul Curren [Atlassian]	↑	OPEN	Unresolved	Feb 24, 2014	Feb 28, 2014
	CONF-14657	Retrieving the global settings in a clustered environment causes a lot of contention	Unassigned	Chris Kiehl [Atlassian]	↑	OPEN	Unresolved	Feb 22, 2009	Nov 09, 2009
	CONF-16419	Installing a font for PDF export in a cluster will not carry to cluster nodes that are down or unavailable.	Unassigned	Charles Miller [Atlassian]	↑	OPEN	Unresolved	Jul 21, 2009	Aug 05, 2009
	CONF-32755	Remote API methods that generate URLs (e.g. export space) are not cluster compatible	Unassigned	Paul Curren [Atlassian]	↑	OPEN	Unresolved	Feb 27, 2014	Feb 28, 2014
	CONF-13421	Layout customisations are not propagated to other cluster nodes	Unassigned	Matt Ryall [Atlassian]	↑	OPEN	Unresolved	Oct 16, 2008	Mar 25, 2014
	CONF-17089	Reindexing in cluster only runs on one node if triggered from web UI	Unassigned	Anatoli Kazatchkov [Atlassian]	↑	OPEN	Unresolved	Oct 01, 2009	Mar 25, 2014
	CONF-9324	Lots of ObjectDeletedException's during cluster builds	Unassigned	Matthew Jensen [Atlassian]	↑	OPEN	Unresolved	Aug 29, 2007	Mar 25, 2014
	CONF-14120	Hibernates UpdateTimestampsCache doesn't handle concurrent writes	Unassigned	Chris Kiehl [Atlassian]	↑	OPEN	Unresolved	Jan 06, 2009	Mar 25, 2014
	CONF-32059	User Macro in cluster environment is broken due cache messageBundles	Unassigned	William Zanchet [Atlassian]	↑	OPEN	Unresolved	Dec 20, 2013	Mar 25, 2014

	CONF-17825	Re-indexing under Confluence Cluster is 30% slower than under standalone	Unassigned	George Barnett [Atlassian]		OPEN	Unresolved	Nov 24, 2009	Mar 25, 2014
	CONF-12614	Intermittent ConcurrentModificationException in cluster	Paul Curren [Atlassian]	Anatoli Kazatchkov [Atlassian]		OPEN	Unresolved	Jul 30, 2008	Apr 10, 2014
	CONF-32768	"Create Decision" button does not work in Confluence Cluster version.	Unassigned	Deividi Luvison [Atlassian]		OPEN	Unresolved	Feb 27, 2014	Apr 15, 2014
	CONF-32420	JIRA Issues macro fails because the confluence-coherence-cache-config-clustered.xml file is not being updated in clustered upgrades	Unassigned	William Zanchet [Atlassian]		OPEN	Unresolved	Jan 29, 2014	Mar 11, 2014
	CONF-33205	Slowness when saving page after update the image on the page	Unassigned	Yilin MO [Atlassian]		OPEN	Unresolved	Apr 08, 2014	Apr 28, 2014
	CONF-23033	Viewfile macro does not work in Confluence Clustered when Office Connector is configured to use Cache in Memory for temporary storage	Unassigned	Roy Hartono [Atlassian]		OPEN	Unresolved	Aug 09, 2011	Apr 29, 2014
	CONF-12486	ClassNotFoundException logged on cluster node startup	Unassigned	Anatoli Kazatchkov [Atlassian]		OPEN	Unresolved	Jul 18, 2008	Apr 08, 2014
	CONF-25211	Plugin installation breaks clustered cache	Unassigned	Thomas Krug		OPEN	Unresolved	Apr 14, 2012	Apr 08, 2014

Showing 20 out of 28 issues

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.

Related

Cluster Safety Mechanism

Multicast Test

This page describes the **Multicast Test**, a Coherence tool for testing multicast traffic from one node to another. You may find this useful when troubleshooting a [clustered installation](#) of Confluence.

In order to run the Multicast test, you need to download the [Coherence for Java](#) from Oracle's website. You will need to sign up for a free Oracle account and sign the license agreement, before downloading the file.

The Multicast Test comes as a script called `multicast-test`, which you will find located in the `bin` folder in the above zip file.

Instructions on how to run this script file can be found in the [Coherence documentation](#). You may like to go straight to the subheading called '**Example**' in the guide, where there is an example on how to use the `multicast-test` script.

The Multicast Test will use the multicast address of 237.0.0.1:9000 by default. Confluence creates a unique address based on the cluster name that you enter during setup. As such, you should include the `-group` flag in your multicast testing to ensure your tests are broadcasting across the same address as your Confluence nodes.

The information on this page *does not apply* to Confluence OnDemand.

RELATED TOPICS

[Cluster Troubleshooting](#)

[Confluence Clustering Overview](#)

Clustering for Scalability vs Clustering for High Availability (HA)

Request for interest

Hi there,

Great to see your interest in Confluence Clustered!

We are currently working on a completely new version of our clustered offering that provides greater scalability as well as High Availability.

If you are interested in joining our closed beta program for the new offering, please, send an email to kelvin@atlassian.com with the subject: **Confluence Clustered** and I will get back to you ASAP.

Cheers,

Kelvin Yap
Enterprise Advocate

People occasionally enquire about setting up High-Availability (HA) Confluence clusters. [Confluence's clustering](#) is designed to solve a different problem, that of scaling under high load. This page explains the difference.

What is High Availability (HA)?

HA means that your application *will be* available, without interruption. It's a very difficult thing to achieve, and is typically what people are talking about when they refer to [five-nines availability](#).

In the context of application clustering, it means that any given node (or combination of nodes) can be shut down, blown up, or simply disconnected from the network unexpectedly, and the rest of the cluster will continue operating cleanly as long as at least one node remains. It requires that nodes can be upgraded individually while the rest of the cluster operates, and that no disruption will result when a node rejoins the cluster. It typically also requires that nodes be installed in geographically separate locations.

 Confluence's clustering is not designed to solve this problem, and does not provide high availability.

What does Confluence's clustering do, then?

Confluence's clustering system allows a single installation to serve a much greater number of concurrent requests than a single server. This is what we refer to as 'scaling under load'.

It does provide a certain amount of resilience, as the death of one node won't bring the other(s) down. However, it requires very low network latency, which rules out geographic separation of the servers, and upgrading can only be performed while the entire cluster is shut down. This doesn't mean that Confluence's clustering is buggy or broken. It simply reflects the difference between the two design aims.

On this page:

- [What is High Availability \(HA\)?](#)
- [What does Confluence's clustering do, then?](#)
- [So what kind of resilience can I build into a Confluence installation?](#)
- [What's the difference between load balancing and failover?](#)
- [What do you mean by 'session affinity'?](#)
 - [RELATED TOPICS](#)

 *The information on this page **does not apply** to Confluence OnDemand.*

So what kind of resilience can I build into a Confluence installation?

It's still entirely possible to build a resilient Confluence installation, using a 'cold-failover' approach in which two (or more) servers share a database and (normally) a network-mounted file system, where no more than one

server is actually running at any given time.

Several different approaches are feasible, but the common elements are:

- a well-configured load balancer (session affinity is irrelevant in this case)
- a reliable monitoring system which can detect and shut down a misbehaving Confluence instance before starting the spare server
- startup scripts with added smarts to check for the presence of another running node before deciding whether to start up a server
- servers with the same view of both the database and the home directory.



It's vital to ensure that only one server is running at any one time, in this kind of setup. If a server starts while another is already running against the same database, the result will be a [cluster panic](#) that shuts down both servers.

A single database becomes the single point of failure in such a system. This can be alleviated by database clustering, or by replication from the 'active' database server to the standby server(s) if you wish to separate the failover systems while keeping database latency to a minimum.

In the same vein, the home directory can be hosted on a shared network system — SAN or NAS, preferably with its own replication/rapid recovery system — though there's a [known issue](#) to consider. Alternatively, to avoid the use of networked file systems, a utility such as `rsync` can be used to periodically bring the spare servers' home directories up to date, so long as you keep the period sufficiently short — probably between one and five minutes, depending on the rate of activity. This can be avoided altogether by keeping attachments in the database; it increases the demands on the bandwidth between the application and database servers, but guarantees that the system is in a consistent state at switchover. If the data is at all sensitive or confidential, it's advisable to run `rsync` over `ssh`, to minimise the opportunity for the data to be captured on its way across the network.

What's the difference between load balancing and failover?

Load balancing means that all servers are active, and new requests are distributed among them. Several strategies are available, but the most common are:

- round-robin — the first request goes to the first server, the second request goes to the second server, and so on. When you run out of servers, the next request goes to the first server, and around it goes again.
- percentage-based — if (for example) you have two servers, and one can handle twice the load of the other, you can tell the load balancer to send two requests to the stronger server for every request that goes to the weaker one.
- availability — the load balancer sends a test query to each of the servers every second or so, and directs each new request to the server that's currently responding the fastest.

Failover means that only one server is active at any given time, and normally involves two servers (any number of servers may be involved, depending on the system). If the active one stops responding, requests are directed to the other server — the system 'fails over' to the second one.

'Cold failover' means that the second server is only started up after the first one has been shut down. This is the case for non-clustered Confluence.

'Hot failover' or 'hot standby' means that all servers are running at all times, and that the load is directed entirely toward one server at any one time.

A load balancer can be used in both scenarios, especially if it's smart enough to keep track of which servers are currently running.

Failover can also be managed via DNS, in a sufficiently well-controlled environment.

What do you mean by 'session affinity'?

Sessions consist of several transmissions in each direction between the client (browser) and the server. Session affinity means that the load balancer keeps track of which server received the initial transmission from a given browser, and that it will then send any subsequent requests from that browser to the same server.

This is necessary with Confluence clustering, in particular, because sessions are not shared across cluster nodes. If you log into one node and then send a request to another, the other node will send you the login

screen because it doesn't recognise your session cookie.

RELATED TOPICS

[Confluence Clustering Overview](#)

Recommended network topology

Atlassian recommends a network topology similar to the one shown below, to get the best results from a Confluence Clustered deployment.

The number of Confluence nodes in the deployment is adjustable — select the number which suits your own requirements.

The most important aspect is that cluster, database and HTTP (client) traffic are all carried on separate subnets. It is possible, on a sufficiently fast network, to carry cluster and database traffic on the same subnet but we do strongly recommend that HTTP traffic be always confined to a separate subnet on production deployments.


Confluence Clustered does not support clustered communication over WAN, VLAN or VPN. All Confluence Clustered nodes must be on the same local subnet, ideally networked via an ethernet hub or simple switch. The cluster communication network must also support multicast IP networking.

 *The information on this page **does not apply** to Confluence OnDemand.*



Use this example as a basis for your own network diagram

When you are considering a Confluence Clustered deployment, you should prepare a network diagram like the one on this page. This will facilitate discussion with [Atlassian Support](#) and help with your own planning. Please refer to the [cluster checklist](#) for more guidance on planning your clustered deployment.



Cluster Administration page

Overview

Any instance of Confluence which uses a clustered license has a Cluster Configuration page which includes information about the active cluster.

To open the Cluster Administration page:

1. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
2. Click '**Cluster Configuration**' in the left-hand menu, in the section called 'Clustering'.

 The information on this page *does not apply* to Confluence OnDemand.

On this page:

- [Overview](#)
- [Availability](#)

Related pages:

- [Overview of Confluence Clusters](#)
- [Confluence Cluster Installation](#)
- [Cluster Troubleshooting](#)

Availability

To access this functionality, you must:

- Be a System Administrator (i.e. have global System Administrator [permissions](#)), and
- be using Confluence 2.3 or later, and
- be using a clustered Confluence license.

The 'Cluster Administration' page shows your cluster configuration, and allows you to start a new Confluence cluster using data from this instance.

Cluster Status indicates whether your cluster is currently running.

Licensed nodes is the maximum number of instances of Confluence your license allows in a cluster.

Active nodes lists the instances of Confluence currently participating in the cluster.

Starting a new cluster will perform the following changes:

- enable a clustered cache
- migrate attachments from file system to the database
- publish database connection information so other nodes can join the cluster.

i All access to Confluence will be locked while this takes place, and you will be forced to restart Confluence afterwards.

Cluster name is a short name for identifying your cluster. Other Confluence instances can join the cluster using this name.

i To join an existing cluster, start a clean copy of Confluence on this node and select 'Join Cluster' during the setup wizard.

Cluster Checklist**Request for interest**

Hi there,

Great to see your interest in Confluence Clustered!

We are currently working on a completely new version of our clustered offering that provides greater scalability as well as High Availability.

If you are interested in joining our closed beta program for the new offering, please, send an email to kelvin@atlassian.com with the subject: **Confluence Clustered** and I will get back to you ASAP.

Cheers,

Kelvin Yap
Enterprise Advocate

It is possible to run Confluence in a clustered environment instead of on a single server. This means that you can run multiple copies of Confluence in a cluster, so that clients (such as a browser) can connect to any copy and see the same information.

Refer to the [clustering overview](#) for more information and a list of related pages about clustering Confluence.

Note: Consider your options carefully before deciding on a clustered installation. While we have tried to make clustering Confluence as easy and administrator-friendly as possible, it is a major architectural change and requires extra planning for deployment and upgrades. Please consider the information below and then consult [Atlassian Sales](#) before making your final decision.

Purpose of this Document

The purpose of this cluster checklist is to help you:

- Decide whether Confluence Clustered is the right solution for you.
- Create a plan for your clustered deployment.

If you need to raise a support request with Atlassian during or after cluster deployment, we will need to ask you questions about your configuration. It will save crucial time if you can provide us with your deployment plan.

For more information about clustering Confluence, refer to the [clustering overview](#).

Assumed Knowledge

In writing this document, we have assumed that our readers have an in-depth knowledge of the following technical areas:

- Database
- Networking
- Application servers
- Load balancers

Before starting a clustered deployment please read the information on this page carefully, as well as the linked documentation, to assess if you have the assumed knowledge.

On this page:

- [Purpose of this Document](#)
- [Assumed Knowledge](#)
- [General Considerations](#)
- [Server Setup](#)
- [Database Setup](#)
- [Network Setup](#)
- [Staging Environment](#)

Related pages:

- [Running Confluence in a Virtualised Environment](#)
- [Confluence Cluster Installation](#)
- [Confluence Clustering Overview](#)
- [Recommended network topology](#)
- [Apache and Tomcat load balancing](#)
- [Confluence Administrator's Guide](#)

 *The information on this page **does not** apply to Confluence OnDemand.*

General Considerations



What will Confluence Clustered do for you?

The points in this section of the page will help you evaluate your reasons for considering a clustered deployment, and then decide whether Confluence Clustered is the right solution for your environment.

Confluence Clustered is designed to scale the number of simultaneously connected users at a much better performance than what a single node can achieve

Confluence Clustered will not improve performance in systems with few users.

Clustering Confluence means that user requests can be served by independent machines. The performance gains are substantial, and have improved a lot further since Confluence 3.0. Clustering is especially great in dealing with spikes to the load, e.g. during certain hours of business. Just note that if rendering a complicated page (e.g. containing many macros or rendering many graphs) takes five seconds on an otherwise idle server, it will not be faster in a clustered environment. Also, the first step when you encounter performance issues is to tune your existing system, make sure you are using the right hardware and have looked at your database.

Confluence Clustered is not a high availability solution.

Confluence Clustered is not designed specifically to provide a [high availability](#) solution.

General availability is higher in a Confluence cluster than on a single installation, you can for example take one node down for minor maintenance tasks e.g. when adding a new CPU or adding RAM. But you still have to bring down all nodes at the same time for software upgrades. Also there are certain conditions, like loss of network connectivity between nodes ('split brain'), that will result in the cluster shutting itself down. Confluence Clustered offers higher reliability, but not high availability.

Confluence Clustered is not for disaster recovery nor for transparent failover.

If one node crashes, there is no transparent failover for the connected client. Also, our network requirements (see [below](#)) make Confluence unsuitable for deployment to different cities or even to different buildings.

Server Setup

The number of supported cluster nodes is limited to four.

 **Not supported.** In theory, you can connect more than four nodes — but that is not covered by Atlassian Support.

All cluster nodes must have the same version of OS, application server, etc.

Confluence requires a homogeneous environment. All Confluence cluster nodes must have the same version of the following:

- Operating system
- CPU
- Installed memory
- Java
- Application server

 Note that 'same version' means '**same to the last digit**'. For example, Java v1.4.2_16 is not the same as v1.4.2_15

 We strongly recommend user to have the same memory configuration (both the JVM and the physical memory) because a cluster uses a replicated cache. A replicated cache requires the same amount of memory on each node in the operating cluster. The memory allocations must be equal.

Use good and up-to-date hardware.

While the details are up to you, we strongly suggest that your servers have at least 4GB of physical RAM. A high number of concurrent users means that a lot of RAM will be consumed. You usually don't need to assign more than 4GB per JVM process, and most of the time even just 1GB or 2GB will be fine, you should just be prepared to fine tune the settings.

Confluence should be the only application on the cluster servers.

No additional applications (other than core operating system services) should be running on the same servers as Confluence.

Since your goal should be increased capacity and performance, you should not risk this by running any other process on the machine with a Confluence Clustered node. While it may be fine to run JIRA, Confluence and Bamboo on a dedicated Atlassian software server for small installations, it is strongly discouraged for clustering Confluence.

Do not upgrade and switch to Confluence Clustered at the same time

If you plan to migrate to a clustered solution, make sure you are migrating within the same version of Confluence. If you plan to upgrade to a higher version of Confluence, do this **before** the migration to the clustered version. For example, if you are currently running Confluence 2.9.2, and want to roll out the clustered version of Confluence 3.0, you must first upgrade to Confluence 3.0 and check that everything works fine (e.g. by running and monitoring your production system for a week). Then you are in a good position to migrate to the clustered version.

Database Setup

Run the database on its own physical server.

You are optimising for performance, so you don't want the database to slow down your application servers, or vice versa. In high load scenarios, the database may need to have better hardware than the application servers to be able to handle all requests. You should find out by performing loadtesting.

Attachments must be stored in a database and not the local file system

Storing attachments in the database is the only supported [attachment storage configuration](#) for clustering Confluence.

Make sure that you use a supported version of a database server to store Confluence's data.

Please check that your intended database is [officially supported](#) by Atlassian Confluence. The load on an average cluster solution is higher than on a single box installation, and it is therefore even more crucial to use the right database vendor and version.

Your database must be provisioned to store a large volume of binary data.

Note that Confluence clustered stores file attachments in the database, and you need an experienced DBA who can monitor and manage the data growth.

You need an experienced DBA available to troubleshoot database performance issues.

Not having an experienced full-time DBA at hand at short notice when entering the realm of high load is dangerous. While small installations of Confluence basically work 'out of the box', anything that involves high load and a lot of database space requires continual monitoring, optimising and fine tuning of the Confluence database. When we ramp up the load on our loadtesting environment, we see that database usage goes up as well. Having powerful hardware in place helps, but if there are queries that become inefficient with your particular load pattern, you need an expert to tune it. As an example, we have seen PostgreSQL switch its internal caching mechanism when a particular table reached a certain size, which resulted in a drop of performance by about 200ms per request. This happened from one second to the other. Being able to troubleshoot and then fix issues like these is important in any enterprise system, but it is even more in a high load scenario.

Network Setup

We recommend hardware load balancers or putting a software loadbalancer onto its on server.

If you use a software load balancer (which is fine except for really extreme installations), it must be deployed on a machine of its own. Running a software load balancer on a cluster node is not supported. If a node unexpectedly got overwhelmed by a spike in load, a load balancer on that node would turn unresponsive. As a result, your whole cluster would be inaccessible even though the other nodes would be available. So using a different server is common practice and common sense.

Use separate network adapters for communication between servers.

The Confluence cluster nodes should have a separate physical network (i.e. separate NICs) for inter-server communication.

This is the best way of getting the cluster to run fast and reliably. Performance problems are likely to occur if you connect cluster nodes via a network that has lots of other data streaming through it.

The switch connecting the Confluence cluster nodes must not be a 'smart switch'.

 **Not supported.** Smart switches are not covered by Atlassian Support for Confluence Clustered.

Do not use smart switches between cluster nodes. Many problems have been reported and attributed to smart switches. They have a tendency to interrupt broadcast or multicast traffic, thus reliably killing a cluster after a certain amount of time has passed. This makes troubleshooting especially complex and tedious.

Cisco switches need additional configuration.

If the switch connecting the Confluence cluster nodes is a Cisco switch then it might need [additional configuration](#) to support Confluence clustering.

Please make sure you find out all the details about your switches before you start the deployment.

It is recommended that the database is on a different physical network from the Confluence server nodes.

Since you want to increase your capacity and performance for high loads, it is recommended to have your database on a different network. Please refer to the [recommended topology diagram](#) for more information.

Minimize the latency between the Confluence cluster nodes and the database.

Even though having the nodes and the database on the same physical network usually suffices, you should take the time to explicitly measure network latency, and make sure it is as close to zero as possible.

Prepare a network diagram.

To facilitate discussion and to ease planning, you should prepare a network diagram like [this example of recommended network topology](#).

If you request [support](#) with Confluence Clustered, we may ask for your network diagram. We recommend that you create one similar to our example before you proceed with the installation.

You need network support staff available to troubleshoot cluster communication issues.

Setting up a cluster is not trivial. Even small problems in network design will be expanded in a clustered installation. (This is true of any kind of software.)

It is absolutely vital that you have dedicated network staff available to track down problems when they arise. A cluster will usually be used by thousands of users, and you don't want to keep them waiting because a network card breaks, or because someone made an undocumented change to the network and you don't have an expert around who can figure it out.

Staging Environment

You need a staging environment that is exactly the same as your production system.

You must be able to test drive any change to the cluster (installing upgrades, installing plugins) and to perform other tests (checking connectivity, debugging problems) on a staging cluster.

The staging environment must be:

- On the same OS, database, and Java version as your production environment.
- Clustered.

If you require support, we may for example ask you to turn off certain third-party plugins. If you can't do this in your production environment and you don't have a staging environment for troubleshooting, we may not be able to help you.

Performance Tuning

This document describes tuning your application for improved performance. It is not a guide to troubleshooting Confluence outages. Check [Troubleshooting Confluence Hanging or Crashing](#) for help if Confluence is crashing.

Description

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. As an example, our internal performance tests show a **20% speed-up** (when viewing pages under load) between Tomcat 6 on Java 6 vs Tomcat 5.5 on Java 5 **out of the box**.

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.

On this page:

- [Description](#)
- [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 indexes](#)
- [Database statistics and query analysers](#)
- [Cache tuning in Confluence and Apache](#)
- [Antivirus software](#)
- [Enabling HTTP compression](#)
- [Virtual operating systems](#)
- [Performance testing](#)
- [Access logs](#)
- [Built-in profiler](#)
- [Application server memory settings](#)
- [Web server configuration](#)
- [Parallel GC](#)
- [Troubleshooting possible memory leaks](#)
- [Plugins](#)

 *The information on this page **does not** apply to Confluence OnDemand.*

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 database** that is provided with Confluence is meant only to be used for evaluation, not for production Confluence sites. After the evaluation finishes, you will certainly need to switch to an external relational database management system. Beyond this, we do not recommend any particular RDBMS over another. 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

Database Connection	Transaction Isolation	Read committed
Database Latency	0	
Confluence Usage		

console.

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 indexes

Especially if you have more than a few thousand active users, and all most obvious measures have been tried out but the database still seems to be under high load, you should consider engaging a database administrator (DBA) to tune the database specifically to the demands that your particular Confluence installation is placing on it. If you do not have a full-time DBA and can't even get one for temporary consulting, you may want to consult the [database indexing advice](#) that we have been gathering from customer reports and our own experience running and developing Confluence. The instructions on that page are for Oracle, but most of the indexes can be applied to (and will help with) any database.

(These database indexes are now created automatically when Confluence is installed, but existing installations upgrading to a more recent version may still need to add them manually)

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.

Virtual operating systems

Virtual Environments such as VMWare can cause Confluence CPU to spike. Run Confluence on a native OS. Refer to the list of supported operating systems for Confluence in the [Supported Platforms](#) topic.

Note: In some situation the VMTools can crash, cause a excessive context switches and interrupts causing the JVM to run slowly and Confluence to start up very slowly.

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.

Parallel GC

If you have multiple CPU's on your server, you can add `-XX:+UseParallelOldGC` to your `JAVA_OPTS` options. This will allow garbage collection of the Tenured Space to happen in parallel with the application and can boost performance and can reduce slow performance spikes. For more information, please refer to our detailed page on [Garbage Collector Performance Issues](#), and Sun's [summary](#) of collectors.

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.

Plugins

Some 3rd-party plugins were not written to scale to large enterprises' needs.

Confluence has been optimized to work under high load and with many pages. Some 3rd party plugins however have been written with small size companies in mind, and can't cope with large numbers of concurrent users, or large numbers of pages and permissions, or large numbers of spaces. It is impossible to tell which ones will fail under which conditions, but it will always help to turn off 3rd-party plugins that are not strictly mission-critical while investigating performance issues.

RELATED TOPICS

[Garbage Collector Performance Issues](#)
[Cache Performance Tuning](#)
[Cache Performance Tuning for Specific Problems](#)
[Working with Confluence Logs](#)
[Operating Large or Mission-Critical Confluence Installations](#)
[Confluence Clustering Overview](#)
[Requesting Performance Support](#)
[Confluence Administrator's Guide](#)
[Configuring Confluence](#)

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 the majority of the cases, cache size is the parameter you would want to change. Fortunately, from Confluence 3.0, it is very easy to [adjust cache sizes](#) through the Administration Console. However, if you need to modify parameters other than a cache size, you would need to modify the relevant configuration files manually.

The cache performance information for your Confluence installation is available under **Administration > Cache Statistics**. For more information about the numbers displayed on that screen, see [Cache Statistics](#).

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](#).
- If you only need to modify Confluence's maximum cache sizes, you can do this through the [Cache Statistics](#) feature of the Administration Console.

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 caches above are of size 1000 (meaning that it can contain up to 1000 objects), which is the default size for caches in the default cache scheme. Refer to [Confluence Cache Schemes](#) for more explanation.

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.

i 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.

On this page:

- [Cache tuning example](#)
- [Finding the configuration file](#)
- [Cache key mappings](#)
- [Standard editions of Confluence](#)
- [Clustered editions of Confluence](#)
- [Reference of Internal names to Human readable names](#)
- [Important caches](#)
- [Cache tuning follow-up](#)
- [Notes](#)

Related pages:

- [Cache Performance Tuning for Specific Problems](#)
- [Confluence Cache Schemes](#)
- [Working with Confluence Logs](#)
- [Operating Large or Mission-Critical Confluence Installations](#)
- [Confluence Clustering Overview](#)
- [Requesting Performance Support](#)
- [Confluence Administrator's Guide](#)
- [Configuring Confluence](#)

 *The information on this page does not apply to Confluence OnDemand.*

Finding the configuration file


The caches are configured in `ehcache.xml` (for standard editions) or `confluence-coherence-cache-config-clustered.xml` (for clustered editions) which is stored in `<confluence-home>/config/`.

i Oracle Coherence Licensing Change:

- Due to a license agreement change, Confluence is now available in two editions:
 - **Standard Edition** — Confluence with Ehcache's caching technology (available to customers with non-clustered Confluence licenses).
 -  **If you are currently running a clustered installation of Confluence, please do not upgrade it with a standard edition of Confluence.**
 - **Clustered Edition** — Confluence with Oracle's Coherence clustering and distributed caching technology (available to customers with Confluence clustered licenses only).
- For more information about these changes, please refer to the [Coherence License Changes](#) document.
- If you have a Confluence clustered license, are running a clustered installation of Confluence and wish to upgrade to Confluence version 2.6 or later, please ensure that you download only a [clustered edition of Confluence](#).

Cache key mappings

The cache configuration file configures caches by their keys. When you move your mouse over the the cache names displayed on the cache statistics page, a tooltip will indicate the actual cache key for that cache name.



Using [our example](#) from the table above, if we were to modify parameters for the Users cache we would need to change the cache with the key `com.atlassian.user.impl.hibernate.DefaultHibernateUser`. Do not get confused with Users (External Mappings) and Users (External Groups) which are in themselves, two separate caches. "Users" is the friendly name for `com.atlassian.user.impl.hibernate.DefaultHibernateUser`.

Standard editions of Confluence

In standard editions of Confluence, the caching layer is [Ehcache](#).

Understanding the Ehcache configuration file

For more information about the Ehcache configuration file and a full reference on Ehcache configuration, please refer to the [Ehcache configuration documentation](#).

Converting your Coherence configuration to Ehcache

 This section only applies to customers who:

- Have an installation of Confluence that was downloaded before the 4th of September 2009.
- Intend to (or have already) upgraded to Confluence 3.0.1 or later (or to Confluence versions 2.6.3, 2.7.4, 2.8.3, 2.9.3 and 2.10.4).
- Will use a non-clustered Confluence license for the Confluence upgrade.
- Have implemented customisations to their Confluence installation's cache configuration file (`confluence-coherence-cache-config.xml`).

To maintain your existing cache configuration file settings, you will need to transfer any cache customisations you have implemented in the Coherence cache configuration file (`confluence-coherence-cache-config.xml`) to the relevant entries in the Ehcache cache configuration file (`ehcache.xml`).

Each cache has a cache-mapping element in the Coherence file (of which there is an equivalent cache element in the `ehcache.xml` file). Unfortunately, copying across your customisations is not quite a straightforward process because the Coherence file defines several 'caching schemes' to store the actual cache values, which in turn are referenced by the cache-mapping elements. In contrast, the `ehcache.xml` file does not support caching schemes and a cache's values are expressed explicitly in separate parameters of a cache element.

To convert your Coherence cache configuration file customisations across to the equivalent Ehcache file:

1. Open both the `confluence-coherence-cache-config.xml` and `ehcache.xml` files in a text editor. These files are located in the `<confluence-home>/config` directory.
 -  If you implemented your customisations in a version of Confluence prior to 3.0, you will most likely find the `confluence-coherence-cache-config.xml` file in the `<confluence-install>/confluence/WEB-INF/classes` directory.
2. In the customised `confluence-coherence-cache-config.xml` file:
 - a. Identify the caching schemes that were customised in this file and make a note of the values of all its child elements.
 -  Typically, each caching scheme is located inside a `local-scheme` element and all of these are enclosed within the `cache-schemes` element, which appears towards the end of this file.
 - b. Note each customised caching scheme by the content of its `scheme-name` element.
 - c. For each cache-mapping element (which typically appears towards the top of this file), identify if it has a `scheme-name` element whose content matches one noted in the previous step and if so,

make a note of its associated cache-name element.

3. In the ehcache.xml file:

- a. Identify each cache element whose 'name' parameter matches the cache-name elements noted in step '2c'.
- b. Using the mappings table below, apply the values noted in step '2a' to the appropriate parameters of the cache elements identified in the previous step ('3a').

Mappings table showing how elements of the Coherence cache configuration file map to parameters of the equivalent Ehcache file.

Coherence Element	Ehcache Attribute
high-units	maxElementsInMemory
expiry-delay > 0s	timeToIdleSeconds - Use this attribute for expiry delays greater than 0s along with the eternal attribute set to 'false'
expiry-delay = 0s	eternal - For expiry delays of 0s, set this attribute to 'true'.

Clustered editions of Confluence

Understanding the Coherence configuration file

The Coherence configuration file is a mapping of *cache keys* to *cache schemes*. Each cache scheme controls the expiry, eviction policy and size of the caches linked to it. A cache scheme can extend another scheme.

For a full reference, see the [Oracle's Coherence cache configuration documentation](#).

Defining caching scheme mappings in Coherence cache config file

If a cache key does not have an explicit definition in the caching scheme mappings (defined in confluence-coherence-cache-config.xml) then it will use the "default" cache-mapping.

In our example, com.atlassian.user.impl.hibernate.DefaultHibernateUser is not explicitly defined in the caching scheme mappings. Hence to increase the expiry-delay to 2 hours, we will need to define the mapping ourselves and add the following within the <caching-scheme-mapping>...</caching-scheme-mapping> tags:

```
<cache-mapping>
<cache-name>com.atlassian.user.impl.hibernate.DefaultHibernateUser</cache-name>
<scheme-name>cache:com.atlassian.user.impl.hibernate.DefaultHibernateUser</scheme-name>
</cache-mapping>
```

Then we will need to define a cache schema with name cache:com.atlassian.user.impl.hibernate.DefaultHibernateUser within <caching-schemes>...</caching-schemes> tags.

```
<local-scheme>
<scheme-name>cache:com.atlassian.user.impl.hibernate.DefaultHibernateUser</scheme-name>
<scheme-ref>default</scheme-ref>
<high-units>10000</high-units>
<expiry-delay>7200</expiry-delay>
</local-scheme>
```

It's possible to define a local-scheme mapping for a cache key without defining certain parameters (e.g. <high-units>). In such a cases, their parameters will be inherited from scheme-ref scheme, which is the default

scheme in our case.

Reference of Internal names to Human readable names

The names in the Cache statistics screen are mapped to internal names (as per the ehcache/coherence-override file) as follows:

bucket.user.persistence.dao.hibernate.BucketUserDAO.findUserByUsername()	Users (Username)
bucket.user.propertyset.BucketPropertySetItem	Object Properties
bucket.user.providers.CachingAccessProvider.handles()	Groups (OSUser)
bucket.user.providers.CachingAccessProvider.inGroup()	User Group Mappings (OSUser)
bucket.user.providers.CachingCredentialsProvider	Users (OSUser Credentials)
com.atlassian.bandana.BandanaPersister	Settings (Persistence)
com.atlassian.confluence.core.BodyContent	Content Bodies
com.atlassian.confluence.core.ContentEntityObject	Content Objects
com.atlassian.confluence.core.ContentEntityObject.attachments	Content Attachments
com.atlassian.confluence.core.ContentEntityObject.bodyContents	Content Body Mappings
com.atlassian.confluence.core.ContentEntityObject.labelings	Content Label Mappings
com.atlassian.confluence.core.ContentEntityObject.outgoingLinks	Content Links (Outgoing)
com.atlassian.confluence.core.ContentEntityObject.permissions	Content Permission Mappings
com.atlassian.confluence.core.ContentEntityObject.previousVersions	Content Versions
com.atlassian.confluence.core.ContentEntityObject.referralLinks	Content Links (Referral)
com.atlassian.confluence.core.ContentEntityObject.trackbackLinks	Content Links (Trackback)
com.atlassian.confluence.diffs	Page Diffs
com.atlassian.confluence.html.diffs	Html Page Diffs
com.atlassian.confluence.plugins.like.notifications.dao.NotificationDao	Likes Notification DAO
com.atlassian.confluence.security.ContentPermission	Content Permissions
com.atlassian.confluence.core.PersistentDecorator	Layouts (Database)
com.atlassian.confluence.labels.Label	Labels
com.atlassian.confluence.labels.Labelling	Label Content Mappings
com.atlassian.confluence.pages.Attachment.labelings	"Attachment Label Mappings"

com.atlassian.confluence.pages.AttachmentDownloadPathCache	Attachment Download Paths
com.atlassian.confluence.pages.templates.PageTemplate.labellings	"Page Template Label Mappings"
com.atlassian.confluence.links.ReferralLink	Links (External)
com.atlassian.confluence.links.TrackbackLink	Links (Trackback)
com.atlassian.confluence.core.ContentEntityObject.comments	Comments
com.atlassian.confluence.pages.Attachment.previousVersions	Attachment Versions
com.atlassian.confluence.pages.Comment.children	Comment Relationships
com.atlassian.confluence.pages.Draft	Drafts
com.atlassian.confluence.pages.Page.ancestors	Page Ancestors
com.atlassian.confluence.pages.Page.children	Page Children
com.atlassian.confluence.pages.templates.PageTemplate.previousVersions	Template Versions
com.atlassian.confluence.pages.attachments.ImageDetailsDto	Image Details
com.atlassian.confluence.security.SpacePermission	Space Permissions (by ID)
com.atlassian.confluence.setup.bandana.ConfluenceBandanaRecord	Settings
com.atlassian.confluence.spaces.Space	Spaces
com.atlassian.confluence.user.persistence.dao.CachingPersonalInformationDao.usernameTold	User Information By Username
com.atlassian.confluence.util.velocity.ConfluenceVelocityResourceCache	UI Templates
com.atlassian.user.impl.hibernate.DefaultHibernateExternalEntity	Users (External Mappings)
com.atlassian.user.impl.hibernate.DefaultHibernateExternalEntity.groups	Users (External Groups)
com.atlassian.user.impl.hibernate.DefaultHibernateGroup	Groups
com.atlassian.user.impl.hibernate.DefaultHibernateGroup.externalMembers	Groups (External Members)
com.atlassian.user.impl.hibernate.DefaultHibernateGroup.localMembers	Groups (Local Members)
com.atlassian.user.impl.hibernate.DefaultHibernateUser	Users
com.atlassian.user.impl.hibernate.DefaultHibernateUser.groups	User Group Mappings
com.atlassian.user.impl.hibernate.CachingExternalEntityDAO.externalEntityName	Users (External Mappings)
com.opensymphony.user.provider.hibernate.impl.HibernateGroupImpl	Groups (OSUser)

com.opensymphony.user.provider.hibernate.impl.HibernateUserImpl	Users (OSUser)
com.opensymphony.user.provider.hibernate.impl.HibernateUserImpl.groups	User Group Mappings (OSUser Hibernate)
net.sf.hibernate.cache.StandardQueryCache	Database Queries
net.sf.hibernate.cache.UpdateTimestampsCache	Object Timestamps
com.atlassian.confluence.lock-cache	Locks
com.atlassian.confluence.rpc.auth.TokenAuthenticationManager.tokens	Remote Auth Tokens
bucket.user.providers.CachingProfileProvider.getPropertySet()	Bucket Property Set
bucket.user.providers.CachingProfileProvider.handle()	Profile Providers Handles
com.atlassian.confluence.cluster.safety.DefaultClusterSafetyManager.safetyNumber	Cluster Safety Numbers
com.atlassian.confluence.security.PermissionCheckDispatcher.isPermitted()	User Authorized URLs
com.atlassian.confluence.security.persistence.dao.hibernate.legacy.HibernateKey	Hibernate Keys
com.atlassian.confluence.security.trust.ConfluenceTrustedApplication	Trusted Applications
com.atlassian.confluence.security.trust.ConfluenceTrustedApplication.restrictions	Trusted Application Restrictions (Foreign Keys)
com.atlassian.confluence.security.trust.TrustedApplicationRestriction	Trusted Application Restrictions (Objects)
com.atlassian.confluence.themes.persistence.hibernate.DefaultPersistentDecoratorDao	Decorators
com.atlassian.confluence.util.i18n.I18NBeanFactory.by.locale	Internationalisation Bean Factories
com.atlassian.confluence.core.CachingInheritedContentPermissionManager.getInheritedContentPermissionSets()	Inherited Content Permissions
com.atlassian.confluence.pages.persistence.dao.PageDao.getPage()	Pages
com.atlassian.confluence.security.CachingSpacePermissionManager.permissions	Space Permissions (by Type, Scope & Entity)
com.atlassian.confluence.spaces.persistence.dao.SpaceDao.getSpace()	Spaces (by key)
com.atlassian.confluence.util.UserChecker	Number Of Registered Users
com.atlassian.confluence.cache.jcaptcha.ConfluenceCachingCaptchaStore	Captchas
com.atlassian.confluence.core.DefaultContentPropertyManager	Content Properties
com.atlassian.confluence.spaces.SpaceGroup	Space Groups

com.atlassian.confluence.hosted.SpaceGroupPermission	Space Group Permissions
com.atlassian.confluence.spaces.persistence.dao.SpaceGroupDao.getSpaceGroup()	Space Groups (Hibernate)
com.atlassian.confluence.core.ContentEntityObject.contentPermissionSets	Permission Set Collections in Content Entity Objects
com.atlassian.confluence.security.ContentPermissionSet	Content Permission Sets
com.atlassian.confluence.security.ContentPermissionSet.contentPermissions	Permissions in Content Permission Sets
com.atlassian.confluence.published-cache	Objects Published to All Cluster Members
com.atlassian.confluence.core.DefaultHeartbeatManager.activities	Page Edit Activities for Heartbeat Tracking
com.atlassian.confluence.pages.Attachment	Attachments
com.atlassian.confluence.pages.attachments.AttachmentCache	Attachment IDs
com.atlassian.confluence.security.persistence.dao.hibernate.AliasedKey	Encryption Keys
com.atlassian.user.impl.hibernate.properties.HibernatePropertySetFactory.propertysets	Hibernate User Properties
com.atlassian.confluence.follow.Connection	Connection
com.atlassian.confluence.user.DefaultUserAccessor.deactivatedUsers	Disabled Users
com.atlassian.confluence.links.DefaultReferralManager.hotReferrers	Hot Referrers
com.atlassian.confluence.extra.jira.OldRssMacro	Old Rss Macro
com.atlassian.confluence.security.login.DefaultLoginManager	Login Manager: Login attempts for unknown users
com.atlassian.confluence.user.persistence.dao.ConfluenceRememberMeToken	RememberMe Tokens
com.atlassian.confluence.locale.requestLang	Browser language cache
com.atlassian.confluence.security.persistence.dao.hibernate.UserLoginInfo	User Login Information
com.atlassian.confluence.like.LikeEntity	Likes
com.atlassian.crowd.integration-groupnames	Crowd Group Names
com.atlassian.crowd.integration-user	Crowd Users
com.atlassian.crowd.integration-group-membership	Crowd Group Memberships
com.atlassian.crowd.integration-all-group-members	Crowd All Group Members
com.atlassian.crowd.integration-groupname-case	Crowd Group Name Case
com.atlassian.crowd.integration-all-memberships	Crowd All Memberships
com.atlassian.crowd.integration-username-case	Crowd User Name Case
com.atlassian.crowd.integration-parentgroup	Crowd Parent Groups

com.atlassian.crowd.integration-usernames	Crowd User Names
com.atlassian.crowd.integration-group	Crowd Groups
com.atlassian.crowd.integration-is-user-or-group	Crowd User Or Group Cache
com.atlassian.crowd.integration-user-with-attributes	Crowd users with Attributes
1. Embedded Crowd cache friendly names	
com.atlassian.crowd.model.user.InternalUserAttribute	Embedded Crowd Internal User Attribute
com.atlassian.crowd.model.user.InternalUser	Embedded Crowd Internal User
com.atlassian.crowd.model.application.ApplicationImpl.directoryMappings	Embedded Crowd Application Directory Mappings
com.atlassian.crowd.model.directory.DirectoryImpl.attributes	Embedded Crowd Directory Attributes
com.atlassian.crowd.model.directory.DirectoryImpl.allowedOperations	Embedded Crowd Directory Allowed Operations
com.atlassian.crowd.model.application.ApplicationImpl	Embedded Crowd Application
com.atlassian.crowd.model.directory.DirectoryImpl	Embedded Crowd Directory
com.atlassian.crowd.model.application.DirectoryMapping	Embedded Crowd Directory Mapping
com.atlassian.crowd.model.group.InternalGroup	Embedded Crowd Internal Group
com.atlassian.crowd.embedded.hibernate2.HibernateMembership	Embedded Crowd Group Membership
com.atlassian.crowd.model.application.DirectoryMapping.allowedOperations	Embedded Crowd Directory Mapping Allowed Operations
com.atlassian.crowd.model.user.InternalUser.credentialRecords	Embedded Crowd Internal User Credential Records
com.atlassian.crowd.model.application.ApplicationImpl.attributes	Embedded Crowd Application Attributes
com.atlassian.crowd.model.application.DirectoryMapping.authorisedGroups	Embedded Crowd Directory Mapping Authorised Groups
com.atlassian.crowd.model.application.ApplicationImpl.remoteAddresses	Embedded Crowd Application Remote Addresses
com.atlassian.crowd.model.user.InternalUserCredentialRecord	Embedded Crowd Internal User Credential Record
com.atlassian.crowd.model.application.GroupMapping	Embedded Crowd Group Mapping
com.atlassian.crowd.model.group.InternalGroupAttribute	Embedded Crowd Internal Group Attribute
com.atlassian.confluence.user.crowd.DefaultApplicationCache	Embedded Crowd Immutable Application
com.atlassian.confluence.user.crowd.CachedCrowdUserDao.USER_CACHE	Embedded Crowd Users
com.atlassian.confluence.user.crowd.CachedCrowdUserDao.ATTRIBUTE_CACHE	Embedded Crowd User Attributes

com.atlassian.confluence.user.crowd.CachedCrowdGroupDao.GROUP_CACHE	Embedded Crowd Groups
com.atlassian.confluence.user.crowd.CachedCrowdGroupDao.ATTRIBUTE_CACHE	Embedded Crowd Group Attributes
com.atlassian.confluence.user.crowd.CachedCrowdMembershipDao.STRING_PARENT_CACHE	Embedded Crowd String Parent Memberships
com.atlassian.confluence.user.crowd.CachedCrowdMembershipDao.GROUP_PARENT_CACHE	Embedded Crowd Group Object Parent Memberships
com.atlassian.confluence.user.crowd.CachedCrowdMembershipDao.GROUP_CHILD_CACHE	Embedded Crowd Group Object Child Memberships
confluence.fifo.buffer.mail	Mail Queue (only flushable individually)
confluence.fifo.buffer.task	Task Queue (only flushable individually)
confluence.fifo.buffer.mail-error	Mail Error Queue (only flushable individually)
com.atlassian.user.impl.hibernate.HibernateGroupManager.hibernateRepository.groups	Hibernate Groups
com.atlassian.user.impl.hibernate.HibernateGroupManager.hibernateRepository.groups_getGroupsForUser	Hibernate User Groups (Group Side)
com.atlassian.user.impl.hibernate.HibernateGroupManager.hibernateRepository.groups_hasMembership	Hibernate Membership
com.atlassian.user.impl.hibernate.HibernateGroupManager.hibernateRepository.repositories	Hibernate Group Repository
com.atlassian.user.impl.hibernate.HibernateUserManager.hibernateRepository.users	Hibernate Users
com.atlassian.user.impl.hibernate.HibernateUserManager.hibernateRepository.groups_getGroupsForUser	Hibernate User Groups (User Side)
com.atlassian.user.impl.hibernate.HibernateUserManager.hibernateRepository.repository	Hibernate User Repository
com.atlassian.user.impl.hibernate.HibernateUserManager.hibernateRepository.users_ro	Hibernate User Read-Only Flags
com.atlassian.user.impl.hibernate.HibernateUserManager.IdapRepository.users	Hibernate-LDAP Users
com.atlassian.user.impl.hibernate.HibernateUserManager.IdapRepository.groups_getGroupsForUser	Hibernate-LDAP Groups for Users (User Side)
com.atlassian.user.impl.Idap.LDAPGroupManagerReadOnly.IdapRepository.groups	LDAP Groups
com.atlassian.user.impl.Idap.LDAPGroupManagerReadOnly.IdapRepository.groups_getGroupsForUser	LDAP User Groups (Group Side)
com.atlassian.user.impl.Idap.LDAPGroupManagerReadOnly.IdapRepository.groups_hasMembership	LDAP Membership
com.atlassian.user.impl.Idap.LDAPGroupManagerReadOnly.IdapRepository.repositories	LDAP Group Repository
com.atlassian.user.impl.Idap.LDAPUserManagerReadOnly.IdapRepository.users	LDAP Users

<code>com.atlassian.user.impl.Ldap.LDAPUserManagerReadOnly.LdapRepository.groups_getGroupsForUser</code>	LDAP User Groups (User Side)
<code>com.atlassian.user.impl.Ldap.LDAPUserManagerReadOnly.LdapRepository.repository</code>	LDAP User Repository
<code>com.atlassian.user.impl.Ldap.LDAPUserManagerReadOnly.LdapRepository.users_ro</code>	LDAP User Read-Only Flags
<code>com.atlassian.crowd.embedded.atlassianuser.EmbeddedCrowdGroupManager.embeddedCrowd.groups</code>	Embedded Crowd Groups
<code>com.atlassian.crowd.embedded.atlassianuser.EmbeddedCrowdGroupManager.embeddedCrowd.groups_getGroupsForUser</code>	Embedded Crowd User Groups (Group Side)
<code>com.atlassian.crowd.embedded.atlassianuser.EmbeddedCrowdGroupManager.embeddedCrowd.groups_hasMembership</code>	Embedded Crowd Membership
<code>com.atlassian.crowd.embedded.atlassianuser.EmbeddedCrowdGroupManager.embeddedCrowd.repositories</code>	Embedded Crowd Group Repository
<code>com.atlassian.crowd.embedded.atlassianuser.EmbeddedCrowdUserManager.embeddedCrowd.users</code>	Embedded Crowd Users
<code>com.atlassian.crowd.embedded.atlassianuser.EmbeddedCrowdUserManager.embeddedCrowd.groups_getGroupsForUser</code>	Embedded Crowd User Groups (User Side)
<code>com.atlassian.crowd.embedded.atlassianuser.EmbeddedCrowdUserManager.embeddedCrowd.repository</code>	Embedded Crowd User Repository
<code>com.atlassian.crowd.embedded.atlassianuser.EmbeddedCrowdUserManager.embeddedCrowd.users_ro</code>	Embedded Crowd User Read-Only Flags
<code>com.atlassian.crowd.embedded.atlassianuser.EmbeddedCrowdPropertySetFactory.propertysets</code>	Embedded Crowd Properties
<code>com.atlassian.confluence.schedule.ScheduledJobStatus</code>	Scheduled Job Status

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.

- **`com.atlassian.confluence.core.ContentEntityObject`** (known as **Content Objects** cache) 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`.
- **`com.atlassian.confluence.core.ContentEntityObject.bodyContents`** (known as **Content Body Mappings** cache) 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`.
- **`com.atlassian.confluence.security.PermissionCheckDispatcher.isPermitted()`** (known as **User Authorized URLs** cache) should be set to at least the number of concurrent users you expect to access Confluence at the same time

- **com.atlassian.crowd.model.user.InternalUser** (known as **Embedded Crowd Internal User** cache) 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';
```

- **com.atlassian.confluence.user.crowd.CachedCrowdUserDao.USER_CACHE** (known as the **Embedded Crowd Users** cache) should be set to the number of rows in the `cwd_user` table.

```
SELECT
COUNT(*)
FROM
cwd_user u;
```

- **com.atlassian.confluence.security.SpacePermission** (known as **Space Permissions (by ID)** cache) 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`.

Cache tuning follow-up

After you have made changes to your cache config, doing a follow up on the changes in the next week or after the expected performance spike would be important.

Make sure that you take a screenshot of the cache statistics before and after the change. Then compare them with the cache statistics in the later period where performance improvement is expected.

Notes

You can monitor what's in the cache by using a JSP included in the Confluence distribution. Browse to `<base-URL>/admin/cachecontents.jsp` to monitor the cache contents.

Cache Performance Tuning for Specific Problems

The following are more specific performance problems that can be resolved from tuning the cache.

"Edit Page" screen takes a long time to load

If your installation of Confluence is suffering from this problem, it may be due to a insufficient SpacePermissions cache size. To address this problem, first determine the number of space permission objects in your Confluence instance. You can do this by running this query against your database:

```
> select count(*) from SPACEPERMISSIONS
```

Now locate the cache entry for SpacePermissions in your `confluence-coherence-cache-config.xml`:

```
<local-scheme>

<scheme-name>cache:com.atlassian.confluence.security.CachingSpacePermissionManager.
permissions</scheme-name>
  <scheme-ref>default</scheme-ref>
  <high-units>10000</high-units>
  <expiry-delay>0s</expiry-delay>
</local-scheme>
```

Adjust the **maxElementsInMemory** or **high-units** property to the number of space permissions you have (in the example above, I've used 10000). Also, just as important, you need to adjust the **timeToLiveSeconds** or **expiry-delay** property to 0.

Note: 10K of space permissions consumes approximately 8MB of memory. Please ensure there is enough memory allocated to your instance to cater for this.

How to set specific cache settings

- Find the cache name from the cache name mappings:
 - For **Confluence 2.5.x and earlier**, the cache name mappings are in file `confluence/WEB-INF/classes/com/atlassian/confluence/admin/actions/cache-name-mappings.properties`.
 - For **Confluence 2.6.0 and later**, you will find the cache name mappings in the file `com/atlassian/confluence/core/ConfluenceActionSupport.properties` which is packed into the `confluence-2.x.*.jar` file.
- Find the appropriate `<cache-mapping>` tag in `confluence-coherence-cache-config.xml` or `confluence-coherence-cache-config-clustered.xml`. If the tag doesn't exist, you can create it within the `<caching-scheme-mapping>` tag.



Attached to this page are corrected copies of [confluence-coherence-cache-config.xml](#) and [confluence-coherence-cache-config-clustered.xml](#). These are updated from a bug [CONF-11857](#).

- The `<scheme-name>` will correspond to a `<local-scheme>` tag below. It refers to a scheme reference. Either change the `high-units` tag in the scheme reference, or add a `high-units` tag to override the scheme reference. For example, the following tag would change the Content Bodies cache from the default 1000 units to 2000 units:

```
<local-scheme>
<scheme-name>cache:com.atlassian.confluence.core.ContentEntityObject.
bodyContents</scheme-name>
<high-units>2000</high-units>
<scheme-ref>default</scheme-ref>
<expiry-delay>0s</expiry-delay>
</local-scheme>
```

Another popular cache to change is the LDAP related User cache:

```
<local-scheme>
<scheme-name>user</scheme-name>
<scheme-ref>default</scheme-ref>
<high-units>5000</high-units>
<expiry-delay>300s</expiry-delay>
</local-scheme>
```

- After updating the appropriate file, you do not need to repack it into the jar to use it. You can simply place the file in your `confluence/WEB-INF/classes/` directory. The file in this directory will override the settings in your jar file. If you want to back out the changes, you only need to remove the file from your `confluence/WEB-INF/classes/` directory — then the default values in the `confluence-coherence-cache-config.xml` located in your jar file will apply.

You can find more information about configuring the Coherence cache in the [Coherence cache documentation](#).

RELATED TOPICS

[Cache Performance Tuning](#)
[Confluence Cache Schemes](#)
[Working with Confluence Logs](#)
[Operating Large or Mission-Critical Confluence Installations](#)
[Confluence Clustering Overview](#)
[Requesting Performance Support](#)
[Confluence Administrator's Guide](#)
[Configuring Confluence](#)

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.

 The information on this page *does not apply to* Confluence OnDemand.

Viewing Cache Statistics and Modifying Cache Sizes

To view the cache statistics:

- Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
- Click '**Cache Statistics**' in the left-hand panel. There you will find a list of all objects cached within Confluence.
- Click the '**Advanced**' tab for more detail. Below is an example for one of the most frequently used caches, the 'Content Object' cache.

Name	Percent Used	Effectiveness	Objects / Size	Hit / Miss / Expiry	Adjust Size	Flush
Content Object	80%	73%	4023 / 5000	374550 / 140460 / 55044	<i>Adjust Size</i>	<i>Flush</i>

About the generated numbers:

Percent Used:	=(Objects)/(Size)
Effectiveness:	$\text{=(Hits)/(Hits + Misses)}$
Objects / Size:	The number of entries in the cache / the number of total possible entries allowed (configurable).

Hit / Miss / Expiry:	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.
Adjust Size	Use this option to specify a different maximum cache size. Enter a new cache size and click the ' Adjust Size ' button to set it.
Flush:	Flushes the cache.

For instance, to calculate **Percent Used**:

$$\text{Percent Used} = \text{Objects} / \text{Size}$$

$$\text{Percent Used} = 4023/5000 = 80\%$$

To calculate **Effectiveness**:

$$\text{Effectiveness} = (\text{Hits}) / (\text{Hits} + \text{Misses})$$

$$\text{Effectiveness} = 374550 / (374550 + 140460) = 73\%$$


The [clustered versions of Confluence](#) use distributed cache called [Tangosol Coherence](#).

Watching the Cache Contents

To see the specific items in the caches, view the cache statistics at `<baseUrl>/admin/cachecontents.jsp`.

Additional Notes about Configurable Caches

Changes to cache size configurations persist across confluence restarts as they are saved in the `<confluence-home>/config/ehcache.xml` file (or `<confluence-home>/config/confluence-coherence-cache-config-clustered.xml` for a clustered instance). In most cases, a Confluence administrator will never need to know about these files. However, if it is necessary to tune cache options other than the maximum cache size, this can be done by manually editing these files. See [Cache Performance Tuning](#) for details.



Important note about clustered Confluence installations

The cache configuration file is stored in a home directory of each cluster node. When a Confluence administrator changes a cache size, all running cluster nodes will automatically update their own configuration files in their respective home directories. However, if a cluster node is not running when an administrator adjusts a cache size, the `/config/confluence-coherence-cache-config-clustered.xml` file in its home directory will not be updated. Since cluster caches are configured by the first node to start, if a node with an outdated cache configuration is the first to start up, the whole cluster would end up using the configuration of that node. However, copying this file from one node to another would resolve this issue.

Performance Tuning

If you need to tune your application when under high usage, you may like to review [this document](#) for suggestions.

Related Topics

Content by label

There is no content with the specified labels



[Administrators Guide Home](#) [Confluence Documentation Home](#)

Confluence Cache Schemes

Default Scheme

If a cache has not been defined, then it will use the default cache size and expiry. As the start of your `confluence-classes/confluence-coherence-cache-config.xml` file you will notice the following:

```
<cache-mapping>
  <cache-name>*</cache-name>
  <scheme-name>default</scheme-name>
</cache-mapping>
```

So basically all caches will default to using the default scheme, which is defined as below:

```
<!-- Default scheme -->
<local-scheme>
  <scheme-name>default</scheme-name>

  <class-name>com.atlassian.confluence.cache.tangosol.ExpiryCountingLocalCache</class-name>
  <high-units>1000</high-units>
  <expiry-delay>3600</expiry-delay>
</local-scheme>
```

I.e. with a size of 1000 Objects and an expiry of 3600 seconds. Other schemes use the above as their default and either override the size of the cache, or the length of the expiry.

 *The information on this page **does not apply** to Confluence OnDemand.*

Common Schemes

In addition to the default scheme, there are also common schemes used in Confluence caches:

```

<!-- Common schemes -->
<local-scheme>
  <scheme-name>large</scheme-name>
  <scheme-ref>default</scheme-ref>
  <high-units>10000</high-units>
</local-scheme>
<local-scheme>
  <scheme-name>medium</scheme-name>
  <scheme-ref>default</scheme-ref>
  <high-units>5000</high-units>
</local-scheme>
<local-scheme>
  <scheme-name>small</scheme-name>
  <scheme-ref>default</scheme-ref>
  <high-units>100</high-units>
</local-scheme>
<local-scheme>
  <scheme-name>large-transient</scheme-name>
  <scheme-ref>default</scheme-ref>
  <high-units>10000</high-units>
  <expiry-delay>300s</expiry-delay>
</local-scheme>
<local-scheme>
  <scheme-name>user</scheme-name>
  <scheme-ref>default</scheme-ref>
  <high-units>5000</high-units>
  <expiry-delay>300s</expiry-delay>
</local-scheme>

```

RELATED TOPICS

[Cache Performance Tuning](#)
[Confluence Cache Schemes](#)
[Cache Performance Tuning for Specific Problems](#)
[Requesting Performance Support](#)
[Confluence Administrator's Guide](#)
[Configuring Confluence](#)

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](#).

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](#)
- [Confluence Administrator's Guide](#)



The information on this page does not apply to Confluence OnDemand.

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 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.

On this page:

- [Basic performance troubleshooting steps](#)
- [Requesting basic performance support](#)
- [Advanced performance troubleshooting](#)

Related pages:

- [Memory Usage and Requirements](#)
- [Operating Large or Mission-Critical Confluence Installations](#)
- [Confluence Administrator's Guide](#)



The information on this page does not apply to Confluence OnDemand.

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'`)

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 [AWStats](#), 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 [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.

 The information on this page *does not apply* to Confluence OnDemand.

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.

Enabling Page-Request Profiling

 To see just the slow performing macros, see [Identifying Slow Performing Macros](#).

From **Confluence 2.7**, you can use the '**Logging and Profiling**' option to enable or disable profiling.

 You need to have [System Administrator](#) permissions in order to perform this function.

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'.

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](#)

 *The information on this page does not apply to Confluence OnDemand.*

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	
	INFO ▼	Add entry

Existing Levels

Class/Package Name	Current Level	New Level	
atlassian.plugin	INFO	INFO ▼	Remove
com.atlassian.confluence.admin.actions	INFO	INFO ▼	Remove

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())
    ...
```

 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 if a user's web browser supports it, Confluence will compress the data it sends to the user. This will speed up Confluence over slow or congested Internet links, and reduce the amount of bandwidth consumed by a Confluence server.

i Gzipping the HTTP Response is available in Confluence 1.4 and later.

You should 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 having to compress each HTTP response may in fact slow Confluence down.

 The information on this page *does not apply to* Confluence OnDemand.

Known issues in Confluence 2.7 and earlier

There are known issues with the GZip filter and memory consumption evident in versions 2.7 of Confluence and earlier ([CONF-9930](#)). If you are running a large instance of Confluence 2.7 or earlier and frequently experiencing 'out of memory' errors, we recommend that you do not enable HTTP compression. These issues have been resolved in Confluence 2.8.

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**'.

In Confluence 2.8 and later, you can 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](#). Generally speaking, it is unlikely that you will need to alter this file.

RELATED TOPICS

[Performance Tuning](#)
[Confluence Administrator's Guide](#)

Garbage Collector Performance Issues

 This document relates broadly to memory management with Oracle's Hotspot JVM. These are recommendations based on Support's successful experiences with customers and their 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.

 The information on this page *does not apply to* Confluence OnDemand.

Use a small heap

Keep your heap as small as possible, without the instance experiencing OutOfMemory errors. The default size for Confluence is 512mb. 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. Generally, the JVM is much better at this than we humans are, and most of the time simply allowing it 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

Enable [Garbage Collection Logging](#), and use [Chewiebug's GCViewer](#) to view the resulting logs.

Confluence Installation and Upgrade Guide

About the Installation and Upgrade Guide

This guide covers how to install and upgrade Confluence.

Information on the features and changes in specific Confluence releases can be found in the [Confluence Release Notes](#).

For information on using and administering Confluence refer to the [Confluence User's Guide](#) and [Confluence Administrator's Guide](#).

- [System Requirements](#)
 - [Server Hardware Requirements Guide](#)
 - [Example Size and Hardware Specifications From Customer Survey](#)
 - [Running Confluence in a Virtualised Environment](#)
- [Confluence Installation Guide](#)
 - [Installing Confluence](#)
 - [Installing the Confluence EAR-WAR Edition](#)
 - [Confluence Cluster Installation](#)
 - [Creating a Dedicated User Account on the Operating System to Run Confluence](#)
- [Confluence Setup Guide](#)
 - [Configuring JIRA Integration in the Setup Wizard](#)
- [Upgrading Confluence](#)
 - [Upgrading Beyond Current Licensed Period](#)
 - [Confluence Post-Upgrade Checks](#)
 - [Upgrading Confluence EAR-WAR Distribution](#)
 - [Migration from Wiki Markup to XHTML-Based Storage Format](#)

Downloads



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

Other Resources

[Confluence Release Notes](#)

[Confluence User's Guide](#)

[Confluence Administrator's Guide](#)

[Confluence Knowledge Base](#)

[Atlassian Answers](#)

- [Migration of Templates from Wiki Markup to XHTML-Based Storage Format](#)
- [Upgrading Confluence Manually](#)
- [Supported Platforms](#)
 - [End of Support Announcements for Confluence](#)
 - [Supported Platforms FAQ](#)
- [Migrating Confluence Between Servers](#)
 - [From Confluence Evaluation through to Production Installation](#)
 - [Migrating from Confluence OnDemand to a Confluence Installed Site](#)

System Requirements

Confluence works with a broad range of operating systems, database systems and application servers. Provided you have the technical knowledge, it is very likely that you will be able to run Confluence with an 8-year-old database or even on some 8-year-old hardware. Realistically, it is not technically feasible for us to provide our legendary support service on all environments available. There can only be a finite number of platforms and release versions of those that we support.

Our rule of thumb when releasing a new version of Confluence is that we will officially support platforms that have been released within the last one to two years (or the latest version of that platform if no new version of it was released in that period). This does not necessarily mean that you will need to upgrade your database or application server every time you upgrade Confluence. However, if you do run into problems with an unsupported version of a database or application server, we may have to ask you to upgrade to something newer.

Please refer to our [Supported Platforms](#) topic for details on platforms that we currently support in this version of Confluence and our [Supported Platforms FAQ](#) topic for details on our support handling procedures.

On this page:

- [Confluence Software Requirements](#)
 - [Operating Systems](#)
 - [Application Servers](#)
 - [Databases](#)
 - [Java](#)
 - [Please Note: Impact of Antivirus Software](#)
- [Confluence Hardware Requirements](#)
- [Atlassian Hosted Solutions – Atlassian OnDemand](#)

Confluence Software Requirements

Please read the [Supported Platforms](#) page for Confluence. That page contains a list of specific software that Confluence will work with.

Operating Systems

Atlassian supports the operating systems listed on the [Supported Platforms](#) page.

If you would like to run Confluence on virtualised hardware, please read our [Running Confluence in a Virtualised Environment](#) document first.

Application Servers

An application server is required to run Confluence. Apache Tomcat is bundled with the distribution.

Atlassian only supports the application servers listed on the [Supported Platforms](#) page, provided they are running on Windows, Linux, or Solaris. If you are using a different application server or earlier version, we may ask you to migrate to one of the supported application servers before we can provide you with further support.

Databases

A database is required to run Confluence. Atlassian supports the databases listed on the [Supported Platforms](#) page.

When evaluating Confluence, you can use the embedded database included in the Confluence installation.

When moving to a production installation, you must set up an external database server. If you have no preference for a particular database, we highly recommend using PostgreSQL. This is a scalable, robust and free database server that is also easy to set up. For database setup information, please refer to [Database Configuration](#).

Java

Confluence requires the Java Runtime Environment (JRE) installed.

If using the zip or archive distribution of Confluence, you will need to install a [supported](#) JRE. The automated installer bundles Java and will install this for you.

For instructions on installing the JRE for Windows and Linux/Solaris, please refer to [Installing Java for Confluence](#).

Please Note: Impact of Antivirus Software

The presence of antivirus software on your operating system running Confluence greatly decreases the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental and may even cause errors in Confluence.

You should configure your antivirus software to ignore the following directories:

- Confluence [home directory](#)
- Confluence's index directory
- All database-related directories

 This recommendation above is particularly important if you are running Confluence on Windows. No matter how fast your hardware is, antivirus software will almost always have a negative impact on Confluence's performance and may render Confluence impossible to use.

Confluence Hardware Requirements

Please be aware that while some of our customers run Confluence on SPARC-based hardware, Atlassian only officially supports Confluence running on x86 hardware and 64-bit derivatives of x86 hardware.

See [Server Hardware Requirements Guide](#) for details.

Refer also to the tips on [reducing out of memory errors](#), in particular the section on [Permanent Generation Size](#).

Atlassian Hosted Solutions – Atlassian OnDemand

If you do not have the resources to set up and maintain a Confluence installation locally, consider Atlassian hosted solutions. Atlassian can run and maintain your installation of Confluence, handling all the testing, monitoring and upgrading processes for you. For more information, please refer to the information about [Confluence OnDemand](#) on our website.

Server Hardware Requirements Guide

Server administrators can use this guide in combination with the free Confluence trial period to evaluate their server hardware requirements. Because server load is difficult to predict, live testing is the best way to determine what hardware a Confluence instance will require in production.

Peak visitors are the maximum number of browsers simultaneously making requests to access or update pages in Confluence. Visitors are counted from their first page request until the connection is closed and if public access is enabled, this includes internet visitors as well as logged in users. Storage requirements will vary depending on how many pages and attachments you wish to store inside Confluence.

Minimum hardware requirements

The values below refer to the minimum available hardware required to run Confluence only, eg the minimum heap size to allocate to Confluence is 512mb. You will need additional physical hardware, of at least the minimum amount required by your Operating System, and any other applications that run on the server. Also please note that these are a guide only, and your configuration may require more.

On small instances, server load is primarily driven by peak visitors.

5 Concurrent Users

- 2GHz+ CPU
- 512MB RAM
- 5GB database space

25 Concurrent Users

- Quad 2GHz+ CPU
- 2GB+ RAM
- 10GB database space

Note: Please be aware that while some of our customers run Confluence on SPARC-based hardware, Atlassian only officially supports Confluence running on x86 hardware and 64-bit derivatives of x86 hardware. Confluence typically will not perform well in a tightly constrained, shared environment - examples include an AWS micro.t1 instance. Please be careful to ensure that your choice of hosting platform is capable of supplying sustained processing and memory capacity for the server, particularly the processing-intensive startup process.

On this page:

- [Minimum hardware requirements](#)
- [Example hardware specifications](#)
- [Server load and scalability](#)
- [Maximum reported usages](#)
- [Hard disk requirements](#)
- [Professional assistance](#)
- [Example - <https://confluence.atlassian.com/>](#)

Related pages:

- [Confluence Installation Guide](#)
- [Operating Large or Mission-Critical Confluence Installations](#)
- [Managing Application Server Memory Settings](#)
- [Confluence Clustering Overview](#)
- [Running Confluence in a Virtualised Environment](#)

Example hardware specifications

These are example hardware specifications for non-clustered Confluence instances. It is not recorded whether the RAM refers to either total server memory or memory allocated to the JVM, while blank settings indicate that the information was not provided.

Accounts	Spaces	Pages	CPUs	CPU (GHz)	RAM (Meg)	Notes
150	30	1,000	1	2.6	1,024	
350	100	15,000	2	2.8	1,536	
5,000	500		4	3	2,024	
10,000	350	16,000	2	3.8	2,024	
10,000	60	3,500	2	3.6	4,048	
21,000	950		2	3.6	4,048	

85,000	100	12,500	4	2.6	4,048	3 machines total: application server, database server, Apache HTTPD + LDAP tunnel server. See Accenture's slides and video for full details (That link isn't working, but the slides can be found here)

Server load and scalability

When planning server hardware requirements for your Confluence deployment, you will need to estimate the server scalability based on peak visitors, the editor to viewer ratio and total content.

- The editor to viewer ratio is how many visitors are performing updates versus those only viewing content
- Total content is best estimated by a count of total spaces

Confluence scales best with a steady flow of visitors rather than defined peak visitor times, few editors and few spaces. Users should also take into account:

- Total pages is not a major consideration for performance. For example, instances hosting 80K of pages can consume under 512 meg of memory
- Always [use an external database](#), and check out the [performance tuning guides](#).

As mentioned on the documentation for [Operating Large or Mission-Critical Confluence Installations](#), some important steps are loadtesting your usecase and monitoring the system continuously to find out where your system could do better and what might need to improve in order to scale further.

Maximum reported usages

These values are largest customer instances reported to Atlassian or used for performance testing. [Clustering for load balancing](#), database tuning and other performance tuning is recommended for instances exceeding these values.

Most Spaces	1700
Most Internal Users	15K
Most LDAP Users	100K
Most Pages	80K

Hard disk requirements

All wiki content is stored in the database, while attachments use [either the database or filesystem](#). For example, the wiki instance you are reading now uses approximately 2.8 GB of database space and 116 GB of disk space. The more attachments you have, the more disk space you will require.

Private and public comparison

Private instances manage their users either internally or through a user repository such as LDAP, while online

instances have public signup enabled and must handle the additional load of anonymous internet visitors. Please keep in mind that these are examples only, not recommendations:

Use Case	Spaces	User Accounts	Editors	Editor To Viewer Ratio	Pages	Page Revisions	Attachments	Comments	Total Data Size (GB)	Notes
Online Documentation	140	11,500	1,000	9%	8,800	65,000	7,300	11,500	10.4	
Private Intranet	130	180	140	78%	8,000	84,000	3,800	500	4.5	
Company-Wide Collaboration	100	85,000	1,000+	1%+	12,500	120,000	15,000			Accenture - see slides and video for full details (That link isn't working, but the slides can be found here .)

Professional assistance

For large instances, it may be worthwhile contacting an [Atlassian Expert](#) for expertise on hardware sizing, testing and performance tuning. Simply [contact a local Expert directly](#) or [email our Experts team](#) for a recommendation.

Example - <https://confluence.atlassian.com/>

Here is a breakdown of the disk usage and memory requirements for this wiki, as at April 2013:

Database size	2827 MB
Home directory size	116 GB
Average memory in use	1.9 GB

Size of selected database tables

Data	Relevant Table	Rows	Size
Attachment metadata	attachments	193903	60 MB
Content and user properties	os_propertyentry (?)	639737	255 MB
Content bodies (incl. all versions of blogs, pages and comments)	bodycontent	517520	1354 MB

Content metadata (incl. title, author)	content	623155	459 MB
Labels	label (5982, 1264 kB), content_label (134151, 46 MB)	140133	47.2 MB
Users	users	38766	6200 kB

Note: not all database tables or indexes are shown, and average row size may vary between instances.

Size of selected home directory components

Data	Files	Size
Attachments (incl. all versions)	207659	105 GB
Did-you-mean search index	10	14 MB
Office Connector cache	3506	456 MB
Plugin files	1851	669 MB
Search index	448	3.9 GB
Temporary files	14232	5 GB
Thumbnails	86516	1.7 GB
Usage index (now disabled)	239	2.6 GB

Note: not all files are shown, and average file size may vary between instances.

Example Size and Hardware Specifications From Customer Survey

Below are the results of a survey conducted by Atlassian in July 2007, showing some capacity statistics for Confluence users. The figures are broken down by industry and number of users.

Num Users	Length of time in production	Database	Application Server	Num CPUs/Cores	Physical Memory/RAM	Operating System	Satisfaction with Confluence Performance
Banking/Finance							
26 - 50	3-6 Months Ago	Microsoft SQL Server	Confluence distribution /Apache Tomcat	2	2G	Windows	Neutral
26 - 50	2 Years Ago	Sybase ASE	Weblogic	>8	>16G	Unix	Satisfied
51 - 250	3-6 Months Ago	Oracle	Confluence distribution /Apache Tomcat	2	4G	Unix	Neutral
501 - 1,000	3-6 Months Ago	Microsoft SQL Server	Websphere	2	2G	AIX	Satisfied

1,001 - 5,000	3-6 Months Ago	Oracle	Confluence distribution /Apache Tomcat	2	4G	Windows	Satisfied
1,001 - 5,000	2 Years Ago	Oracle	Websphere	4	>16G	Solaris	Extremely Satisfied
5,001 - 10,000	10-12 Months Ago	Microsoft SQL Server	Confluence distribution /Apache Tomcat	4	16G	Linux	Satisfied
Education							
1-25	2 Years Ago	DB2	Confluence distribution /Apache Tomcat	2	2G	Linux	Satisfied
26 - 50	10-12 Months Ago	MySQL	Confluence distribution /Apache Tomcat	2	2G	Linux	Extremely Satisfied
51 - 250	<3 Months Ago	Oracle	Confluence distribution /Apache Tomcat		1G	Windows	Unsatisfied
51 - 250	10-12 Months Ago	Oracle	Confluence distribution /Apache Tomcat	1	2G	Unix	Extremely Satisfied
Engineering/Aerospace							
251 - 500	7-9 Months Ago	Oracle	Confluence distribution /Apache Tomcat	1	1G	Mac OS X	Satisfied
1,001 - 5,000	7-9 Months Ago	Microsoft SQL Server	JBoss	2	4G	Linux	Satisfied
Entertainment							
1,001 - 5,000	10-12 Months Ago	PostgreSQL	Confluence distribution /Apache Tomcat	2	8G	Linux	Extremely Satisfied
Government							
51 - 250	2 Years Ago	MySQL	Confluence distribution /Apache Tomcat	2	2G	Mac OS X	Extremely Satisfied
Technology							

501 - 1,000	7-9 Months Ago	MySQL	Confluence distribution /Apache Tomcat	1	2G	Linux	Satisfied
Telecommunications & Media							
1-25	3-6 Months Ago	Confluence distribution /HSQL	Confluence distribution /Apache Tomcat	1		Linux	Satisfied
1-25	7-9 Months Ago	MySQL	Confluence distribution /Apache Tomcat	1	2G	Linux	Satisfied
26 - 50	10-12 Months Ago	MySQL	Confluence distribution /Apache Tomcat	2	2G	Linux	Satisfied

Running Confluence in a Virtualised Environment

This page provides pointers for things to look at when running Confluence on virtualised hardware.

Summary

Running Confluence in a virtual machine (VM) requires specialised skills to set up and manage the virtualised environment. In particular, the performance of Confluence can be affected by the activity of other VMs running on the same infrastructure, as well as how you configure the Confluence VM itself.

Atlassian supports Confluence sites running on a virtualised environment, but we can only offer support for problems which are unrelated to the environment itself. You will need to understand and be prepared to manage your own virtualised environment if you wish to run Confluence on such a platform.

On this page:

- [Summary](#)
- [Recommendations](#)
- [Further help](#)

Related pages:

- [Server Hardware Requirements Guide](#)
- [Operating Large or Mission-Critical Confluence Installations](#)
- [Confluence Installation Guide](#)

Recommendations

The following recommendations come from our experience in running and testing Confluence in virtualised environments like VMWare and KVM, and our experience in working with customers running on these platforms.

- **Know your platform.** Consult the documentation for your operating system and your chosen virtualisation technology, for details on setting up a reliable VM (virtual machine) image.
- **Allocate enough memory.** As a Java web application, Confluence requires a relatively large memory allocation, compared to some other web technologies. Ensure that your VM images have enough physical memory allocated to run Confluence without swapping.
- **Handle high I/O.** Under normal usage, Confluence requires a significant number of input/output (I/O) operations to the database and home directory for each web request. Ensure that you use the correct drivers and consider how you make storage available to your VMs to optimise this access.
- **Handle peak CPU and memory usage.** For certain operations (including PDF export, Office document processing, and displaying large pages) Confluence requires a significant amount of CPU and memory. Ensure that your virtualisation infrastructure has the flexibility and capacity to deal with peak load, not just

idle load.

- **Synchronise time correctly.** Some customers have had problems with time synchronisation between the VM and the host system. This causes problems in Confluence due to irregularities in the execution of scheduled tasks. We strongly recommend checking your VM time sync if you have issues with scheduled tasks in a virtualised environment.

Further help

For further assistance in setting up a virtualised environment for running Confluence, you may want to consult an [Atlassian Expert](#). Several experts have experience with installation and performance tuning, and can help you with your Confluence configuration.

Confluence Installation Guide

Prerequisites

Before beginning to install [Confluence](#), please check that:

- Your system meets the [minimum system requirements](#) to run Confluence.
- This version of the Confluence documentation matches the version of Confluence that you are installing. The Confluence documentation version you are currently viewing is indicated toward the top of the page tree on the left or in the 'breadcrumb trail' in the top banner of this page. If you need to access a different version of the Confluence documentation, use the control at the top of the page tree on the left or you can access it from the [documentation home page](#).

If you have chosen a package that includes add-ons such as Team Calendars or the SharePoint Connector you will need to install these from within Confluence after your setup is complete. See [Finding new add-ons](#) for information on how to find and install the add-ons.

Choose the Confluence Installation Type

Choose the type of Confluence installation you'd like from the table below, and follow the link(s) to the installation instructions.

Installation Type	Description
• Installing Confluence on Windows • Installing Confluence on Linux	Install Confluence via the Atlassian installer. This is the easiest method of installing Confluence.  This is the best option for evaluators.
• Installing from a Zip File on Windows • Installing From an Archive File on Linux	This option requires you to manually carry out installing the files and configuring system properties.  Use this option if there is no specific installer for your operating system.
EAR/WAR distribution (Zip Archive)	This distribution allows you to deploy Confluence onto your own existing application server, instead of the Apache Tomcat server bundled with the regular distribution.
Confluence Clusters (Zip Archive)	Install Confluence as a series of clusters, to improve performance or availability. Please read the Confluence Clustering Overview and the Cluster Checklist before you consider installing Confluence in a cluster.

Please read [Running Confluence in a Virtualised Environment](#) if you are interested in running Confluence in a virtual machine.

If you wish to upgrade Confluence, see [Upgrading Confluence](#).

Related Topics

[Upgrading Confluence](#)
[System Requirements](#)

Installing Confluence

Choose the type of Confluence installation you'd like from the table below and follow the link to the installation instructions. When you have finished the installation phase, you will be prompted to start the setup phase.

Installation Type	Description
Installing Confluence on Windows Installing Confluence on Linux	Install Confluence via the Atlassian installer. This is the easiest method of installing Confluence.  This is the best option for evaluators.
Installing from a Zip File on Windows Installing From an Archive File on Linux	This option requires you to manually carry out installing the files and configuring system properties.  Use this option if there is no specific installer for your operating system.

 If you have not already done so, please verify that this version of the Confluence documentation matches that of the Confluence version you are installing. The Confluence documentation version you are currently viewing is indicated toward the top of the page tree on the left or in the 'breadcrumb trail' in the top banner of this page. If you need to access a different version of the Confluence documentation, use the control at the top of the page tree on the left or you can access it from the [documentation home page](#).

Take me back to the [Confluence Installation Guide](#).

Installing Confluence on Windows

This guide describes how to install a new Confluence installation on Windows using the automated 'Windows Installer'. You can also install Confluence from a 'zip' archive — see [Installing Confluence on Windows from Zip File](#) for details.

If you are upgrading Confluence, please refer to the [Upgrading Confluence](#) guide.

Please Note:

- Some anti-virus or other Internet security tools may interfere with the Confluence installation process and prevent the process from completing successfully. If you experience or anticipate experiencing such an issue with your anti-virus/Internet security tool, disable this tool first before proceeding with the Confluence installation.
- Before you begin installing Confluence, please read the [System Requirements](#) page.

On this page:

- [Using the Installation Wizard](#)
 - [1. Download and Run the Confluence 'Windows Installer'](#)
 - [2. Starting Confluence](#)
 - [3. Run the Setup Wizard](#)
 - [4. Next Steps](#)
- [Performing an Unattended Installation](#)
 - [Download and Run the Confluence 'Windows Installer' in Unattended Mode](#)

Using the Installation Wizard

Use the installation wizard if you are installing Confluence on your server for the first time or you wish to specify your installation options.

If you have previously installed Confluence using the installation wizard and wish to re-install Confluence again with the same installation options, you can re-install Confluence in 'unattended mode' without any user input required (see [below](#) for details).

1. Download and Run the Confluence 'Windows Installer'

✓ To install Confluence as a service, the Windows Installer must be run using a Windows administrator account. While you can run the Windows Installer with a non-administrator account, your installation options will be much more limited.

1. Download the **Confluence 'Windows Installer'** (.exe) file from the [Confluence Download](#) page.
2. Run the installer file to start the installation wizard.
 - ❗ If a Windows 7 (or Vista) 'User Account Control' dialog box requests if you want to allow the installation wizard to make changes to your computer, click 'Yes'. If you do not, the installation wizard will have restricted access to your operating system and any subsequent installation options will be limited.
3. Choose between the '**Express Install**' or '**Custom Install**' options:
 - a. **Express Install** — If you choose this option, Confluence will be installed with default settings which are shown in the next step of the installation wizard. If you want to customise any of these options, click the '**Back**' button and choose the '**Custom Install**' option instead.
 - b. **Custom Install** — If you choose this option, Confluence will prompt you to specify the following options (which are presented during subsequent steps of the installation wizard and pre-populated with default values):
 - The 'Destination Directory' in which to [install Confluence](#).
 - The [Confluence Home Directory](#) (which must be unique for each Confluence installation).
 - The Windows 'Start' menu folder options.
 - The *TCP ports* (i.e. an HTTP connector port and a control port) that Confluence will operate on.
 - If you are running the installer using an administrator account, you will be prompted to 'Install Confluence as a service' (recommended). You can also do this manually later, as described in [Start Confluence Automatically on Windows as a Service](#).
4. The installation wizard will install Confluence onto your operating system and will start Confluence automatically when the wizard finishes. Confluence will also be launched automatically in your browser window if you chose this option.

❗ Please Note:

- If you chose to install Confluence as a service, the Confluence service will be run as the Windows 'SYSTEM' user account. To change this user account, see [Changing the Windows user that the Confluence service uses](#).
- If you do not install Confluence as a service, then once started, Confluence will be run as the Windows user account under which Confluence was installed.
- If you use Confluence running on a Windows Server in production, we strongly recommend creating a dedicated user account (e.g. with username 'confluence') for running Confluence.
 - For more information about creating a dedicated user account and defining which directories this account should have write access to, refer to our [guidelines](#).
 - ✓ If your Windows Server is operating under Microsoft Active Directory, ask your Active Directory administrator to create a dedicated user account that you can use to run Confluence (with no prior privileges).
 - If Confluence is installed as a service, do not forget to [change the user account that runs the Confluence service](#) to your dedicated user account for running Confluence.

2. Starting Confluence

If Confluence is not already started, you can start Confluence using the appropriate [Windows 'Start' menu shortcut](#) or [command prompt option](#).

Once Confluence is started, you can access Confluence from the appropriate [Windows 'Start' menu shortcut](#) or a [browser](#) on any computer with network access to your Confluence server.

2.1 Windows 'Start' Menu Shortcuts

The Installer will have created the following Windows 'Start' menu shortcuts:

- **Access Confluence** — opens a web browser window to access your Confluence application.
 - ❗ Your Confluence server must have been started for this shortcut to work.
- **Start Confluence Service** — starts up the Apache Tomcat application server which runs your

Confluence installation, so that you can access Confluence through your web browser.

- **Stop Confluence Service** — stops the Apache Tomcat application server which runs your Confluence installation. You will not be able to access Confluence through your web browser after choosing this shortcut.
- **Uninstall Confluence** — [uninstalls Confluence](#) from your Windows operating system.

2.2 Starting and Stopping Confluence from a Command Prompt

Enter the `bin` subdirectory of your Confluence installation directory and run the appropriate file:

- `start-confluence.bat` (to start Confluence)
- `stop-confluence.bat` (to stop Confluence)

 If you followed our [guidelines](#) for running Confluence with a dedicated user account, then to run Confluence as this user account (e.g. 'confluence'), use the `runas` command to execute `start-confluence.bat`. For example:

- `> runas /env /user:<DOMAIN>\confluence start-confluence.bat`
(where `<DOMAIN>` is your Windows domain or computer name.)

2.3 Accessing Confluence from a Browser

You can access Confluence from any computer with network access to your Confluence server by opening a [supported web browser](#) on the computer and visiting this URL:

- `http://<computer_name_or_IP_address>:<HTTP_port_number>`

where:

- `<computer_name_or_IP_address>` is the name or IP address of the computer on which Confluence is installed and
- `<HTTP_port_number>` is the HTTP port number specified when you installed Confluence ([above](#)).

 If Confluence does not appear in your web browser, you may need to [change the port that Confluence runs on](#).

3. Run the Setup Wizard

See the [Confluence Setup Guide](#).

4. Next Steps

- See [Confluence 101](#).
- If you did not install Confluence as a service, you will need to [start Confluence manually](#) every time you restart your computer. To change your Confluence installation to run as a service, please see [Start Confluence Automatically on Windows as a Service](#).
- To get the most out of Confluence, please see [Performance Tuning](#).

Performing an Unattended Installation

If you have previously installed Confluence using the installation wizard ([above](#)), you can use a configuration file from this Confluence installation (called `response.varfile`) to re-install 'unattended mode' without any user input required.

Installing Confluence in unattended mode saves you time if your previous Confluence installation was used for testing purposes and you need to install Confluence on multiple server machines based on the same configuration.

Please Note:

- The `response.varfile` file contains the options specified during the installation wizard steps of your previous Confluence installation. Hence, do not uninstall your previous Confluence installation just yet.
- If you intend to modify the `response.varfile` file, please ensure all directory paths specified are absolute, for example, `sys.installationDir=C:\\Program Files\\Atlassian\\Confluence`. Unattended installations will fail if any relative directory paths have been specified in this file.

Download and Run the Confluence 'Windows Installer' in Unattended Mode

1. Download the **Confluence 'Windows Installer'** (.exe) file from the [Confluence Download Center](#) to a suitable location.
2. Open the Windows command prompt and perform the remaining steps in the command prompt.
3. Copy the response.varfile file located in the .install4j subdirectory of your previous Confluence installation directory, to the same location as the downloaded 'Windows Installer' file.
 You can [uninstall your previous Confluence installation](#) after this step. Save your response.varfile if you need to install Confluence on multiple machines.
4. Change directory (cd) to the location of the 'Windows Installer' file and run the following command:

```
atlassian-confluence-X.Y.exe -q -varfile response.varfile
```

Where:

- X.Y — refers to the version of Confluence you are about to install.
 - -q — instructs the installer to operate in unattended mode (i.e. 'quietly').
 - -varfile response.varfile — specifies the configuration file containing the configuration options used by the installer. The location and name of the configuration file should be specified after the -varfile option.
5. Confluence will start automatically when the silent installation finishes. Continue from step 2 [Starting Confluence](#) (above).

Installing Confluence on Windows from Zip File

These instructions apply to:

- Confluence distributed as an archive file. This distribution includes Apache Tomcat as the application server.
- **Windows** systems. For other operating systems please refer to the [Confluence Installation Guide](#).
- Manual installation and configuration using a **zipped download file**. For a simpler installation process, please use the [Confluence Installer](#) instead.

Also, please check that the version of Confluence which you are installing coincides with the version that this documentation is written for.

On this page:

1. Before you Start

Please check the following points:

1. Ensure that your system meets the minimum requirements to run Confluence. For more information, please refer to our [Supported Platforms](#) topic and for further details, our [System Requirements](#) topic.
2. Have your Confluence license key ready. You can [obtain a trial, free or commercial license](#) now, or [retrieve your existing license key](#).

2. Install Java

Please refer to [Installing Java for Confluence](#). If you are certain that this has already been installed and that the JAVA_HOME environment variable has been correctly configured, then proceed to the next step.

3. Download the Confluence Installation File

1. If you have not downloaded Confluence already, [download the zip file](#).
2. Please check your unzip program before extracting the downloaded zip file. You should use a third-party unzip program like 7Zip or Winzip. If you do not have one, please download and install one of these before continuing:
 - [7Zip](#) (recommended). If in doubt, download the '32-bit.exe' version.
 - [Winzip](#).
3. Use your unzip program to unzip the installation file to a directory such as c:\confluence.
 - Do not use spaces in your directory path.

i The directory into which you unzipped the Confluence installation is called the [Confluence Installation directory](#). Next, you will define the [Confluence Home directory](#).

4. Define your Confluence Home Directory

Now you need to define the [Confluence Home directory](#). This is where Confluence will store its configuration information, indexes and attachments.

✓ Tip: Another term for 'home directory' would be 'data directory'.

We suggest using different paths for your installation and home directories. This will facilitate easier upgrades.

i **Examples of Installation and Home Directories**
Installation directory: c:\confluence\confluence-vX.X
Home directory: c:\confluence\data

1. Open your Confluence Installation directory (created when you unzipped Confluence — see [above](#)).
2. Under the Installation directory, open this file: confluence\WEB-INF\classes\confluence-init.properties in a text editor such as Notepad.
3. Scroll to the bottom of the text and find this line:

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

4. Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with confluence.home
5. If you decide to change the Confluence Home directory from the default, please note the following:
 - Avoid spaces in the directory path or file name.
 - Use forward slashes '/' to define the path.

For example:

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

5. Check the Ports

If you have another application running on your machine which is using the same ports that Confluence uses by default, you may need to change the port which Confluence will use. For example, if you have an installation of [JIRA](#) running on this machine, JIRA might be already using the port which Confluence requests by default.

By default, Confluence listens on port '8090'. If this port is already in use in your installation, follow these instructions to change the ports:

- To change the ports for Confluence, open the file conf/server.xml under your [Confluence Installation directory](#). The first four lines of the file look like this:

Default conf/server.xml

```
<Server port="8000" shutdown="SHUTDOWN" debug="0">
  <Service name="Tomcat-Standalone">
    <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
port="8090" minProcessors="5" maxProcessors="75"
      enableLookups="true" redirectPort="8443" acceptCount="10"
debug="0" connectionTimeout="20000" useURValidationHack="false"/>
    ...
```

You need to modify both the **server** port (default is 8000) and the **connector** port (default is 8090) to

ports that are free on your machine. The server port is required by Tomcat but is not user facing in any way. The connector port is what your users will use to access Confluence, eg in the snippet above, the URL would be <http://example.com:8090>.

 **Hint:** You can use `netstat` to identify free ports on your machine. See more information on using `netstat` on Windows or on Linux.

For example, here are the first four lines of a modified `server.xml` file, using ports '8020' and '8099':

Modified conf/server.xml using ports 8020 and 8099

```
<Server debug="0" shutdown="SHUTDOWN" port="8020">
  <Service name="Tomcat-Standalone">
    <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
port="8099" minProcessors="5" maxProcessors="75"
      enableLookups="true" redirectPort="8443" acceptCount="10"
debug="0" connectionTimeout="20000" useURIVValidationHack="false"/>
    ...
```

To access Confluence in this configuration, point your web browser to <http://localhost:8099/>.

 You should also ensure at this point that if you are using a firewall, it is configured to allow http/https traffic over the port you have chosen.

Once this is working, if this is the URL your users will use, remember to [update your Base URL](#) to point to the new URL.

You will find more information on [this page](#).

6. Select an External Database

This step is optional for users evaluating Confluence. However, if you are installing Confluence for production purposes, this step is mandatory. Please refer to the [database](#) requirements listed on our [System Requirements](#) topic for help in choosing an external database.

 External databases are those listed on our [Supported Platforms](#) topic, excluding HSQLDB, which is bundled with Confluence and should not be used in production.

When you have chosen your external database, follow the the appropriate [database setup guide](#) to set up your database to work with Confluence.

You can learn more about migration from an existing installation or use of the evaluation database [here](#). You will continue to use the database setup guide during the Confluence Setup Wizard. (See [step 8](#) below.)

7. Start Confluence

1. Go to your Confluence Installation directory (created when you unzipped Confluence — see [above](#)).
2. Under your Confluence Installation directory, open the `bin` directory and run the startup script: `startup.bat`. A command prompt window should appear.

 Please do not close this command prompt window. If you do so, Confluence will stop running.

Troubleshooting

If the window closes immediately when started, this means that an error is preventing Confluence from starting. To view this error:

- a. Open a command prompt: Click on your 'Start' menu, then click 'Run'. In the Run box, type `cmd` and click 'OK'.
- b. From the command prompt, go to your Confluence Installation directory.
- c. Go into the `bin` subdirectory.
- d. Run `catalina.bat run`.

 You should not run `startup.bat` at this point, because that would still produce a

- popup window that would close straight away.
- e. Read the error message.
 - f. Find the solution to that error in the [Installation Troubleshooting](#) section of the Confluence Knowledge Base.
3. Once Confluence is running, open a web browser and visit <http://localhost:8090/>.
- ✔ If you changed the port earlier, use the port you specified in [step 5 above](#).
 - i If your web browser window shows an error, try waiting for 30 seconds or so and then refresh the browser page.

8. Next Step is the Confluence Setup Wizard

The Confluence Setup Wizard should appear in your web browser, prompting you to enter your license key. Follow the instructions on the screens, and read more guidelines on the [Confluence Setup Wizard](#).

9. Start Confluence automatically on Windows as a Service

Confluence should be [run as a service](#).

Related Topics

[Change listen port for Confluence](#)

[Running Confluence Over SSL or HTTPS](#)

[Confluence Setup Guide](#)

[Configuring Confluence](#)

Uninstalling Confluence from Windows

This page describes the procedure for uninstalling an instance of Confluence which has been installed using the [Windows Installer](#).

To uninstall Confluence from Windows:

1. Log in to Windows as the same user that was used to install Confluence with the [Windows Installer](#).
2. Start the uninstaller by doing either of the following:
 - Click the Windows **Start Menu** > **All Programs** > **Confluence** > **Uninstall Confluence**
 - OR**
 - Open the Windows Control Panel, choose **Add or Remove Programs** (on Windows XP) or **Programs and Features** on (Windows 7, Vista) and then select **Confluence X.Y** from the list of applications and click **Uninstall/Change**.
 - OR**
 - Open the Windows command prompt and do the following:
 - a. Change directory to your Confluence installation directory
 - b. Run the `uninstall.exe` file
3. Follow the prompts to uninstall Confluence from your computer.

i Please note:

- The uninstaller will not delete the Confluence Home Directory.
- All log files that were generated while Confluence was running will not be deleted.
- All files within the Confluence Installation Directory will be deleted (with the exception of the Tomcat log folder located in the Confluence Installation Directory).
- The uninstaller can be made to operate in unattended mode by specifying the `-q` option at the Windows command prompt — i.e. `uninstall -q`
- If you wish to re-install Confluence in 'unattended mode', do not uninstall your previous installation of Confluence just yet. See [Using the Silent Installation Feature](#) for more information.

Installing Confluence on Linux

This guide describes how to install a new Confluence installation on Linux using the automated 'Linux Installer'. You can also install from a 'zip' archive — see [Installing Confluence on Linux from Archive File](#) for details. If you are upgrading Confluence, please see [Upgrading Confluence](#).

⚠ Please Note:

- It is possible that any anti-virus or other Internet security tools installed on your Linux operating system may interfere with the Confluence installation process and prevent the process from completing successfully. If you experience or anticipate experiencing such an issue with your anti-virus/Internet

- security tool, disable this tool first before proceeding with the Confluence installation.
- You may also wish to consider [disabling the Linux OutOfMemory Killer \(OOM Killer\)](#). This is a Linux function that can kill processes when memory on the server becomes low, and sometimes targets Confluence.
- Before you begin installing Confluence, please read the [System Requirements](#) page.

Using the Console Wizard

Use the console wizard if you are installing Confluence on your server for the first time or you wish to specify your installation options.

If you have previously installed Confluence using the installation wizard and wish to re-install Confluence again with the same installation options, you can re-install Confluence in 'unattended mode' without any user input required (see [below](#) for details).

1. Download and Install the Confluence 'Linux Installer'

✔ If you execute the Linux Installer with 'root' user privileges, the installer will create and run Confluence using a dedicated user account. You can also execute the Linux Installer without 'root' user privileges, although your installation options will be much more limited and a dedicated user account (to run Confluence) will not be created. To run Confluence as a service, the Linux Installer must be executed with 'root' user privileges.

- Download the appropriate **Confluence 'Linux 64-bit / 32-bit Installer'** (.bin) file from the [Confluence Download](#) page.

Please Note:

- To access the 32-bit installer, you may need to click the '**Show all**' link on the 'Confluence Download' page to access the other installation packages.
- The difference between the 64-bit / 32-bit .bin installers relates to their bundled Java platforms that run Confluence. Bear in mind that a Confluence installation installed using the 64-bit installer may require additional memory (to run at a similar level of performance) to a Confluence installation installed using the 32-bit installer. This is because a 64-bit Java platform's object references are twice the size as those for a 32-bit Java platform.

- Open a Linux console and change directory (cd) to the '.bin' file's directory.

⚠ If the '.bin' file is not executable after downloading it, make it executable, for example:

```
chmod a+x atlassian-confluence-X.Y.bin
```

(where X.Y represents your version of Confluence)

- Execute the '.bin' file to start the console wizard.

- When prompted to choose between '**Express Install**', '**Custom Install**' or '**Upgrade an existing Confluence installation**', choose either the '**Express Install**' or '**Custom Install**' options:

- Express Install** — If you choose this option, Confluence will be installed with default settings which are shown in the next step of the console wizard.

Please Note:

- If you are running the installer with 'root' user privileges, Confluence will be installed as a service.
- If you want to customise any of these options:
 - Enter '**e**' to exit the console wizard.
 - Execute the console wizard again (step 3 above).
 - Choose the '**Custom Install**' option instead.
- Custom Install** — If you choose this option, Confluence will prompt you to specify the following options (which are presented during subsequent steps of the console wizard and pre-populated with default values):
 - The 'Destination Directory' in which to [install Confluence](#).
 - The [Confluence Home directory](#) (which must be unique for each Confluence installation).
 - The *TCP ports* (i.e. an HTTP and a Control port) that Confluence will run through.
 - If you are running the installer with 'root' user privileges, you will be prompted to 'Run Confluence as a service' (recommended). You can also do this manually later, as described in [Start Confluence Automatically on Linux](#).

- The console wizard will install Confluence onto your operating system and will start Confluence automatically when the wizard finishes.

Please Note:

If you executed the Linux Installer with 'root' user privileges, the Linux Installer creates a dedicated Linux user account with username 'confluence', which is used to run Confluence. This account has only:

- Full write access to your [Confluence Home Directory](#).
- Limited write access to your [Confluence Installation Directory](#).

If you executed the Linux Installer without 'root' user privileges, be aware that Confluence can still be run with 'root' privileges. However, to protect the security of your operating system, *this is not recommended*.

2. Start Confluence

If Confluence is not already started, you can start Confluence using the appropriate [command](#) at the Linux console.

Once Confluence is started, you can access Confluence from a [browser](#) on any computer with network access to your Confluence server.

2.1 Starting and Stopping Confluence manually

In the Linux console, enter the bin subdirectory of your Confluence installation directory and execute the appropriate file:

- `start-confluence.sh` (to start Confluence)
- `stop-confluence.sh` (to stop Confluence)

Confluence will be ready to access (from a browser window) when the following message appears in the application's [log file](#):

```
*****
... You can now access Confluence through your web browser.
*****
```

2.2 Accessing Confluence from a Browser

You can access Confluence from any computer with network access to your Confluence server by opening a [supported web browser](#) on the computer and visiting this URL:

- `http://<computer_name_or_IP_address>:<HTTP_port_number>`

where:

- `<computer_name_or_IP_address>` is the name or IP address of the computer on which Confluence is installed and
- `<HTTP_port_number>` is the HTTP port number specified when you installed Confluence ([above](#)).

 If Confluence does not appear, you may need to [change the port that Confluence runs on](#).

Note: Application server logs (i.e. for Apache Tomcat) will be written to `logs/catalina.out`.

3. Run the Setup Wizard

See the [Confluence Setup Guide](#).

4. Next Steps

- See [Confluence 101](#).
- If you did not install Confluence to run as a service, you will need to [start Confluence manually](#) every time you restart your computer. To change your Confluence installation to run as a service, please see [Start Confluence Automatically on Linux](#).
- To get the most out of Confluence, please see [Performance Tuning](#).

Performing an Unattended Installation

If you have previously installed Confluence using the console wizard ([above](#)), you can use a configuration file from this Confluence installation (called `response.varfile`) to re-install Confluence in 'unattended mode' without any user input required.

Installing Confluence in unattended mode saves you time if your previous Confluence installation was used for testing purposes and you need to install Confluence on multiple server machines based on the same configuration.

 Please Note:

- The `response.varfile` file contains the options specified during the installation wizard steps of your previous Confluence installation. Hence, do not uninstall your previous Confluence installation just yet.
- If you intend to modify the `response.varfile` file, please ensure all directory paths specified are absolute, for example, `sys.installationDir=/opt/atlassian/confluence`
Unattended installations will fail if any relative directory paths have been specified in this file.

Download and Run the Confluence 'Linux Installer' in Unattended Mode

1. Download the **Confluence 'Linux Installer'** (.bin) file from the [Confluence Download Center](#) to a suitable location.
2. Open a Linux console.
3. Copy (cp) the file `.install4j/response.varfile` located in your previous Confluence installation directory, to the same location as the downloaded 'Linux Installer' file.
 You can [uninstall your previous Confluence installation](#) after this step. Save your `response.varfile` if you need to install Confluence on multiple machines.
4. Change directory (cd) to the location of the 'Linux Installer' file and execute the following command:

```
atlassian-confluence-X.Y.bin -q -varfile response.varfile
```

Where:

- X.Y — refers to the version of Confluence you are about to install.
 - -q — instructs the installer to operate in unattended mode (i.e. 'quietly').
 - -varfile `response.varfile` — specifies the configuration file containing the configuration options used by the installer. The location and name of the configuration file should be specified after the -varfile option.
5. Confluence will start automatically when the silent installation finishes. Continue from the step above, [Starting Confluence](#).

Installing Confluence on Linux from Archive File

 These instructions apply to:

- Confluence distributed as an archive file. The distribution includes Apache Tomcat as the application server.
- **Linux or Solaris** systems. If you are installing Confluence on a different system, please refer to [Installing Confluence](#).

Also, please check the version of Confluence which you are installing. Refer to the [documentation home page](#) to verify the latest Confluence version and to find documentation for older versions.



Hint: If you are evaluating Confluence on Solaris or you are unsure which version to install, this is the one to use.

On this page:

- [1. Before you Start](#)
- [2. Install Java](#)
- [3. Download and Extract the Confluence Installation File](#)
- [4. Define your Confluence Home Directory](#)
- [5. Check the Ports](#)
- [6. Select an External Database](#)

- [7. Start Confluence](#)
- [8. Confluence Setup Wizard](#)

1. Before you Start

Please check the following points:

1. Ensure that your system meets the minimum requirements to run Confluence. For more information, please read the detailed [System Requirements](#).
2. Have your Confluence license key ready. You can [obtain a trial, free or commercial license](#) now, or [retrieve your existing license key](#).
3. You must be able to use a command prompt and install Java to continue. If not, please contact your system administrator to assist you or consider the [Confluence Hosted evaluation option](#).
4. Make sure that you use a Gnu version of zip application - Solaris and AIX are known to have [problems](#) with zip, because they use their own (old) versions instead of the Gnu version.

2. Install Java

Please refer to the [Supported Platforms](#) for the required version of Java. (OpenJDK is currently **not** supported. A JIRA issue to request support for this JDK has [been created](#).)

1. If you are not sure whether you have Java installed correctly, please confirm by doing the following:
 - a. Open a shell console.
 - b. Type `echo $JAVA_HOME` in the shell console and then press Enter
 - c. View the result:
 - If a line is displayed such as `/opt/jdk1.6.0_12` or `/usr/lib/jvm/java-6-sun`, then Java is installed and properly configured.
 - If nothing is displayed, then you either need to install Java or set the `$JAVA_HOME` environment variable. You can set this environment variable in your user account's 'profile' file. Alternatively, you can set this after installing Confluence ([in step 4 below](#)) by defining this path in your Confluence installation's `setenv.sh` file, usually located in the Confluence `bin` directory.
 - If you have installed an unsupported JDK and you want to use SSL then you need to install the [Sun JSSE package](#).
2. If you need to install Java, follow these instructions:
 - Go to the [Java download page](#).
 - Download the latest JRE or JDK that is listed on the Confluence [Supported Platforms](#) page. (Confluence works with either the JDK or the JRE.)
 - When the download has finished, run the Java installer. Detailed installation instructions are provided on [Oracle's website](#).

Note: you will be asked to choose an installation directory. Make a note of this directory for use later.

3. Download and Extract the Confluence Installation File

1. If you have not downloaded Confluence already, [download the TAR.GZ file](#).

Use your unzip program to unzip the installation file to a directory such as `/home/jsmith/confluence-2.7.0-std/`.

Most Linux/Solaris users can use any unzip program (such as [GNU Tar](#)) to extract the Confluence installer. However, Solaris users should not use the Solaris Tar program due to a [known issue](#) associated with its use in extracting Confluence. Use another application such as GNU Tar instead.

For example, change directory to your home directory in Linux and enter the following commands in the shell console:

- `gunzip confluence-<version>-std.tar.gz`
 - `tar -xf confluence-<version>-std.tar`
- (where `<version>` refers to the Confluence version you downloaded.)

 As usual on Linux/Solaris-based operating systems, avoid using spaces in your directory path. The directory into which you unzipped the Confluence installation is called the [Confluence Installation directory](#). Next you will define the [Confluence Home directory](#).

4. Define your Confluence Home Directory

Now you need to define the [Confluence Home directory](#). This is where Confluence will store its configuration information, indexes and attachments.

✔ Tip: Another term for 'Home directory' would be 'data directory'.

We suggest using different paths for your installation and home directories. This will facilitate upgrades. Examples of Installation and Home Directories:

- Installation directory: `/usr/local/confluence/`

If you wish to install or maintain multiple versions of Confluence, you can add a version number to the Confluence installation directory name like `/usr/local/confluence-3.1-std/` and optionally, create the symbolic link `/usr/local/confluence/` that points to `/usr/local/confluence-3.1-std/`

- Home directory: `/usr/local/confluence-data/`

1. Open your Confluence Installation directory (created when you unzipped Confluence — see [above](#)).
2. Under the Installation directory, find this file: `confluence/WEB-INF/classes/confluence-init.properties`
3. Open the `confluence-init.properties` file in a text editor.
4. Scroll to the bottom and find this line:

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

5. Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with `confluence.home`
6. If you decide to change the Confluence Home directory from the default, use an absolute path rather than a symbolic link to specify the path and file name. For example:

```
confluence.home=/home/jsmith/confluence-data/
```

5. Check the Ports

If you have another application running on your machine which is using the same ports that Confluence uses by default, you may need to change the port which Confluence will use. For example, if you have a installation of [JIRA](#) running on this machine, JIRA might be already using the port which Confluence requests by default.

By default, Confluence listens on port '8090'. If this port is already in use in your installation, follow these instructions to change the ports:

- To change the ports for Confluence, open the file `conf/server.xml` under your [Confluence Installation directory](#). The first four lines of the file look like this:

Default conf/server.xml

```
<Server port="8000" shutdown="SHUTDOWN" debug="0">
  <Service name="Tomcat-Standalone">
    <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
port="8090" minProcessors="5" maxProcessors="75"
      enableLookups="true" redirectPort="8443" acceptCount="10"
debug="0" connectionTimeout="20000" useURValidationHack="false"/>
    ...
```

You need to modify both the **server** port (default is 8000) and the **connector** port (default is 8090) to ports that are free on your machine. The server port is required by Tomcat but is not user facing in any way. The connector port is what your users will use to access Confluence, eg in the snippet above, the URL would be <http://example.com:8090>.

✔ **Hint:** You can use `netstat` to identify free ports on your machine. See more information on using `netstat` on Windows or on Linux.

For example, here are the first four lines of a modified `server.xml` file, using ports '8020' and '8099':

Modified `conf/server.xml` using ports 8020 and 8099

```
<Server debug="0" shutdown="SHUTDOWN" port="8020">
  <Service name="Tomcat-Standalone">
    <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
      port="8099" minProcessors="5" maxProcessors="75"
      enableLookups="true" redirectPort="8443" acceptCount="10"
      debug="0" connectionTimeout="20000" useURIVValidationHack="false"/>
    ...
```

To access Confluence in this configuration, point your web browser to <http://localhost:8099/>.



You should also ensure at this point that if you are using a firewall, it is configured to allow `http/https` traffic over the port you have chosen.

Once this is working, if this is the URL your users will use, remember to [update your Base URL](#) to point to the new URL.

You will find more information on [this page](#).

6. Select an External Database

This step is optional for users evaluating Confluence. However, if you are installing Confluence for production purposes, this step is mandatory. Please refer to the [database](#) requirements listed on our [System Requirements](#) topic for help in choosing an external database.

i External databases are those listed on our [Supported Platforms](#) topic, excluding HSQLDB, which is bundled with Confluence and should not be used in production.

When you have chosen your external database, follow the appropriate [database setup guide](#) to set up your database to work with Confluence.

You can learn more about migration from an existing installation or use of the evaluation database [here](#). You will continue to use the Database Setup Guide during the Confluence Setup Wizard. (See [step 8](#) below.)

7. Start Confluence

1. Go to your Confluence Installation directory (created when you unzipped Confluence — see [above](#)).
2. Under your Confluence Installation directory, open the `bin` directory and run the startup script: `start-confluence.sh`.
3. Once Confluence is running, open a web browser and visit <http://localhost:8090/>.



Hint: If you changed the port earlier, use the port you specified in [step 6](#) above.

8. Confluence Setup Wizard

The Confluence Setup Wizard should appear in your web browser, prompting you to enter your license key. Follow the instructions on the screens, and read more guidelines on the [Confluence Setup Wizard](#).

Related Topics

[Change listen port for Confluence](#)

[Running Confluence Over SSL or HTTPS](#)

[Confluence Setup Guide](#)

[Configuring Confluence](#)

[Documentation Home](#)

[Uninstalling Confluence from Linux](#)

This page describes the procedure for uninstalling Confluence, which had been installed using the [Linux Installer](#).

To uninstall Confluence from Linux:

1. Open a Linux console.
2. Change directory (cd) to your Confluence installation directory.
3. Execute the command `uninstall`. This command must be executed as the same user account that was used to install Confluence with the [Linux Installer](#).
4. Follow the prompts to uninstall Confluence from your computer.

Please note:

- The uninstaller will not delete the Confluence Home Directory.
- All log files that were generated while Confluence was running will not be deleted.
- All files within the Confluence Installation Directory will be deleted (with the exception of the Tomcat log folder located in the Confluence Installation Directory).
- The uninstaller can be made to operate in unattended mode by specifying the `-q` option — i.e. `uninstall -q`
- If you wish to re-install Confluence in 'unattended mode', do not uninstall your previous installation of Confluence just yet. See [Using the Silent Installation Feature](#) for more information.

Change listen port for Confluence

Problem

This page tells you what to do if you get errors like the following when starting Confluence, when you can't access Confluence on port **8090**.

If you see this error:

```
java.net.BindException: Address already in use: JVM_Bind:8090
```

This means you are running other software on Confluence's default port of **8090**. This may be another other process running on the same port. It may also be a previous instance of Confluence that hasn't been shut down cleanly.

To find out what process is listening on that port, load a command prompt and type: **netstat -an**

```
-a : Displays all active TCP connections and the TCP and UDP ports on which the computer is listening.  
-n : Displays active TCP connections, however, addresses and port numbers are expressed numerically and no attempt is made to determine names.
```

There is also [Process Explorer tool](#) available to determine what is binding port **8090**.

Solution: Change the Ports which Confluence Listens On

To change the ports for Confluence, open the file `conf/server.xml` under your [Confluence Installation directory](#). The first four lines of the file look like this:

Default conf/server.xml

```
<Server port="8000" shutdown="SHUTDOWN" debug="0">
  <Service name="Tomcat-Standalone">
    <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
      port="8090" minProcessors="5" maxProcessors="75"
      enableLookups="true" redirectPort="8443" acceptCount="10" debug="0"
      connectionTimeout="20000" useURValidationHack="false"/>
    ...
  </Service>
</Server>
```

You need to modify both the **server** port (default is 8000) and the **connector** port (default is 8090) to ports that are free on your machine. The server port is required by Tomcat but is not user facing in any way. The connector port is what your users will use to access Confluence, eg in the snippet above, the URL would be <http://example.com:8090>.

✔ Hint: You can use [netstat](#) to identify free ports on your machine. See more information on using [netstat](#) on Windows or on Linux.

For example, here are the first four lines of a modified `server.xml` file, using ports '8020' and '8099':

Modified conf/server.xml using ports 8020 and 8099

```
<Server debug="0" shutdown="SHUTDOWN" port="8020">
  <Service name="Tomcat-Standalone">
    <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
      port="8099" minProcessors="5" maxProcessors="75"
      enableLookups="true" redirectPort="8443" acceptCount="10" debug="0"
      connectionTimeout="20000" useURValidationHack="false"/>
    ...
  </Service>
</Server>
```

To access Confluence in this configuration, point your web browser to <http://localhost:8099/>.



You should also ensure at this point that if you are using a firewall, it is configured to allow http/https traffic over the port you have chosen.

Once this is working, if this is the URL your users will use, remember to [update your Base URL](#) to point to the new URL.

NOTES

[1] For more information on [netstat](#), see [using netstat on Windows](#), or [netstat man page](#) (Linux).

[2] The JIRA distribution runs on port **8080** by default. If you're looking to change the port of the JIRA distribution, see [Changing JIRA Standalone's port](#).

RELATED PAGES

[Installing Confluence Documentation Home](#)

Installing the Confluence EAR-WAR Edition

The Confluence [EAR-WAR](#) distribution is intended for deployment into an existing J2EE application server.

To use this method of installation, you need to know how to deploy a web application on an existing application server. If not, please use the [Confluence distribution](#) instead.

On this page:

- [Step 1. Check the System Requirements and Known Issues](#)

- [Step 2. Download and Extract EAR-WAR Installation File](#)
- [Step 3. Review Application Server Memory Allocation](#)
- [Step 4. Configure confluence-init.properties](#)
- [Step 5. Edit Tomcat Context Descriptors](#)
- [Step 6. Add UTF-8 Encoding](#)
- [Step 7. \(Optional\) Configure Tomcat to Run on a Different Port](#)
- [Step 8. \(Optional\) Configure Confluence to Run as a Windows Service](#)
- [Step 8. Run the Confluence Setup Wizard](#)
- [Notes](#)

Step 1. Check the System Requirements and Known Issues

1. Please check the [Confluence system requirements](#).
2. In addition to the above requirements, the EAR-WAR distribution requires the Apache Tomcat application server. For more information on Confluence's supported application servers, please refer to our [Supported Platforms](#) page.
3. If deploying as an unexploded WAR, Ant 1.3 or later is required. This is bundled with the WAR download.
4. Confluence, the database and application server must [use the same character encoding](#). UTF-8 is recommended.
5. Deploying multiple Atlassian applications in a single Tomcat container is **not supported**. We do not test this configuration and upgrading any of the applications (even for point releases) is likely to break it. There are also a number of known issues with this configuration (see [this FAQ](#) for more information).

We also do not support deploying multiple Atlassian applications to a single Tomcat container for a number of practical reasons. Firstly, you must shut down Tomcat to upgrade any application and secondly, if one application crashes, the other applications running in that Tomcat container will be inaccessible.

Finally, we recommend not deploying *any other applications* to the same Tomcat container that runs Confluence, especially if these other applications have large memory requirements or require additional libraries in Tomcat's `lib` subdirectory.

6. Read through the [Known Issues for Apache Tomcat](#).

Step 2. Download and Extract EAR-WAR Installation File

This section gives detailed instructions for installing Confluence EAR-WAR edition on an **Apache Tomcat 5.5, or 6** server.

1. Download the [Confluence EAR/WAR zip file](#). (You need to click the '**Show all**' link to see the EAR/WAR zip file.)
2. Please check your unzip program before extracting the downloaded zip file. Some archive-extract programs cause errors when unzipping the Confluence zip file:
 - Windows users must avoid the Windows built-in unzip utility, as it doesn't extract all the files. Use a third-party unzip program like [7Zip](#) or [Winzip](#).
 - Solaris users will need to use [GNU tar](#) to handle the long file names.
3. Extract the downloaded zip file.
4. You have now unzipped your *Confluence installation directory*, which should contain the version number e.g. `confluence-4.0.1` or `confluence-4.0.2`. This directory will be later referred to as the *Confluence installation directory*. Inside is a `confluence` subdirectory, referred to later as the (*Exploded*) *Confluence WAR directory*. Record the absolute path to the *Confluence WAR directory*.

Step 3. Review Application Server Memory Allocation

Confluence requires a maximum heap allocation (Xmx) of at least 256 MB for normal operation. Also, remember to set the maximum PermGen memory allocation (XX:MaxPermSize). See [Increasing Application Server Memory](#).

 Do not configure a heap allocation so large that it does not allow enough remaining physical memory for your operating system and other applications on the server. The heap allocation should be large enough for Confluence, but not so large that the memory would be [paged to disk](#) during normal operation.

Step 4. Configure confluence-init.properties

1. Inside the *Confluence installation directory*, edit `...confluence/WEB-INF/classes/confluence-init.properties` in a text editor.
2. Now define your Confluence Home directory, by setting the `confluence.home` property to a directory of your choosing.

We suggest using different paths for your installation and home directories. This will facilitate upgrades. This is the directory that will contain all of Confluence's configuration, backup and attachment files.

✓ Tip: Another term for 'Home directory' would be 'data directory'.

⚠ Make sure the user that runs Tomcat has full write access to the Confluence Home directory

Step 5. Edit Tomcat Context Descriptors

1. Create a file called `confluence.xml` and save in the `conf/Catalina/localhost` sub-directory of Tomcat. If these directories don't exist you can create them manually.
2. Open your new `confluence.xml` file and add these lines:

```
<Context path="/confluence"
docBase="<CONFLUENCE_INSTALLATION_DIRECTORY_PATH>/confluence" debug="0"
reloadable="true">
</Context>
```

More on Context Path

To run Confluence *without* a context path, change the path in the Context tag to an empty string (`""`). **If not using a context path, your config will need to be saved as `R00T.xml` rather than `confluence.xml`.**

In Tomcat, a context path name follows the name of its xml file (except for `R00T.xml` where no context path is used). Hence if you wish to change the context path to a different name, change both the context path and the name of the xml file. eg. `"/wiki"` context path should be saved in file `wiki.xml`.

3. For `docBase`, specify the value you noted down earlier.
4. Restart Tomcat, and Confluence should be accessible under `/confluence/` on your Tomcat server.
5. Follow the link below to proceed with the setup wizard.

Step 6. Add UTF-8 Encoding

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="8080"/>
```

2. Add a `URIEncoding="UTF-8"` property to the connector:

```
<Connector port="8080" URIEncoding="UTF-8"/>
```

Step 7. (Optional) Configure Tomcat to Run on a Different Port

See [Switching to Apache Tomcat](#).

Step 8. (Optional) Configure Confluence to Run as a Windows Service

Confluence can be [run as a service](#).

Step 8. Run the Confluence Setup Wizard

Once Confluence is running, open a web browser and visit <http://localhost:8080/> (Tomcat default port).

If you changed the port earlier, use the port you specified. ⓘ Note that the Confluence installer normally uses port 8090 as the default, to avoid conflicts with JIRA (using port 8080).

The Confluence Setup Wizard should appear in your web browser, prompting you to enter your license key. Follow the instructions on screen, and read more guidelines on the [Confluence Setup Wizard](#).

Notes

- Tomcat users, take care not to unzip the Confluence installation into your Tomcat webapps folder, as this may cause Confluence to be deployed more than once. It may cause a [Cluster Panic](#) error.
- If you deploy Confluence on an unsupported server, server-related issues cannot be covered by Atlassian technical support. You can try [Atlassian Answers](#) for assistance instead.

Known Issues for Apache Tomcat

On this page:

- [Supported Application Servers](#)
- [Tomcat Documentation](#)
- [Known Issues](#)

Supported Application Servers

Check the list of supported application servers on the [Supported Platforms](#) topic.

Tomcat Documentation

An excellent resource for Tomcat configuration is the [Apache documentation](#).

Known Issues

- 📖 [Setup Fails Creating MySQL Schema Due to Tomcat Incompatibility](#)
- 📖 [Unable to Install Service on Windows Vista](#)
- 📖 [Confluence will not update or install add-ons](#)
- 📖 [Installing UPM 2.7.2 or later on versions of Tomcat prior to 6 fails](#)
- 📖 ['All threads \(150\) are currently busy, waiting. Increase maxThreads \(150\) or check the servlet status' Due to High Volume Transactions](#)
- 📖 [Unable to Start Tomcat after Confluence User Management Delegation to JIRA](#)
- 📖 [Fix 'Not supported by BasicDataSource' Setup or Startup Error](#)
- 📖 [Confluence Does Not Start due to 'Error deploying configuration descriptor'](#)
- 📖 [Tomcat 6.0.26 or higher Shutdown Reports 'A web application created a ThreadLocal ThreadLocal has been forcibly removed'](#)
- 📖 [Confluence Can't Start and Doesn't Create Logfiles due to CATALINA_HOME Being Set](#)
- 📖 [Confluence Startup Referencing a Different Tomcat](#)
- 📖 [Unable to Enable Workbox's Notifications and Tasks - Host Plugin Due to NoClassDefFoundError](#)
- 📖 [NotSerializableException on Shutdown](#)
- 📖 [Login Fails After Upgrade](#)
- 📖 [Confluence Deadlocks when Running under Tomcat 6.0.24-6.0.32](#)
- 📖 [Tomcat fails to start with "The system could not find the environment option that was entered."](#)
- 📖 [HTML Macros Fail after Upgrading to 3.4 or Later Due to External URL References to Local Resources](#)
- 📖 [Universal Plugin Manager stops working after upgrade to v4.3.5 or v4.3.6](#)
- 📖 [Confluence Does Not Start Due to NullPointerException in FelixOsgiContainerManager](#)
- 📖 ["NoSuchMethodError: javax.servlet.ServletContext.getContextPath\(\)" when starting Confluence](#)
- 📖 [How to Determine Your Version of Tomcat and Java](#)
- 📖 [Application Servers Troubleshooting](#)
- 📖 [Confluence Menus Do Not Work, or Confluence Fails to Start when Running in the Same Application Server](#)

as JIRA 4.0, 4.0.1 or Crowd 2.0.x

-  [Unable to Configure Confluence to Run as a Service on Tomcat 5](#)
-  [Installation or Upgrade of Confluence 4.0 EAR-WAR Fails on Red Hat or CentOS](#)
-  [Slow Page Rendering of Large Pages Due to HTTP POST Limitations](#)

RELATED TOPICS

[Running Confluence behind Apache](#)

[Configuring a MySQL Datasource in Apache Tomcat](#)

Installing Java for Confluence

This page contains instructions for installing a Java Development Kit (JDK). This is a manual step that is only required for Confluence installations where you are installing from a zip or archive file.

 If you are using the [automated installer](#), the required Java files are bundled and will be automatically put in place, hence you will not need to follow the instructions on this page.

Please refer to our [Supported Platforms](#) topic for details of the Java versions that are supported for Confluence.

Installing the JDK

A JDK (Java Development Kit) needs to be installed on the same server machine that will have Confluence installed.

▼ [For Windows: \(click to expand\)](#)

Installing the JDK on Windows

1. If you are not sure whether you have a JDK installed, please confirm by doing the following:
 - Check **Control Panel > Programs and Features** in Windows 7 (just Programs on older version of Windows).
 - Java should appear as a line item in the list. If not, you do not have Java installed.
2. To install the JDK, follow these instructions:
 - Go to the [Java download page](#).
 - Download the version entitled 'Java SE Update XX (JDK)', where 'XX' stands for some number. (The latest version will be available on that page.)
 - When the download has finished, run the Java installer. At one point, you will be asked to choose a directory to install to. Copy or write this directory down for use later.
3. Check that the JAVA_HOME environment variable has been set correctly.
 - Open the **Start** menu, choose **Run**, type cmd in the **Run** dialog box and click **OK**.
 - In the command prompt window, type echo %JAVA_HOME% and then press *Enter*.
 - View the result:
 - If a directory path is displayed that looks similar to one of the following examples, with the letters 'JDK' immediately preceding a series of version numbers, and this path matches the location where you installed the JDK in step 2, then your JDK has been successfully installed and your JAVA_HOME environment variable has been set correctly.
Examples of typical JAVA_HOME environment variable values:
 - C:\Program Files\Java\JDK7
 - C:\Progra~1\Java\JDK7
 - C:\Java\JDK7
 - C:\JDK7
 - If nothing is displayed or you do not see 'JDK' immediately followed by a series of version numbers (like one of the examples above), then you need to set the JAVA_HOME environment variable. Please follow [these instructions](#) to set your JAVA_HOME environment variable to the directory you where you have just installed the JDK. By default, this directory is under C:\Program Files\Java.

Note: Any Java or JDK version numbers on this page are **examples only**. Please refer to the [Supported Platforms](#) page for the supported versions of Java.

▼ [For Linux: \(click to expand\)](#)

Installing the JDK on Linux

1. If you are not sure whether you have JDK installed correctly, please confirm by doing the following:
 - a. Open a shell console.
 - b. Type `echo $JAVA_HOME` in the shell console and then press Enter
 - c. View the result:
 - If a line is displayed such as `/opt/JDK7` or `/usr/lib/jvm/java-7`, then your JDK is installed and properly configured.
 - If nothing is displayed, then you either need to install the JDK or set the `$JAVA_HOME` environment variable. You can set this environment variable in your user account's 'profile' file. Alternatively, you can set this after installing Confluence, by defining this path in your Confluence installation's `setenv.sh` file, usually located in the Confluence `bin` directory.
 - If you have installed [an unsupported JDK](#) and you want to use SSL then you need to install the [Sun JSSE package](#).
2. If you need to install the JDK, follow these instructions:
 - Go to the [Java download page](#).
 - Download the version entitled 'Java SE Update XX (JDK)', where 'XX' stands for some number. (The latest version is available on that page.)
 - When the download has finished, run the Java installer. Detailed installation instructions are provided on [Oracle's website](#).

Note: Any Java or JDK version numbers on this page are **examples only**. Please refer to the [Supported Platforms](#) page for the supported versions of Java.

Setting the JAVA_HOME Variable in Windows

 This information is only relevant if you are installing Confluence on a Windows server.

After you have [installed the Java Runtime Environment \(JRE\)](#) in Windows, you must set the `JAVA_HOME` environment variable to point to the JRE installation directory.

Stage 1. Locate the JRE Installation Directory

If you already know the installation path for the Java Runtime Environment, go to Stage 2 below. Otherwise, find the installation path by following these instructions:

1. If you didn't change the installation path for the Java Runtime Environment during installation, it will be in a directory under `C:\Program Files\Java`. Using Explorer, open the directory `C:\Program Files\Java`.
2. Inside that path will be one or more subdirectories such as `C:\Program Files\Java\jre6`.

Stage 2. Set the JAVA_HOME Variable

Once you have identified the JRE installation path:

1. Right-click the **My Computer** icon on your desktop and select **Properties**.
2. Click the **Advanced** tab.
3. Click the **Environment Variables** button.
4. Under **System Variables**, click **New**.
5. Enter the variable name as `JAVA_HOME`.
6. Enter the variable value as the installation path for the Java Development Kit.
 - If your Java installation directory has a space in its path name, you should use the shortened path name (e.g. `C:\Progra~1\Java\jre6`) in the environment variable instead.

Note for Windows users on 64-bit systems

Progra~1 = 'Program Files'
Progra~2 = 'Program Files(x86)'

7. Click **OK**.
8. Click **Apply Changes**.
9. Close any command window which was open before you made these changes, and open a new command window. There is no way to reload environment variables from an active command prompt. If the changes do not take effect even after reopening the command window, restart Windows.
10. If you are running the Confluence EAR/WAR distribution, rather than the regular Confluence distribution,

you may need to restart your application server.

Related Topics

[Starting Tomcat as a Windows Service](#)

[Installing Confluence in Linux](#)

Confluence Cluster Installation

Request for interest

Hi there,

Great to see your interest in Confluence Clustered!

We are currently working on a completely new version of our clustered offering that provides greater scalability as well as High Availability.

If you are interested in joining our closed beta program for the new offering, please, send an email to kelvin@atlassian.com with the subject: **Confluence Clustered** and I will get back to you ASAP.

Cheers,

Kelvin Yap
Enterprise Advocate

Overview

There are two methods of installing Confluence in a cluster, depending on whether you have existing data. **This page describes a fresh installation with no existing data.**

See also [Confluence Cluster Installation with Existing Data](#).

Oracle Coherence Licensing Change:

- Due to a license agreement change, Confluence is now available in two editions:
 - **Standard Edition** — Confluence with Ehcache's caching technology (available to customers with non-clustered Confluence licenses).
 **If you are currently running a clustered installation of Confluence, please do not upgrade it with a standard edition of Confluence.**
 - **Clustered Edition** — Confluence with Oracle's Coherence clustering and distributed caching technology (available to customers with Confluence clustered licenses only).
- For more information about these changes, please refer to the [Coherence License Changes](#) document.
- If you have a Confluence clustered license, are running a clustered installation of Confluence and wish to upgrade to Confluence version 2.6 or later, please ensure that you download only a [clustered edition of Confluence](#).

Installation with no existing data

To get Confluence running in a two-node cluster, you must do the following:

1. Ensure you meet the clustering [requirements](#), including obtaining a clustered license key from Atlassian for each node.
2. Install Confluence on a single node, configuring an external database and a cluster name.
3. Load test the single node installation, see whether clustering is required.
4. Shut down the first node, copy the Confluence application and Confluence home directory to the second node.
5. Start the first node, wait until it is running, then bring up the second node and it will automatically join the cluster.
6. Test the cluster is working correctly.
7. Configure a load balancer in front of the two clustered nodes.

Each of these steps will be described in detail below.

1. Clustering requirements

Your Confluence cluster installation must meet *all* the following criteria for clustering:

- You must have a clustered license.
- You must use an external database.
- You must use a load balancer with session affinity in front of the cluster.

Clustered commercial licenses may be purchased through [Confluence website](#). Clustered evaluation licenses may be obtained by emailing sales@atlassian.com.

A cluster can run using two copies of Confluence. However, cluster administrators must understand how to configure an application server and web server with load balancing, so we recommend you are comfortable installing Confluence as a EAR/WAR in your application server before proceeding with a clustered installation.

2. Installation on first node

Cluster administrators should already be comfortable with the [normal installation method](#), so it won't be repeated here. There are two differences in the Confluence Setup Wizard from a normal installation:

- You must use an external database.
- You must enter a cluster name.



Enter a cluster name to create a new cluster



Technical note

The cluster name will be converted into a unique multicast IP address and port for your Confluence cluster. UDP multicast traffic is used for Confluence to automatically discover other nodes in the cluster when they start up.

3. Load test the single node

Most Confluence installations do not need to be clustered. Ensure you have tested your single node installation with the number of users you expect to host before going ahead with the additional complexity of clustering.

Check out our [performance tuning tips](#) for ways to improve the performance of a single instance of Confluence.

You can upgrade your single node to a multi-node cluster at any time by resuming this guide from step 4 below.

4. Copy Confluence to second node

Confluence clusters must use the same JDK, application server and application. The easiest way to ensure this is to shut down Confluence on the first node, then copy its web application and home directory to the second node:

1. Shut down Confluence on node #1.
2. Shut down your application server on node #2, or stop it automatically loading web applications.
3. Copy the Confluence web application from node #1 to node #2.
4. Copy the Confluence home directory from node #1 to node #2.
5. If the node #1 and node #2 filesystem structures are different, update the `/confluence/WEB-INF/classes/confluence-init.properties` file on in the web application directory of node #2 to point to the Confluence home directory path on node #2.

Copying the web application ensures any modifications you have made to the application itself, custom LDAP settings (`atlassian-user.xml`), and any other advanced configuration are copied to node #2.

Copying the home directory ensures the Confluence search index (the `index/` directory), the database and cluster configuration (`confluence.cfg.xml`), and any other home directory settings are copied to node #2.

5. Start Confluence on the first node, wait, then start Confluence on second node


For the most stable start-up process, it is important to start Confluence one server at a time.

1. Start Confluence on node #1.
2. Wait for Confluence to become available on node #1.
3. Start Confluence on node #2.
4. Wait for Confluence to become available on node #2.

6. Test cluster connectivity

The [Cluster Administration page](#) (Administration, Cluster Configuration) includes information about the active cluster. When the cluster is running properly, this page displays:

- a correct count of the nodes in the cluster
- a status display for each node in the cluster
- an uptime for each node that is accurate.



Cluster Administration page

A simple process to ensure your cluster is working correctly is:

1. Create a new document on node #1.
2. Ensure the new document is visible by accessing it directly on node #2.
3. Wait one minute (Confluence does batch indexing once per minute).
4. Search for the new document on node #1, ensure it appears.
5. Search for the new document on node #2, ensure it appears.



Technical note

If Confluence detects more than one instance accessing the database but not in a working cluster, it will shut itself down in a *cluster panic*. This can be fixed by [troubleshooting the network connectivity of the cluster](#).

7. Configure load balancer

For the moment, configuring the load balancer is outside the scope of this document.

However, a simple [Apache and Tomcat load-balancing configuration](#) is available, which includes sample configuration for the Apache Tomcat and the Apache web server, using its load-balancing JK connector.

Troubleshooting

If you have problems with the above procedure, please see our [Cluster Troubleshooting guide](#).

Upgrading a cluster

It is important that upgrades follow the procedure for [Upgrading a Confluence Cluster](#).

Related documentation

[Overview of Confluence Clusters](#)
[Clustering in Confluence](#)
[Confluence Cluster Installation with Existing Data](#)
[Confluence Installation Guide](#)
[Upgrading a Confluence Cluster](#)
[Cluster Administration page](#)

Confluence Cluster Installation with Existing Data

Overview

There are two methods of installing Confluence in a cluster, depending on whether you have existing data. **This page describes how to upgrade an existing Confluence instance into a cluster.**

See also [Cluster installation without existing data](#).

Oracle Coherence Licensing Change:

- Due to a license agreement change, Confluence is now available in two editions:
 - **Standard Edition** — Confluence with Ehcache's caching technology (available to customers with non-clustered Confluence licenses).
 **If you are currently running a clustered installation of Confluence, please do not upgrade it with a standard edition of Confluence.**
 - **Clustered Edition** — Confluence with Oracle's Coherence clustering and distributed caching technology (available to customers with Confluence clustered licenses only).
- For more information about these changes, please refer to the [Coherence License Changes](#) document.
- If you have a Confluence clustered license, are running a clustered installation of Confluence and wish to upgrade to Confluence version 2.6 or later, please ensure that you download only a [clustered edition of Confluence](#).

Cluster installation from an existing copy of Confluence

BEFORE ATTEMPTING THIS, PLEASE MAKE A BACKUP. To upgrade an existing copy of Confluence to run in a two-node cluster, you must do the following:

1. Ensure that your version of the Confluence distribution has been upgraded to the version you want to run the Cluster on. **Do not upgrade your version of Confluence and switch to the clustered version at the same time.** First upgrade your system (e.g. from Confluence 2.5.8 to 2.7.1) and make sure everything works fine (e.g. for a week) before switching (e.g. from Confluence 2.7.1 to 2.7.1 Clustered)
2. Ensure you meet the clustering requirements, including obtaining a clustered license key from Atlassian for each node
3. Due to [CONF-8959](#), you need to perform [attachment migration](#) to the database before you change your license to a clustered license
4. Upgrade the existing Confluence instance to a clustered license. Do this by going to Admin> Licence Details. Confluence should warn you that this version of Confluence is not capable of clustering.
5. Shutdown Confluence. Deploy a clustered version of Confluence (Do not attempt to install any version of Confluence that is not the Clustered equivalent to your current release). Edit `confluence-init.properties` (`confluence-ver-clustered/confluence/WEB-INF/classes/confluence-init.properties`) to set `confluence.home` to the same path as the old home. Start the first node, and verify that things are working correctly.
6. Shut down the first node, copy the Confluence application and Confluence home directory to the second node
7. Start the first node, wait until it is running, then bring up the second node and it will automatically join the cluster
8. Test the cluster is working correctly
9. Configure a load balancer in front of the two clustered nodes.

Each of these steps will be described in detail below.

1. Clustering requirements

Your Confluence cluster installation must meet *all* the following criteria for clustering:

- you must be running Confluence 2.3 or later
- you must have a clustered license
- you must use an external database
- you must use a load balancer with session affinity in front of the cluster.

Clustered commercial licenses may be purchased through [Confluence website](#). Clustered evaluation licenses may be obtained by emailing sales@atlassian.com.

A cluster can run using two copies of the Confluence distribution. However, cluster administrators must understand how to configure an application server and web server with load balancing, so we recommend you are comfortable installing Confluence as a EAR/WAR in your application server before proceeding with a clustered installation.

You can follow the instructions to [Migrate Confluence to an external database](#).

2. Upgrade existing instance to clustered license

Once you've obtained your clustered license from Atlassian, you can simply update the license in your running Confluence instance:

1. Go to 'Administration'.
2. Go to 'License Details', and paste in the new license.
3. Click 'Save'.

When you enter a clustered license, you will see a new line appear on this page: *Licensed Clustered Nodes*. This tells you how many nodes your Confluence license will allow.

Organisation	Atlassian
Date Purchased	Aug 15, 2006
License Type	Confluence: Commercial Server
Licensed Users	Unlimited
Licensed Clustered Nodes	8 nodes (2 nodes currently clustered).

License Details page shows the number of cluster nodes permitted

3. Migrate your attachments to the Database

You can do this by navigating to Admin> Attachment Storage > Edit, and changing it to "Database".

4. Copy Confluence to second node

For the remaining steps in setting up a cluster with existing data, please continue from [step 4 in the normal Confluence cluster installation guide](#).

5. Start Confluence on the first node, wait, then start Confluence on second node

See comment in step 4.

6. Test cluster connectivity

See comment in step 4.

7. Configure load balancer

See comment in step 4.

Troubleshooting

If you have problems with the above procedure, please see our [Cluster Troubleshooting guide](#).

Upgrading a cluster

It is important that upgrades follow the procedure for [Upgrading a Confluence Cluster](#).

Related documentation

[Overview of Confluence Clusters](#)

[Confluence Cluster Installation](#)

[Confluence Installation Guide](#)

[Upgrading a Confluence Cluster](#)

[Confluence User Guide](#)

Apache and Tomcat load balancing

Overview

The following is a description of how to set up a Confluence Cluster on a Windows machine using Apache and `mod_jk` to handle the load-balancing.

The characteristics of this cluster are:

- Session affinity: sessions are associated with single servers.
- Failover: if a server dies, a connection will be directed to the nearest available server. (NOTE: sessions are not replicated)
- Failback: when a server comes back online, it will rejoin the cluster.
- Weighted load balancing: the load balancing can be controlled to take into account machine differences. (See the `mod_jk` documentation for details on this.)

What do you need?

1. Download and install one copy of [Apache httpd](#). Do not install Apache as a service, but set it to listen on port 8080. (Tested with Apache httpd 2.0.55.)
2. Download the latest version of [mod_jk](#). Copy this file into the Apache `modules/` directory and rename it to `mod_jk.so`. (Tested with JK-1.2.19.)
3. Download and extract one copy of the [ZIP distribution of Apache Tomcat](#). (Tested with Tomcat 5.5.)
4. Download JDBC drivers for the external database you will be using. Put the drivers in Tomcat's `common/lib/` directory. (Tested with PostgreSQL 8.1, `postgresql-8.1-404.jdbc3.jar`).

Apache configuration

Edit the main Apache config file, `conf/httpd.conf`:

- add the following immediately after the other `LoadModule` directives:

```
LoadModule jk_module modules/mod_jk.so
```

- add the following just before the end of the file:

```
JkWorkersFile conf/workers.properties
JkLogFile logs/mod_jk.log
JkLogLevel info

JkMount /confluence loadbalancer
JkMount /confluence/* loadbalancer
```

Create a `workers.properties` file in the Apache `conf/` directory. This version of the `workers.properties` file is configured to use 2 Tomcat instances: `tomcat1` and `tomcat2`.


```
worker.list=loadbalancer

worker.tomcat1.port=18081
worker.tomcat1.host=localhost
worker.tomcat1.type=ajp13
worker.tomcat1.lbfactor=1

worker.tomcat2.port=28081
worker.tomcat2.host=localhost
worker.tomcat2.type=ajp13
worker.tomcat2.lbfactor=1

worker.loadbalancer.type=lb
worker.loadbalancer.balance_workers=tomcat1, tomcat2
worker.loadbalancer.method=Busyness
```

Tomcat configuration

The Tomcat configuration below will run multiple instances from the same binaries in the main Tomcat directory. For complete documentation of this configuration, see the `RUNNING.txt` file in the Tomcat distribution.

Create instance home directories

Create a directory for each instance of Tomcat, somewhere outside where you installed Tomcat. For example, if you extracted Tomcat to `/opt/apache/tomcat-5.5`, your instances could be in `/var/tomcat-instances/tomcat1`, `/var/tomcat-instances/tomcat2`. These folders will be referred to as the *instance home directories*.

Copy the following folders from the Tomcat installation directory into each instance home directory. Some of the folders may be empty, but copy them anyway.

- `conf`
- `logs`
- `shared`
- `webapps`

Configure `server.xml` in each instance

Edit `conf/server.xml` in the instance home directories to include the Confluence application and have distinct listen ports for Server, HTTP Connector and AJP13 Connector. All nodes can use the same Confluence webapp as long as you set `confluence.home` via a system property (see startup scripts below).

Attached are two sample configurations:

- [tomcat1/conf/server.xml](#) - listens on port 18080 (http) and 18081 (ajp13)
- [tomcat2/conf/server.xml](#) - listens on port 28080 (http) and 28081 (ajp13)

To use these sample config files, you will need to edit them to set the Confluence web-app location and the data source configuration.

If editing the configuration files yourself, the points to note are:

- 'Server' port must be distinct
- 'Connector' for HTTP must be uncommented and use a distinct port. Use this port for testing the node individually.
- 'Connector' for AJP13 must be uncommented and use a distinct port. This port must match the port of the worker in the `Apache workers.properties`.
- 'Engine' for localhost must have `jvmRoute` matching the name of the worker in `Apache's workers.properties`.
- 'Context' for Confluence must be added inside the 'Host' tag, and include a 'Resource' for the datasource, as per normal Confluence installation under Tomcat.

Create a startup script for each instance

The startup scripts for each instance must set the CATALINA_BASE environment variable and confluence.home system property. The variables in the sample scripts below should reference:

- CATALINA_HOME - Tomcat installation directory
- CATALINA_BASE - Tomcat instance home directory (distinct for each node)
- JRE_HOME - Java runtime directory
- JAVA_OPTS - include a confluence-home system property (distinct for each node)

tomcat1/startup.bat:

```
set CATALINA_HOME=C:\home\mryall\opt\apache\apache-tomcat-5.5.16
set CATALINA_BASE=C:\home\mryall\var\tomcat-instances\tomcat1
set JRE_HOME=C:\Java\jre1.5.0_06
set JAVA_OPTS=-Dconfluence.home=C:\home\mryall\data\confluence\cluster\tomcat1
-Xmx512m
%CATALINA_HOME%\bin\startup.bat
```

tomcat2/startup.bat:

```
set CATALINA_HOME=C:\home\mryall\opt\apache\apache-tomcat-5.5.16
set CATALINA_BASE=C:\home\mryall\var\tomcat-instances\tomcat2
set JRE_HOME=C:\Java\jre1.5.0_06
set JAVA_OPTS=-Dconfluence.home=C:\home\mryall\data\confluence\cluster\tomcat2
-Xmx512m
%CATALINA_HOME%\bin\startup.bat
```

Continue setting up Confluence

Follow the [Confluence Cluster Installation](#) procedure with the steps following the app server setup.

Troubleshooting

General advice

The above tomcat configurations enable HTTP connectors on each Tomcat instance so that you can connect to the nodes individually. To check whether the load balancer (Apache & mod_jk) is causing the problem, try connecting to the individual Tomcat instances. Please note that you should not allow users to directly access individual nodes in production mode: You don't want people to bookmark nodes since the node details might change, or single nodes may be taken out of the cluster for maintenance while the cluster itself is still available.

Session-affinity doesn't seem to be working?

Ensure the name you use for your worker in workers.properties (e.g. tomcat1) matches the jvmRoute attribute of the engine tag in your Tomcat server.xml. For an example, search for 'Engine' in the [attached sample config](#).

For troubleshooting your Confluence cluster, see [Cluster Troubleshooting](#).

References

General

<http://raibledesigns.com/tomcat/> <http://httpd.apache.org/>

Tomcat Clustering support

<http://tomcat.apache.org/tomcat-5.0-doc/cluster-howto.html> <http://tomcat.apache.org/tomcat-5.0-doc/balancer-howto.html> http://tomcat.apache.org/tomcat-3.3-doc/mod_jk-howto.html

[Clustering and Load Balancing in Tomcat 5, Part 1](#)

[Clustering and Load Balancing in Tomcat 5, Part 2](#)

Upgrading a Confluence Cluster

This page contains instructions for **upgrading an existing Confluence cluster** to a new version of Confluence.

If you are not running a clustered instance of Confluence and wish to, see [Confluence Cluster Installation with Existing Data](#).

 Oracle Coherence Licensing Change:

- Due to a license agreement change, Confluence is now available in two editions:
 - **Standard Edition** — Confluence with Ehcache's caching technology (available to customers with non-clustered Confluence licenses).
 **If you are currently running a clustered installation of Confluence, please do not upgrade it with a standard edition of Confluence.**
 - **Clustered Edition** — Confluence with Oracle's Coherence clustering and distributed caching technology (available to customers with Confluence clustered licenses only).
- For more information about these changes, please refer to the [Coherence License Changes](#) document.
- If you have a Confluence clustered license, are running a clustered installation of Confluence and wish to upgrade to Confluence version 2.6 or later, please ensure that you download only a [clustered edition of Confluence](#).

You can download the latest version of Confluence from [here](#).

Overview

The steps involved in upgrading a multi-node Confluence cluster are:

1. Backup your confluence instance.
2. Read the Release Notes for this version and check you have the required expertise to perform the upgrade.
3. Stop each node in the cluster.
4. Install the new version into the application server on the first node.
5. Install the new version into the application server onto the remaining nodes.

Step One: Backing up

 We highly recommend that you backup your Confluence home and install directories and your database before proceeding.

For specific files to backup see [Upgrading Confluence](#).

Step Two: Things you need to check ...

- Always check the [release-notes](#) for the version of Confluence you are installing for upgrade instructions specific to that version.
- To perform this upgrade you must be familiar with the usage of the application server running your Confluence Cluster, and the web server load balancing it.
- Check the [Configuring Confluence](#) for your application server and database, to make sure there isn't anything extra you need to do to get Confluence running.
- Check that you know what configurations or customisations have been made to your Confluence instance. These may include specialised user management configurations and changes to Confluence's Java classes and Velocity templates.

Step Three: Stopping the cluster

 It is vital that all nodes in the cluster are running the same version of Confluence. That's why the first step is to stop all the nodes.

Stop the Confluence application on each node using your application server.

Step Four: Upgrading the first node

 We advise configuring your load balancing web server to redirect traffic away from Confluence until the upgrade is complete on multiple nodes.

Upgrading a cluster node uses the same process as [Upgrading Confluence](#).

1. Unzip the new version.
2. Edit its `confluence-init.properties` to point to the existing home directory.
3. Port any immediately required customisations from the old version to the new one. Eg `atlassian-user.xml`.
4. Install the new version into the application server. Eg for Tomcat edit `confluence.xml` or `server.xml` to point to the new location, and restart Tomcat.
5. Wait for the Node to finish upgrading and confirm that you can log in and view pages before continuing to Step Five.
6. Port any additional customisations from the old version to the new version. Eg modifications to Java classes or Velocity templates.

Step Five: Upgrading other nodes

Copy the confluence installation, complete with customisations, to the next node.

1. Edit its `confluence-init.properties` to point to the existing home directory.
2. Install the new version into the application server. Eg for Tomcat 5 edit `confluence.xml` to point to the new location, and restart Tomcat.
3. Wait for the Node to finish upgrading and confirm that you can log in and view pages before continuing with the next node.

Troubleshooting

For suggested troubleshooting techniques, see our [Cluster Troubleshooting](#) page.

Related documentation

[Overview of Confluence Clusters](#)
[Confluence Installation Guide](#)
[Cluster Troubleshooting](#)
[Confluence Cluster Installation](#)
[Confluence Cluster Installation with Existing Data](#)
[Confluence User Guide](#)

Creating a Dedicated User Account on the Operating System to Run Confluence

 This step is optional if you are evaluating Confluence, but should be mandatory for Confluence installations used in production. If you have used the Confluence installer on Linux, this user will be created automatically.

A dedicated user should be created to run Confluence, because Confluence runs as the user it is invoked under and therefore can potentially be abused. For example:

- If your operating system is *nix-based (for example, Linux or Solaris), type the following in a console:

```
$ sudo /usr/sbin/useradd --create-home --comment "Account for running Confluence" --shell /bin/bash confluence
```
- If your operating system is Windows:
 1. Create the dedicated user account by either:
 - Typing the following at the Windows command line:

```
> net user confluence mypassword /add /comment:"Account for running Confluence"
```

 (This creates a user account with user name 'confluence' and password 'mypassword'. You should choose your own password.)
 - Opening the Windows 'Computer Management' console to add your 'confluence' user with its own password.
 2. (Optional) Use the Windows 'Computer Management' console to remove the 'confluence' user's membership of all unnecessary Windows groups, such as the default 'Users' group.
-  If Windows is operating under Microsoft Active Directory, ask your Active Directory administrator to create your 'confluence' account (with no prior privileges).

Ensure that *only the following directories* can be written to by this dedicated user account (e.g. 'confluence'):

- The following subdirectories of your [Confluence Installation Directory](#):
 - logs
 - temp
 - work
- Your [Confluence Home Directory](#).

 Do not make the [Confluence Installation Directory](#) itself writeable by the dedicated user account.

 See also [Best Practices for Configuring Confluence Security](#).

Confluence Setup Guide

Before running the Confluence Setup Wizard, as described below, you should have already completed [installing Confluence](#).

When you access Confluence in your web browser for the first time, you will see the **Confluence Setup Wizard**. This is a series of screens which will prompt you to supply some default values for your Confluence site. It will also offer some more advanced options for setting up data connections and restoring data from a previous installation.

1. Start the Setup Wizard

1. If Confluence is not already running, start it now:
 - If you are running the Confluence distribution on Windows, click **Start > Programs > Confluence > Start Confluence Server**.
 - Or, run the start-up script found in the bin folder of your installation directory:
 - start-confluence.bat for Windows.
 - start-confluence.sh for Linux-based systems.
2. Go to the following web address in your web browser: <http://localhost:8090>
The above web address uses port '8090'. If you chose a different port during installation, change '8090' to the number you chose. If an error message appears, check that you are using the port which you specified during installation.

On this page:

- [1. Start the Setup Wizard](#)
- [2. Choose your Installation Type](#)
- [3. Enter your License Key](#)
- [4. Production Installation: Database Configuration](#)
- [5. Production Installation: External Database](#)
- [6. Production Installation: Load Content](#)
- [7. Production Installation: Restore Data from Backup](#)
- [8. Set Up User Management](#)
- [9. Connect to JIRA](#)
- [10. Set Up System Administrator](#)
- [11. Setup is Complete](#)
- [12. Install add-ons \(optional\)](#)

2. Choose your Installation Type

In this step, you will choose whether you want a trial or a production installation.

Option 1: Trial Installation - Set up Confluence with the embedded [HSQLDB database](#) and default settings. This option will also generate an evaluation license and install a Demonstration space with some example content. You can upgrade to another type of database later on.

Choose this option if:

- You want to evaluate Confluence or if you are new to Confluence.
- You do not have a Confluence licence.



For production use, we strongly recommend that you connect to an external database rather than using the embedded database. The evaluation installation is therefore not suitable for production environments.

Option 2: Production Installation - Customise your Confluence instance to use your own database and your own data.

The production installation offers the following options:

- Connect Confluence to an external database. **Recommended for Confluence used in production environments.**
- Restore data from an existing Confluence database.
- Install Confluence without the demonstration content.

3. Enter your License Key

Trial installation: Follow the prompts in the setup wizard to generate an evaluation license. Once your license has been generated go to [step 8 below](#).

Production installation: Find your Confluence license key and paste it into the **License Key** field, shown on the screenshot above.

If you already have a license key, you can retrieve it from the [Atlassian website](#).

If you do not already have a Confluence license, you can obtain one now:

- To get a free evaluation license:
 1. Click **generate an evaluation license online** on the setup wizard, shown on the screenshot above.
 2. Follow the prompts to generate your license key and insert it into the setup wizard's licensing screen automatically.
- To get a commercial, academic, non-profit or open source license:
 1. Copy your **Server ID** from the setup wizard's licensing screen, shown on the screenshot above.
 2. Choose the license type you need from the list on the [Atlassian website](#).
 3. Complete the [online order form](#).

4. Production Installation: Database Configuration

Choose a Database Configuration

Choose where Confluence should store its data

Embedded Database

The embedded database will allow Confluence to operate without an external database.

We **strongly recommend against** using this on a **production** server. This is **recommended for evaluating and demonstrating Confluence only**. Production systems should consider an external database for improved scalability and reliability. A guide on how to migrate your data can be found [here](#).

Embedded Database

External Database

If you wish to store your Confluence data in an external database, choose it from the list of supported databases. This is **recommended for production systems**.

Confluence supports a number of databases, and does not strongly recommend any one database. Atlassian provides only limited support for maintaining and tuning databases, so if you already have an established, supported database of choice within your organisation it is advisable to use this database. The benefit of having an expert to diagnose operational and performance issues far outweighs any differences between the databases themselves.

If you have no established database and do not have a strong preference for any of our supported databases, we recommend the latest supported version of [PostgreSQL](#), which is free and thoroughly tested against. Please check the latest supported version in our [documentation](#).

If your database is not listed in the menu, you may configure an "Unsupported Database", but be aware that Confluence may not be fully tested against this database.

PostgreSQL



External Database

Screenshot above: Database configuration

The above screen appears if you have chosen a production installation of Confluence. You can choose to use the embedded database supplied with your Confluence installation, or to connect to an external database.

- **Option 1: Embedded Database** — If you select this option, Confluence will use an embedded [HSQLDB database](#). You should only select this option for the purposes of **evaluating** or **demonstrating** the use of Confluence. You can migrate to an external database later on if you wish.
- **Option 2: External Database** — If you wish Confluence to use an external database, select your database type from the dropdown list and then click the '**External Database**' button.
 - For production purposes, you **should use an external database** to ensure your data is kept safe and consistent.
 - If you choose **PostgreSQL**, please make sure that the version you install is supported by Atlassian. It is possible that we do not yet support the latest version of PostgreSQL.
 - Read the page about [supported platforms](#) for more information about which databases are supported. For details about choosing an external database, refer to the page on [system requirements](#). For information about configuring an external database, see [Database Configuration](#).

5. Production Installation: External Database




Before you Start

- **Character encoding:**
 - We strongly recommend that character encoding is consistent across your database, application server and web application, and that you use **UTF-8** encoding.

- Before setting up your database, please read about [configuring character encoding](#).
- **Database name:** When creating a new external database, give it the name '**confluence**'.

You can choose how you wish Confluence to connect to your database - via a direct JDBC connection or via a server-managed datasource connection. Choose one of the two options below.

Option 1: Direct JDBC — This uses a standard JDBC database connection. Connection pooling is handled within Confluence.



The screenshot shows a web form titled "Setup Database". It contains four input fields: "Driver Class Name*" with the value "org.postgresql.Driver", "Database URL*" with the value "jdbc:postgresql://localhost:5432/confluence", "User Name*", and "Password". Below these fields is a "Next" button.

Screenshot above: Standard (JDBC) connection


Supply the following information:

- **Driver Class Name** – The Java class name for the appropriate database driver. This will depend on the JDBC driver, and will be found in the documentation for your database. Note that Confluence bundles some database drivers, but you'll need to install the driver yourself if it is not bundled. See [Database JDBC Drivers](#) for details.
- **Database URL** – The JDBC URL for the database you will be connecting to. This will depend on the JDBC driver, and will be found in the documentation for your database.
- **User Name** – A valid username which Confluence will use to access your database.
- **Password** – The password corresponding to the above username.

You will also need to know:

- The size of the connection pool Confluence should maintain. If in doubt, just go with the default provided.
- What kind of database you're connecting to, so you can tell Confluence which dialect it needs to use.

Option 2: Datasource — This asks the Java application server for a database connection. You will need to have configured a datasource in your application server. For information about configuring an external database, see [Database Configuration](#).



The screenshot shows a web form titled "Setup Datasource Connection". At the top, it says "Provide the JNDI name of the Datasource Confluence should use to connect to the database." Below this is a text input field labeled "Datasource Name" with the value "java:comp/env/jdbc/". Below the field is a "Next" button. A note above the field says: "If 'java:comp/env/jdbc/DatasourceName' doesn't work, try 'jdbc/DatasourceName' (or vice versa)".

Screenshot above: Datasource connection

Supply the following information:

- **Datasource Name** — The JNDI name of the datasource, as configured in the application server. Note: Some servers will have JNDI names like `jdbc/datasourcename`; others will be of the form `java:comp/env/jdbc/datasourcename`. Consult your application-server documentation.

You will also need to know:

- What kind of database you're connecting to, so you can tell Confluence which dialect it needs to use.

6. Production Installation: Load Content

Load Content

If you are evaluating or demonstrating Confluence, or are introducing Confluence to users who are new to the idea of a Wiki, we **recommend the example site** as the best way to become acquainted with what Confluence can do for you. More experienced users will want to start with an empty site, or restore a backup of their own.

Example Site

Recommended: Load the 'Demonstration Space' to begin working with Confluence immediately.

Example Site

Empty Site

Start with an empty site. After finishing the setup you will need to create at least one space before you can add any content of your own.

Empty Site

Restore From Backup

Use data from another installation of Confluence. If you are migrating to another database or replicating Confluence you will probably want to select this option.

Restore From Backup

Screenshot above: Load content

Select one of the following options:

- **Example Site** - This option will load Confluence's 'Demonstration Space'. Select this if you are using Confluence for the first time, or if you want the Demonstration Space for your other Confluence users. The Demonstration Space helps to familiarise you with Confluence and what it can do for you. You can then continue using your Confluence deployment as normal - there's no need to reinstall later.
- **Empty Site** - Select this option if you are already familiar with Confluence. You will need to create at least one space before you can start adding content to the site.
- **Restore from Backup** - Select this option if you want to use Confluence data from a previous installation.

7. Production Installation: Restore Data from Backup

This option allows you to reload your data from an existing Confluence installation into your new Confluence site during the initial setup procedure. You can choose to upload data from a zipped backup file, or to restore from a backup file on your file system.

Option 1: Upload a zipped backup to Confluence - This option will load the data from a zipped backup file. (To create a backup file from your existing version of Confluence, go to the '**Backup & Restore**' section of your **Administration Console**.)

To restore from a zipped backup:

1. Browse for the relevant daily backup file or a file you have created via a [manual backup](#).
2. Check '**Build Index**' to build the data index, used for the search.
3. Click the '**Upload and Restore**' button.


Option 2: Restore a backup from the filesystem - This option is recommended if you have a very large daily backup file (greater than 100MB), or a daily backup file that is already on the server and doesn't require uploading.

1. Copy the XML backup file into the restore directory inside your confluence [Home directory](#) and then refresh the page. You should now see your backup file appear on the '**Restore Data**' screen (pictured

- above), in the box beneath the heading '**Restore a backup from the filesystem**'.
2. Check '**Build Index**' to build the data index, used for the search.
 3. Click the '**Restore**' button.

When the restore process has finished, you are ready to log in to Confluence. The system administrator account and all other information has been transferred from your previous Confluence installation.

8. Set Up User Management



Manage Users and Groups within Confluence

Confluence will maintain its own database of users and groups. If you are unsure, choose this option. You can always change the user management configuration after Confluence has been set up.

If you are planning to connect to an LDAP or Crowd server, you should choose this option, and complete your user management configuration after Confluence is set up. [Learn more about configuring LDAP integration.](#)

Manage users and groups within Confluence

Manage Users and Groups with JIRA

If you have JIRA 4.3 or later, Confluence can use JIRA for user management. This is not recommended for more than 500 users. [Learn more about configuring JIRA integration.](#)

Connect to JIRA

Screenshot above: User management

You can choose to manage Confluence's users and groups inside Confluence or in JIRA.

- If you do not have Atlassian JIRA installed, or if you would prefer to set up external user management later, choose **Manage users and groups within Confluence**.
- If you have JIRA installed, the setup wizard gives you the opportunity to configure the JIRA connection automatically. This is a quick way of setting up your JIRA integration with the most common options. It will configure a JIRA user directory for Confluence, and set up application links between JIRA and Confluence for easy sharing of data. Choose **Connect to JIRA**.

9. Connect to JIRA

Connect to JIRA

Fill in the details of the JIRA server you wish to retrieve user and group information from. You will need a valid administrator account on that JIRA server. [Learn more about configuring JIRA integration.](#)

JIRA Server Location

JIRA Base URL *

For example: `http://jira.mycompany.com` or `http://mycompany.com/jira`

JIRA Administrator Login

Confluence will need to log in to JIRA with administrative privileges to retrieve user information.

Username *

This username must have system administrator rights on your JIRA server.

Password *

Advanced Options

Confluence Base URL *

JIRA will use this URL to access your Confluence server. If Confluence is behind a proxy, you may need to change the URL given here.

User Groups *

Users in these groups will have access to Confluence. Comma-separated.

Admin Groups *

Users in these groups will have administrator access to Confluence. Comma-separated.

[Return to User Management Selection](#)

Screenshot above: Connecting to JIRA in the Confluence setup wizard

Enter the following information:

- **JIRA Base URL**– The web address of your JIRA server. Examples:

`http://www.example.com:8080/jira/`
`http://jira.example.com`

- **JIRA Administrator Login: Username** – Enter the username of a user with the 'JIRA System Administrators' global permission in JIRA.
- **JIRA Administrator Login: Password** – Enter the password that the above user uses to sign in to JIRA.
- **Confluence Base URL** – JIRA will use this URL to access your Confluence server. The URL you give here will override the base URL specified in your Confluence administration console, for the purposes of the JIRA connection.
- **User Groups** – Specify one or more JIRA groups whose members should be able to use Confluence. The default group is `jira-users`. (These groups will receive the 'can use' permission in Confluence.)
- **Admin Groups** – Specify one or more JIRA groups whose members should have administrative access to Confluence. The default group is `jira-administrators`. (These groups will receive the 'Confluence system administrator' and 'Confluence administrator' permissions in Confluence.)

For full details and a troubleshooting guide, see [Configuring JIRA Integration in the Setup Wizard](#).

10. Set Up System Administrator

The [system administrator](#) has full administrative power over your Confluence instance. This person will be able to add more users, create spaces, and set further Confluence options. Please refer to the [overview of global permissions](#) for more information.



Hint: If you are evaluating Confluence, set yourself as the administrator.

1. Enter the following information to set up your system administrator's user account:
 - **Username** — The username under which the system administrator will log in to Confluence, e.g. 'jsmith'.
 - **Password** — The password which the system administrator will use to log in.
 - **Confirm** — Enter the same password again.
 - **Name** — The system administrator's full name, e.g. 'John Smith'.
 - **Email** — The system administrator's email address, e.g. 'jsmith@example.com'.
2. Click 'Next'.

11. Setup is Complete

Congratulations! You have installed and set up Confluence. Click **Start using Confluence** to open the **Demonstration space** in your Confluence wiki. This space contains some sample content and ideas, to help you get started quickly.

Click **Further Configuration** if you want to go directly to the Administration Console and complete administrator's tasks including configuring a mail server, adding users, changing the base URL and more. Refer to the [Confluence Administrator's Guide](#) for more information.

12. Install add-ons (optional)

If you have chosen a package that includes add-ons such as Team Calendars or the SharePoint Connector you can now install these from within Confluence.

In the Administration Console go to **Find New Add-ons** and search for your add-on. Follow the prompts to install the add-on.

See [Finding new add-ons](#) for more information on how to find and install the add-ons.

Configuring JIRA Integration in the Setup Wizard

This page describes the **Connect to JIRA** step in the Confluence setup wizard.

Overview

You can connect your application to a JIRA server, to manage your users via JIRA and share information with JIRA. When you are installing the application, the setup wizard gives you the opportunity to configure the JIRA connection automatically. This is a quick way of setting up your JIRA integration with the most common options.

You can also configure the JIRA connections via the application administration screens. In that case, you will need to set up connections individually. There are two parts to the integration process:

- A peer-to-peer link between JIRA and the application for sharing information and facilitating integration features. This link is set up via Application Links.
- A client-server link between the application and JIRA for delegating user and group management to your JIRA server.

Requirements: You need JIRA 4.3 or later.

On this page:

- Overview
- Connecting to JIRA in the Setup Wizard
- Troubleshooting
 - Solution 1: Removing a Partial Configuration – The Easiest Way
 - Solution 2: Removing a Partial Configuration – The Longer Way
- Notes

Connecting to JIRA in the Setup Wizard

Connect to JIRA

Fill in the details of the JIRA server you wish to retrieve user and group information from. You will need a valid administrator account on that JIRA server. [Learn more about configuring JIRA integration.](#)

JIRA Server Location

JIRA Base URL *

For example: `http://jira.mycompany.com` or `http://mycompany.com/jira`

JIRA Administrator Login

Confluence will need to log in to JIRA with administrative privileges to retrieve user information.

Username *

This username must have system administrator rights on your JIRA server.

Password *

Advanced Options

Confluence Base URL *

JIRA will use this URL to access your Confluence server. If Confluence is behind a proxy, you may need to change the URL given here.

User Groups *

Users in these groups will have access to Confluence. Comma-separated.

Admin Groups *

Users in these groups will have administrator access to Confluence. Comma-separated.

[Return to User Management Selection](#)

Screenshot above: Connecting to JIRA in the Confluence setup wizard

Enter the following information:

- **JIRA Base URL**– The web address of your JIRA server. Examples:

`http://www.example.com:8080/jira/`
`http://jira.example.com`

- **JIRA Administrator Login: Username** – Enter the username of a user with the 'JIRA System Administrators' global permission in JIRA.
- **JIRA Administrator Login: Password** – Enter the password that the above user uses to sign in to JIRA.

- **Confluence Base URL** – JIRA will use this URL to access your Confluence server. The URL you give here will override the base URL specified in your Confluence administration console, for the purposes of the JIRA connection.
- **User Groups** – Specify one or more JIRA groups whose members should be able to use Confluence. The default group is `jira-users`. (These groups will receive the 'can use' permission in Confluence.)
- **Admin Groups** – Specify one or more JIRA groups whose members should have administrative access to Confluence. The default group is `jira-administrators`. (These groups will receive the 'Confluence system administrator' and 'Confluence administrator' permissions in Confluence.)

Troubleshooting

This section describes the possible problems that may occur when integrating your application with JIRA via the setup wizard, and the solutions for each problem.

Symptom	Cause	Solution
The setup wizard displays one of the following error messages: • Failed to create application link from JIRA server at <URL> to this <application> server at <URL>. • Failed to create application link from this <application> server at <URL> to JIRA server at <URL>. • Failed to authenticate application link from JIRA server at <URL> to this <application> server at <URL>. • Failed to authenticate application link from <application> server at <URL> to this JIRA server at <URL>.	The setup wizard failed to complete registration of the peer-to-peer application link with JIRA. JIRA integration is only partially configured.	Remove the partial configuration if it exists, try the 'Connect to JIRA' step again, and then continue with the setup. Detailed instructions are below.
The setup wizard displays one of the following error messages: • Failed to register <application> configuration in JIRA for shared user management. Received invalid response from JIRA: <response> • Failed to register <application> configuration in JIRA for shared user management. Received: <response>	The setup wizard failed to complete registration of the client-server link with JIRA for user management. The peer-to-peer link was successfully created, but integration is only partially configured.	Remove the partial configuration if it exists, try the 'Connect to JIRA' step again, and then continue with the setup. Detailed instructions are below.
The setup wizard displays the following error message: • Error setting Crowd authentication	The setup wizard successfully established the peer-to-peer link with JIRA, but could not persist the client-server link for user management in your <code>config.xml</code> file. This may be caused by a problem in your environment, such as a full disk.	Please investigate and fix the problem that prevented the application from saving the configuration file to disk. Then remove the partial configuration if it exists, try the 'Connect to JIRA' step again, and then continue with the setup. Detailed instructions are below.

The setup wizard displays the following error message: Error reloading Crowd authentication	The setup wizard has completed the integration of your application with JIRA, but is unable to start synchronizing the JIRA users with your application.	Restart your application. You should then be able to continue with the setup wizard. If this solution does not work, please contact Atlassian Support.
The setup wizard displays the following error message: An error occurred: java.lang.IllegalStateException : Could not create the application in JIRA/Crowd (code: 500). Please refer to the logs for details.	The setup wizard has not completed the integration of your application with JIRA. The links are only partially configured. The problem occurred because there is already a user management configuration in JIRA for this <application> URL.	Remove the partial configuration if it exists, try the 'Connect to JIRA' step again, and then continue with the setup. Detailed instructions are below.
No users can log in after you have set up the application with JIRA integration.	Possible causes: - There are no users in the group that you specified on the 'Connect to JIRA' screen. - For FishEye: There are no groups specified in the 'groups to synchronize' section of your administration console. - For Stash: You may not have granted any JIRA groups or users permissions to log in to Stash.	Go to JIRA and add some usernames to the group. - For FishEye: Go to the FishEye administration screens and specify at least one group to synchronize. The default is 'jira-users'. - For Stash: Grant the Stash User permission to the relevant JIRA groups on the Stash Global permissions page. If this solution does not work, please contact Atlassian Support.

Solution 1: Removing a Partial Configuration – The Easiest Way

If the application's setup wizard fails part-way through setting up the JIRA integration, you may need to remove the partial configuration from JIRA before continuing with your application setup. Please follow the steps below.

Remove the partial configuration if it exists, try the 'Connect to JIRA' step again, and then continue with the setup wizard:

- Log in to JIRA as a user with the '**JIRA System Administrators**' global permission.
- Click the '**Administration**' link on the JIRA top navigation bar.
- Remove the application link from JIRA, if it exists:
 - Click **Application Links** in the JIRA administration menu. The 'Configure Application Links' page will appear, showing the application links that have been set up.
 - Look for a link to your application. It will have a base URL of the application linked to JIRA. For example:
 - If you want to remove a link between JIRA and FishEye, look for the one where the **Application URL** matches the base URL of your FishEye server.
 - If you want to remove a link between JIRA and Confluence, look for the one where the **Application URL** matches the base URL of your Confluence server.
 - If you want to remove a link between JIRA and Stash, look for the one where the **Application URL** matches the base URL of your Stash server.
 - Click **Delete** next to the application link that you want to delete.
 - A confirmation screen will appear. Click **Confirm** to delete the application link.
- Remove the user management configuration from JIRA, if it exists:
 - Go to the JIRA administration screen for configuring the applications that have been set up to use JIRA for user management:
 - In JIRA 4.3: Click '**Other Applications**' in the '**Users, Groups & Roles**' section of the JIRA administration screen.
 - In JIRA 4.4: Select '**Administration**' > '**Users**' > '**JIRA User Server**'.
 - Look for a link to your application. It will have a name matching this format:

```
<Type> - <HostName> - <Application ID>
```

For example:

```
FishEye / Crucible - localhost -  
92004b08-5657-3048-b5dc-f886e662ba15
```

Or:

```
Confluence - localhost - 92004b08-5657-3048-b5dc-f886e662ba15
```

If you have multiple servers of the same type running on the same host, you will need to match the application ID of your application with the one shown in JIRA. To find the application ID:

- Go to the following URL in your browser:

```
<baseUrl>/rest/applinks/1.0/manifest
```

Replace <baseUrl> with the base URL of your application.

For example:

```
http://localhost:8060/rest/applinks/1.0/manifest
```

- The application links manifest will appear. Check the application ID in the <id> element.
 - In JIRA, click **Delete** next to the application that you want to remove.
- Go back to the setup wizard and try the 'Connect to JIRA' step again.

Solution 2: Removing a Partial Configuration – The Longer Way

If solution 1 above does not work, you may need to remove the partial configuration and then add the full integration manually. Please follow these steps:

- Skip the 'Connect to JIRA' step and continue with the setup wizard, to complete the initial configuration of the application.
- Log in to JIRA as a user with the '**JIRA System Administrators**' global permission.
- Click the '**Administration**' link on the JIRA top navigation bar.
- Remove the application link from JIRA, if it exists:
 - Click **Application Links** in the JIRA administration menu. The 'Configure Application Links' page will appear, showing the application links that have been set up.
 - Look for a link to your application. It will have a base URL of the application linked to JIRA. For example:
 - If you want to remove a link between JIRA and FishEye, look for the one where the **Application URL** matches the base URL of your FishEye server.
 - If you want to remove a link between JIRA and Confluence, look for the one where the **Application URL** matches the base URL of your Confluence server.
 - If you want to remove a link between JIRA and Stash, look for the one where the **Application URL** matches the base URL of your Stash server.
 - Click **Delete** next to the application link that you want to delete.
 - A confirmation screen will appear. Click **Confirm** to delete the application link.
- Remove the user management configuration from JIRA, if it exists:
 - Go to the JIRA administration screen for configuring the applications that have been set up to use JIRA for user management:
 - In JIRA 4.3: Click '**Other Applications**' in the '**Users, Groups & Roles**' section of the JIRA administration screen.
 - In JIRA 4.4: Select '**Administration**' > '**Users**' > '**JIRA User Server**'.

- b. Look for a link to your application. It will have a name matching this format:

```
<Type> - <HostName> - <Application ID>
```

For example:

```
FishEye / Crucible - localhost -  
92004b08-5657-3048-b5dc-f886e662ba15
```

Or:

```
Confluence - localhost - 92004b08-5657-3048-b5dc-f886e662ba15
```

If you have multiple servers of the same type running on the same host, you will need to match the application ID of your application with the one shown in JIRA. To find the application ID:

- Go to the following URL in your browser:

```
<baseUrl>/rest/applinks/1.0/manifest
```

Replace `<baseUrl>` with the base URL of your application.

For example:

```
http://localhost:8060/rest/applinks/1.0/manifest
```

- The application links manifest will appear. Check the application ID in the `<id>` element.
 - In JIRA, click **Delete** next to the application that you want to remove.
6. Add the application link in JIRA again, so that you now have a two-way trusted link between JIRA and your application:
- Click **Add Application Link**. Step 1 of the link wizard will appear.
 - Enter the **server URL** of the application that you want to link to (the 'remote application').
 - Click **Next**.
 - Enter the following information:
 - Create a link back to this server** – Check to add a two-way link between the two applications.
 - Username and Password** – Enter the credentials for a username that has administrator access to the remote application.
Note: These credentials are only used to authenticate you to the remote application, so that Application Links can make the changes required for the new link. The credentials are not saved.
 - Reciprocal Link URL** – The URL you give here will override the base URL specified in your remote application's administration console, for the purposes of the application links connection. Application Links will use this URL to access the remote application.
 - Click **Next**.
 - Enter the information required to configure authentication for your application link:
 - The servers have the same set of users** – Check this box, because the users are the same in both applications.
 - These servers fully trust each other** – Check this box, because you trust the code in both applications and are sure both applications will maintain the security of their private keys.
For more information about configuring authentication, see [Configuring Authentication for an Application Link](#).
 - Click **Create**.
7. Configure a new connection for user management in JIRA:
- Go to the JIRA administration screen for configuring the applications that have been set up to use JIRA for user management:
 - In JIRA 4.3: Click **'Other Applications'** in the **'Users, Groups & Roles'** section of the JIRA

- administration screen.
- In JIRA 4.4: Select '**Administration**' > '**Users**' > '**JIRA User Server**'.
- b. **Add** an application.
- c. Enter the **application name** and **password** that your application will use when accessing JIRA.
- d. Enter the **IP address** or addresses of your application. 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.
- 8. Set up the JIRA user directory in the application.
 - For Confluence:
 - a. Go to the **Confluence Administration Console**.
 - b. Click '**User Directories**' in the left-hand panel.
 - c. **Add** a directory and select type '**Atlassian JIRA**'.
 - d. Enter the following information:
 - **Name** – Enter the name of your JIRA server.
 - **Server URL** – Enter web address of your JIRA server. Examples:


```
http://www.example.com:8080/jira/
http://jira.example.com
```
 - **Application name** and **Application password** – Enter the values that you defined for Confluence in the settings on JIRA.
 - e. Save the directory settings.
 - f. Define the **directory order** by clicking the blue up- and down-arrows next to each directory on the '**User Directories**' screen.
For details see [Connecting to Crowd or JIRA for User Management](#).
 - For FishEye/Crucible:
 - a. Click **Authentication** (under 'Security Settings').
 - b. Click **Setup JIRA/Crowd authentication**. Note, if LDAP authentication has already been set up, you will need to remove that before connecting to JIRA for user management.
 - c. Make the following settings:

Authenticate against	Select a JIRA instance
Application name and password	Enter the values that you defined for your application in the settings on JIRA.
JIRA URL	The web address of your JIRA server. Examples: <pre>http://www.example.com:8080/jira/ http://jira.example.com</pre>
Auto-add	Select Create a FishEye user on successful login so that your JIRA users will be automatically added as a FishEye user when they first log in.
Periodically synchronise users with JIRA	Select Yes to ensure that JIRA will synchronize all changes in the user information on a regular basis. Change the value for Synchronise Period if required.

When Synchronisation Happens	Select an option depending on whether you want to allow changes to user attributes from within FishEye.
Single Sign On	Select Disabled . SSO is not available when using JIRA for user management and if enabled will make the integration fail.

- d. Click **Next** and select at least one user group to be synchronised from JIRA. If necessary, you could create a new group in JIRA, such as 'fisheye-users', and select this group here.
 - e. Click **Save**.
 - For Stash:
 - a. Go to the Stash administration area.
 - b. Click **User Directories** in the left-hand panel.
 - c. **Add** a directory and select type **Atlassian JIRA**.
 - d. Enter the following information:
 - **Name** – Enter the name of your JIRA server.
 - **Server URL** – Enter web address of your JIRA server. Examples:


```
http://www.example.com:8080/jira/
http://jira.example.com
```
 - **Application name and Application password** – Enter the values that you defined for Stash in the settings on JIRA.
 - e. Save the directory settings.
 - f. Define the directory order by clicking the blue up- and down-arrows next to each directory on the 'User Directories' screen.
- For details see [Connecting Stash to JIRA for user management](#).

Notes

- When you connect to JIRA in the setup wizard, the setup procedure will configure **Trusted Applications authentication** for your application. Please be aware of the following security implications:
 - Trusted applications are a **potential security risk**. When you configure Trusted Applications authentication, you are allowing one application to access another as any user. This allows all of the built-in security measures to be bypassed. Do not configure a trusted application unless you know that all code in the application you are trusting will behave itself at all times, and you are sure that the application will maintain the security of its private key.
- In the next step, you will specify the username and password of your **Confluence system administrator**. If you have connected to JIRA, the setup wizard will add the Confluence administrator's username and password to both JIRA and Confluence. This is done so that you can still access Confluence even if JIRA is down. Please note that the password in Confluence is not linked to the password in JIRA. If you subsequently change the administrator's password, only the password in JIRA will change. This is because the JIRA user directory is placed first in the list of user directories. See [Managing Multiple Directories](#).

Related Topics

[User Management Limitations and Recommendations](#)
[Confluence Setup Guide](#)

Upgrading Confluence

This document describes the procedure for upgrading to the latest version of Confluence on Windows or Linux.

Before you start

- Check your Confluence **licence is valid**.
 To check go to Confluence Admin > License Details and make sure the license support period has not expired. If your support period has expired renew your licence and reapply it before proceeding with the upgrade.

- Read the [Release Notes](#) and [Upgrade Notes](#) for both the version you are upgrading to, and any versions you are skipping.
- Check that your **Java version, application server, operating system, database and browsers** are supported.
See [Supported Platforms](#) and [End of Support Announcements for Confluence](#) to confirm latest requirements.
- Check the latest **database setup guide** for your database and ensure that the database is configured correctly. There may be new configuration requirements. See [Database Configuration](#).
- Check the **compatibility of any add-ons**.
Some add-ons may not yet be compatible with the latest version of Confluence. See [Checking Add-on Compatibility with Application Updates](#) to find out how to check this in the Universal Plugin Manager.
- Check for any **known issues** that might affect your instance.
See the [Confluence Knowledge Base](#) for known issues for the version you are upgrading to and [Database Troubleshooting](#) for known issues related to supported external databases.
- Make a note of any **modifications** to your Confluence instance (for example layouts or a custom theme). Any customisation you wish to maintain will need to be reapplied after upgrading. See [Additional steps where customisations are present](#).

Step 1 Determine your upgrade path and method

Find the upgrade path that works for your current version of Confluence and the version you plan to upgrade to.

The following table will help you to determine the most efficient upgrade path from your current version to the latest versions of Confluence. To use the table find your current installed version of Confluence in the left column and follow the suggested path.

Your Version	Recommended upgrade path to Confluence 5
2.7 or earlier	Upgrade to 2.7.4 then upgrade to 3.5 , and follow paths below.
2.8 to 3.4	Upgrade to 3.5.17 , and follow paths below.
3.5	Upgrade to 5.0.3 then upgrade to the latest version of Confluence 5 .
4.0 to 4.3	Upgrade directly to the latest version of Confluence 5 .
5.0 to 5.5	Upgrade directly to the latest version of Confluence 5 .

There are several factors that will determine the upgrade method you should use. If you:

- Are upgrading from a version **earlier than Confluence 3.5** you will need to [upgrade manually](#). You can contact [Support](#) for help determining an appropriate upgrade path.
- Are moving to a **different operating system or file location** you should [upgrade manually](#).
- Are moving to a **different database** you should upgrade using the installer first and then follow the procedure outlined in [migrating to a different database](#).
- Have a **clustered** instance of Confluence you should follow the procedure outlined in [Upgrading a Confluence Cluster](#).
- Are running the **EAR-WAR edition** you should check [Installing the Confluence EAR-WAR Edition](#) to see if any additional steps are required.
- Are using the **embedded HSQLDB database** you should [migrate to a different database before](#) upgrading. This database is supplied for evaluation purposes only and is not recommended for production environments. See [embedded HSQLDB database](#) for more information.

otherwise you should follow the instructions below and use the Windows or Linux installer to upgrade Confluence.

Step 2 Upgrade Confluence in a test environment

We strongly recommend you recreate your production instance and test the upgrade in this cloned environment.

1. Create a snapshot of your current production Confluence environment on a test server - see [Moving Confluence Between Servers](#) for how to do this.
2. Follow the steps below to perform the upgrade on your cloned environment.
3. Test all your [unsupported add-ons \(plugins\)](#) and any customisation (for example custom themes and layouts) with the new version before proceeding with the upgrade in your production environment.

Step 3 Back up

Before you begin the Confluence upgrade you must back up:

- **your external database**
You must perform a manual backup of your external database and confirm that the backup was created properly. If you are unfamiliar with the backup-restore facilities of your database, you can simply restore the backup to a different system to ensure the backup worked before proceeding.
- **your Confluence Home directory**
The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. 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.
if you store attachments outside the Confluence Home directory, you should also backup your attachments directory.
- **the Confluence installation directory** or Confluence webapp (if you are using the EAR-WAR edition)
This is where the Confluence application files and libraries were unpacked (unzipped) when Confluence was originally installed. Confluence does not modify or store any data in this directory.

The installation wizard will back up your Confluence directories as part of the installation process, but you should also back these directories up manually before starting the upgrade.

Step 4 Upgrade Confluence in your production environment

1. Download the appropriate Windows or Linux installer from the [Confluence Download Center](#).
2. Start the installer:
 - Windows Users: run the .exe file.
If prompted to allow the upgrade wizard to make changes to your computer, choose 'Yes'. If you do not, the installation wizard will have restricted access to your operating system and any subsequent installation options will be limited.
 - Linux users: open a Linux console and change directory (cd) to the '.bin' file directory and execute the '.bin' file.
If the '.bin' file is not executable after downloading it, make it executable, for example `chmod a+x atlassian-confluence-5.4.1-x64.bin` (specify the exact filename of the installer you downloaded).
3. The installation wizard will guide you through the upgrade process. Some things to note:
 - a. When prompted choose **Upgrade an existing Confluence installation** (for Linux users this is option 3).
 - b. Verify that the **Existing Confluence installation directory** suggested by the wizard is correct. This is especially important if you have multiple Confluence installations running on the same machine.
 - c. At the 'Back up Confluence directories' step, ensure '**Back up Confluence home**' is selected. This will create a .zip backup of the Confluence home and installation directories. **This is strongly recommended.**
 - d. The installation wizard will notify you of customisations in the Confluence Installation directory. Make a note of these before proceeding as you will need to manually reapply these customisations after the upgrade is complete.
 - e. If you have not already done so, the wizard will prompt you to backup your external database and check plugin compatibility. If your database does not support online backups you will need to stop the installation wizard at this point.
4. The wizard will shut down your Confluence instance and proceed with the upgrade. Once complete, it will restart Confluence and you can then launch Confluence in your browser to confirm the upgrade was

successful.

During the upgrade the wizard will migrate following from your existing Confluence installation:

- TCP port values in your `server.xml` file.
- Custom values in your `confluence-init.properties` (`confluence.home` property) and `setenv.sh` / `setenv.bat` files (`JAVA_OPTS` parameters)

⚠ Other configurations or customisations (including any other modifications in the `server.xml` file) are **not migrated** during the upgrade and need to be reapplied manually. See below for more information.

Additional steps when customisations are present

The installation wizard's ability to notify you about customisations will depend on how your existing Confluence instance was installed:

- If your current Confluence instance was installed using the installer, the wizard will check the entire Confluence Installation directory.
- If your current Confluence instance was installed manually it will only check the `confluence` subdirectory of the Confluence Installation directory. The installation wizard will **not** notify you of modifications in any other directory, for example modifications to start-up scripts under the `bin` directory or modifications to the `server.xml` file (such as an [SSL configuration](#)).

If customisations are present you will need to perform the following steps after the upgrade is complete:

1. Stop the upgraded Confluence instance.
2. Reapply the customisations to the relevant files in the upgraded Confluence Installation directory.
3. Restart the upgraded Confluence instance.

We **strongly recommend** you test your customisations in a test instance prior to upgrading your production instance as changes may have been made to Confluence that make your customisations unsuable.

Troubleshooting

Did something go wrong?

If you need to retry the upgrade, **you must restore your pre-upgrade backups first**. Do not attempt to run an upgrade again, or start the older version of Confluence again after an upgrade has failed.

Some common issues encountered while upgrading:

- **Cannot proceed with upgrade because license has expired**
If your licence has expired and was not renewed and reapplied before upgrading you will receive errors during the upgrade process. See [upgrading beyond current license period](#) for information on how to resolve this problem.
- **Unable to proceed with upgrade because of a conflict with anti virus**
Some anti-virus or other Internet security tools may interfere with the Confluence upgrade process and prevent the process from completing successfully, particularly if you run Confluence as a Windows service. If you experience or anticipate experiencing such an issue with your anti-virus / Internet security tool, disable this tool first before proceeding with the Confluence upgrade.
- **Database does not support online backups**
The upgrade wizard will prompt you to backup your database using your database's backup utilities. If your database does not support online backups, stop the upgrade process, shut down Confluence, perform your database backup and then run the installer again to continue with the upgrade.
- **Upgrade is taking a very long time**
If you have a very large database (i.e. database backups take a very long time to complete), setting the `confluence.upgrade.recovery.file.enabled` [system property](#) to `false` will speed up the upgrade process. It should be used only when there is a process to back up database and verify the backup before performing an upgrade.

You can also refer to the [Upgrade Troubleshooting](#) guide in the Confluence Knowledge Base, or check for answers from the community at [Atlassian Answers](#).

Upgrading Beyond Current Licensed Period

This page explains the recovery process should you mistakenly try to upgrade your Confluence installation to a version beyond your current license entitlement.

Related pages:

- [Upgrading Confluence](#)
- [Working with Confluence Logs](#)
- [Confluence Administrator's Guide](#)

License warnings

During an upgrade an obvious indication that your license has expired can be found in your [log file](#). You will see a 'WARN' level entry similar to this:

```
[confluence.upgrade.impl.DefaultUpgradeManager] isUpgradeAllowed
Your license is now outside of it's support period.
You need to renew the license before you can upgrade to this version
of Confluence.
```

When you try to access the Confluence site in your browser, you will see the following warning screen:

You cannot access Confluence at present. Look at the table below to identify the reasons

Time	Level	Type	Description	Exception
2008-02-04 10:51:04	null	(EventType: upgrade)	Cannot proceed with upgrade. Your current license does not entitle you to upgrade to this version of Confluence. Please check that the support period of your license has not expired or that you have the correct partner license. If you wish to renew your license, please contact sales@atlassian.com. If you have a new license, please enter it on this page and restart .	fatal

Updating the Confluence license

1. Contact [Atlassian Sales](#) to arrange for a new license to be issued, as instructed on the warning screen illustrated above.
2. Once you have received a suitable license, supply the license key to Confluence:
 - Click link given on the license warning screen, illustrated above.
 - You will first be asked to log in as a Confluence administrator.
 - Then you will be presented with a simplified license administration screen. Enter the credentials of a Confluence system administrator.
 - Copy the license key into the **License** field and choose **Save**.

Please enter your new license.

You must authenticate as a Confluence administrator to do so.

Support Period	Your commercial Confluence support has ended on 07/09/07 15:00 . Confluence updates created after 07/09/07 15:00 are not available under this license. If you wish to renew your license please contact us .
Username	
Password	
License	

3. Restart Confluence to continue the upgrade.

Confluence Post-Upgrade Checks

This article provides a list of items for Confluence Administrators to check after a [Confluence upgrade](#) to ensure that it has completed successfully. This list is not exhaustive, but it does cover common upgrade mistakes.

Before You Begin

After you have completed an upgrade, you should see the following message in the `atlassian-confluence.log` file:

```
2010-03-08 08:03:58,899 INFO [main]
[atlassian.confluence.upgrade.AbstractUpgradeManager] upgradeFinished Upgrade
completed successfully
```

If you do not see the line in your log similar to the one above, this means that your upgrade has not completed successfully. Please check our [Upgrade Troubleshooting](#) documentation to check for a suitable recommendation or fix. If there are no errors logged or if none of the errors are referenced in the the Troubleshooting Upgrades documentation, please contact Atlassian [Support](#).

Upgrade Checklist

Below is a recommended list of items to check after completing an upgrade.

1. Layout and Menu

Visit the [Confluence dashboard](#) and check that it is accessible and displays as expected. Test the different Internet browsers that you have in use in your environment. In addition, confirm that the [layout](#) appears as expected and that the menus are clickable and functioning.

2. Search

Try [searching for content](#), for example pages, attachments or user names. Check that the expected results are returned.

3. Permissions

Confirm that you can visit a page that has [view restrictions](#), but you have permission to view. Confirm that you can edit a page that has edit restrictions but you have permission to edit. Make sure that the permissions of child pages are functioning as well. Involve as many space administrators as possible to confirm they are working. Confirm that anonymous or forbidden users cannot access or modify restricted pages.

4. Attachments

Confirm that [attachments](#) are accessible and searchable.

5. Plugins

Outdated third-party [plugins](#) can cause upgrade failure. Quite often, they will just be incompatible and simply do not work anymore. If you discover that your plugin is no longer working, please check for the latest version for your plugin in the [Atlassian Plugin Exchange](#) or check for compatibility in the [Universal Plugin Manager](#).

Upgrading Confluence EAR-WAR Distribution

This document tells you how to upgrade from one version of Confluence to a later version. These instructions apply to the **EAR-WAR Distribution** of Confluence, deployed on your own existing application server.

If you want to upgrade the regular Confluence distribution, which includes Apache Tomcat as the application server, please refer to [Upgrading Confluence](#) instead.

Please also check the following before you start using this guide:

- The **version of Confluence** that you will be upgrading to. Refer to the [documentation home page](#) to verify the latest Confluence version and to find documentation for older versions.
- The **supported platforms** for the version that you will be upgrading to. Please see the [Supported Platforms](#) page for the version of Confluence that you will be upgrading to, as well as the [End of Support Announcements for Confluence](#).

- If you are running **Confluence on a cluster**, please see [Upgrading a Confluence Cluster](#) instead of this document.

Upgrading to Confluence 5.4?

If so, please review the [Confluence 5.4 Release Notes](#) for important information about this version of Confluence. Ensure that you have read the [Confluence 5.3 Known Issues](#) in the Confluence Knowledge Base.

Also, we strongly recommend that you check the upgrade notes for every major version of Confluence that you are skipping, since there might be specific changes between Confluence versions that could affect your Confluence installation. The upgrade notes for recent major versions of Confluence are accessible from the [Upgrade Notes Overview](#) page.

Finally, please check the [Supported Platforms](#) page to ensure that your Java version, operating system, application server, database and browser are supported for this release of Confluence. The [End of Support Announcements for Confluence](#) page has important information regarding supported platforms.

On this page:

- [Before you Start](#)
- [Backing Up](#)
- [Testing the Upgrade in a Test Environment](#)
- [Performing the Upgrade](#)
- [Reapplying Customisations to your New Confluence](#)
- [Checking for Known Issues and Troubleshooting the Confluence Upgrade](#)

Before you Start

1. If you are planning to change to a different database, we recommend that you complete the Confluence upgrade first. Then follow the instructions on [migrating to a different database](#).
2. Note that you need current [software maintenance](#) to perform the upgrade.
3. Confirm that your [license support period](#) is still valid before you try to upgrade.
4. If your current license has expired but you have a new license with you, please [update your license](#) in Confluence before performing the upgrade.
 If you forget to do this and your license has expired, you will receive errors during the upgrade process. Refer to the instructions on [upgrading beyond current license period](#).
5. Check the [release notes](#) for the new version of Confluence you are installing, plus the upgrade notes for any major versions you are skipping. It is important to read these upgrade notes as there might be specific changes between Confluence versions that could affect your Confluence instance. The upgrade notes pages for recent major versions of Confluence are accessible from the [Upgrade Notes Overview](#) page. (Each upgrade notes page is a 'child' of its respective release notes page.)
6. Make sure that your environment (e.g. the database system, the operating system, the application server and so on) still complies with the Confluence [System Requirements](#). A newer version of Confluence may have different requirements than the previous version.
7. If you are using Confluence EAR-WAR edition, check [Installing the Confluence EAR-WAR Edition](#) to see if there is anything extra you will need to do to get Confluence running.
8. 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 account) has sufficient permissions to modify the database schema.
9. Note which plugins (add-ons) are installed and enabled on your current Confluence site. Please verify whether a compatible version of the plugin is available in the version of Confluence you are upgrading to. This information is available via the **Confluence Upgrade Check** in the plugin administration section of Confluence. See the documentation: [Checking Add-on Compatibility with Application Updates](#). You can also check the respective home pages for these plugins on the [Atlassian Plugin Exchange](#). Once you have confirmed the availability of compatible versions, you should upgrade your plugins after successfully upgrading Confluence. Please test these first by applying them to the latest Confluence version in a test environment.
10. If you have made any customisations to Confluence, please verify their compatibility in the latest version. For example, if you have modified any layouts or are using your own custom theme, please test these first by applying them to the latest Confluence version in a test environment. You can [see the customisations applied to your Confluence installation](#).
11. Some anti-virus or other Internet security tools may interfere with the Confluence upgrade process and

prevent the process from completing successfully. If you experience or anticipate experiencing such an issue with your anti-virus/Internet security tool, disable this tool first before proceeding with the Confluence upgrade.

12. After upgrading, Confluence may need to rebuild its indexes. If this happens, there may be some extra load placed on the server following the upgrade. Make sure to schedule any upgrade of production Confluence outside of hours where people need to use it.

Backing Up

Before you begin the Confluence upgrade, you must back up the following:

1. **Back up your Confluence Home directory.** The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you are using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.
 **Tip:** Another term for 'Home directory' would be 'data 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.  The Confluence installer will automatically prompt you to run a backup, storing the files in a .zip archive at the same level as your Confluence Home directory.
2. **Back up your database.** Perform a manual backup of your external database before proceeding with the upgrade, and double check that the backup was actually created properly. If you are not a database expert, or unfamiliar with the backup-restore facilities of your database, simply restore the backup to a different system to ensure the backup worked before proceeding. This recommendation is generally a good best practice. Surprisingly, many companies get in trouble for broken database backups because they skip this basic but vital "smoke test" of the operation.
 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 Confluence home directory. You should not be using this database for production systems at all, so if you happen to be using HSQLDB in a production system, please migrate to a proper database **before** the upgrade. Read about [the various shortcomings of HSQLDB](#).
3. **Back up your Confluence Installation directory or your Confluence webapp** (if you are using Confluence EAR-WAR edition).  The Confluence installer will automatically back up these files, storing the files in a .zip archive at the same level as your Confluence installation directory. The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.

Testing the Upgrade in a Test Environment

 **Be sure to test the upgrade in a test environment before proceeding on your production server.**

1. Create a snapshot of your current production Confluence environment on a test server, as described in the page on [Moving Confluence Between Servers](#).

XML imports

 Importing an old XML backup file to a new major version (for example, Confluence 3.5 to Confluence 4.0) **is not recommended**. Please recreate your production instance in a test environment first.

2. Perform the upgrade on your cloned environment.
3. Test all your unsupported plugins (add-ons) and any customisations with the new version before proceeding on your production server. You can read more about [supported and unsupported plugins](#).

Performing the Upgrade

 If you are [migrating servers](#) or [migrating databases](#), perform those operations in separate steps.

The upgrade process allows you to unzip the new Confluence installation into a directory of your choice and then edit the configuration files to point your new installation to your existing data files. Follow these instructions:

1. Shut down your existing Confluence instance.
2. Download the Confluence EAR-WAR zip file: Go to the [Download Center](#), and click '**Show all**' to find the EAR-WAR zip file.
3. **If you are on Windows**, please check your unzip program before extracting the downloaded zip file. Some archive-extract programs cause errors when unzipping the Confluence zip file. You should use a third-party unzip program like 7Zip or Winzip. If you do not have one, please download and install one before continuing:
 - [7Zip](#) — Recommended. If in doubt, download the '32-bit.exe' version
 - [Winzip](#)
4. Use your unzip program to unzip the installation file. You should now have a new directory called `confluence-<version>`.
 - In the rest of this document, we will refer to this as the `<Installation-Directory>`.
 - Do not use spaces in your directory path.
 - You can read more about the [Confluence Installation directory](#).
5. Edit the `confluence-init.properties` file found at: `<Installation-Directory>\confluence\WEB-INF\classes\confluence-init.properties` and update '`confluence.home`' to point to your existing [Confluence Home directory](#).
 - Make sure you have first backed up your Home directory.
 - Open the `confluence-init.properties` file in a text editor such as Notepad.
 - Scroll to the bottom and find this line:

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

- Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with `confluence.home`.
 - Update the directory name after the = sign, to point to your existing [Confluence Home directory](#).
6. If you are using Tomcat, you need to update either your `confluence.xml` or `server.xml` (depending on where you have defined the Confluence [context descriptor](#)) to point to the location of the new Confluence installation (also remember to copy over any customisations such as a **tomcat datasource** if you have one).
 7. If you have delegated your user management to JIRA, LDAP or any other external user management system, copy the following files from your old Confluence installation to your new Confluence installation:
 - `<Installation-Directory>/confluence/WEB-INF/classes/osuser.xml`.
 - `<Installation-Directory>/confluence/WEB-INF/classes/atlassian-user.xml` (if you are upgrading from Confluence 2.2 or later).

**Upgrading to Confluence 3.5+ and using JIRA user management?**

Please review our KB article first: [Upgrade to Confluence 3.5 with JIRA User Management Fails](#)



If you are upgrading from an earlier version of Confluence (2.5.5 and earlier) and are copying your existing `atlassian-user.xml` file from your previous instance, please ensure that the hibernate cache parameter in this file has been enabled, to avoid performance related issues. (NOTE: If you use Crowd for your user management, you do not need to do this.):

```
<hibernate name="Hibernate Repository"
key="hibernateRepository" description="Hibernate
Repository" cache="true" />
```

8. If you have delegated your user management to Crowd, you will also need to copy the Crowd client library and configuration files from your old Confluence installation to your new Confluence installation: `<Installation-Directory>/confluence/WEB-INF/classes/crowd`

lation-Directory>/confluence/WEB-INF/lib/crowd-integration-client-X.X.X.jar and <Installation-Directory>/confluence/WEB-INF/classes/crowd.properties. If you need more information, please refer to the [Crowd documentation](#).

9. Restart your application server and start Confluence.

i Please note that Confluence will need to re-index attachments and this can take 5-10 minutes. Please wait until Confluence has finished indexing the attachments before trying to access Confluence via your web browser. (There is no easy and quick way to determine if the indexing process is completed. Please wait for approximately 10 minutes after the server start up before accessing Confluence via a web browser.)

10. During the startup process Confluence will create any missing database indexes. If you created any database indexes on your own, please check those afterwards and remove those that duplicate the indexes added by Confluence. Just in case you run into any errors which prevent Confluence from starting up, you can set the system property `hibernate.hbm2ddl.skip_creating_missing_indexes` to `true` to skip automatic index creation.
11. Visit Confluence in your web browser and log in using a username from your previous Confluence installation. You should be able to log in immediately, without seeing the Setup Wizard.
12. Take a quick look around your Confluence site to confirm that all your spaces and pages are present and everything looks normal. You should see the new Confluence version number in the page footer.
13. Consider any adjustments you need to make to customisations and special configurations, as described [below](#).

Reapplying Customisations to your New Confluence



Hint: The steps below are for advanced Confluence users, who have applied special settings to their Confluence server and/or Confluence look and feel

After upgrading your Confluence installation to a later version of Confluence, you need to consider any customisations you have applied to your system and other special configurations:

- If you had previously installed **Confluence/Tomcat as a Windows service**, uninstall the service (to ensure that the old Confluence cannot start automatically when the server restarts) and reinstall the new one. For details please see [Start Confluence Automatically on Windows as a Service](#).
- If you are using the Confluence distribution 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 had previously connected your Confluence installation to an external database via a JNDI datasource or you implemented **SSL**, edit your new `web.xml` file and copy over any relevant modifications from your old `web.xml` file, which relate to these customisations.
- If you were previously running **Confluence on a non-standard port**, edit your new <Installation-Directory>\conf\server.xml file as described in [Change listen port for Confluence](#).
- If you had previously defined a **Tomcat datasource**, edit your new <Installation-Directory>\conf\server.xml and copy over the datasource definition from your old server.xml.
- If you were previously using any **plugins**, install the latest compatible version and disable any plugins that are incompatible with your new version of Confluence. The easiest way to do this is to use the [Plugin Repository](#) in the Confluence Administration Console.
- 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](#).
- If you had previously modified the Confluence **source code**, you will need to reapply your changes to the new version.
- 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 had previously modified the **memory flags** (Xms and Xmx) in either the <Installation-Directory>\bin\setenv.sh or the <Installation-Directory>\bin\setenv.bat file, you may want to make the modifications in your new installation. The parameters are specified in the JAVA_OPTS variable.
- If you had changed the [Confluence interface text](#), you will need to pull over the ConfluenceActionSupport.properties file.

- If you were using a custom SSO authenticator, change seraph-config.xml to the correct authenticator.

Checking for Known Issues and Troubleshooting the Confluence Upgrade

After you have completed the steps required to upgrade your Confluence installation, check all the items on the [Confluence post-upgrade checklist](#) to ensure that everything works as expected. If something is not working correctly, please check for known Confluence issues and try troubleshooting your upgrade as described below:

- **Check for known issues.** Sometimes we find out about a problem with the latest version of Confluence after we have released the software. In such cases we publish information about the known issues in the Confluence Knowledge Base. Please check the known issues for the relevant release on this page of the [Knowledge Base](#) and follow the instructions to solve the problem.
- **Check for answers from the community.** Other users may have encountered the same issue. You can check for answers from the community at [Atlassian Answers](#).
- **Did you encounter a problem during the Confluence upgrade?** Please refer to the guide to [troubleshooting upgrades](#) in the Confluence Knowledge Base.

RELATED TOPICS

[Upgrading Confluence](#)
[Upgrading Confluence](#)
[Confluence Installation Guide](#)
[Important Directories and Files](#)
[Site Backup and Restore](#)
[Database Configuration](#)

Migration from Wiki Markup to XHTML-Based Storage Format

If you are upgrading **to Confluence 4.0 or later** from an older version (From Confluence 3.5.x or earlier) then as part of the upgrade an automatic migration of your content will take place. This is a non-destructive process. Your existing content is not overwritten. Instead, the migration process will create a new version of each wiki markup page. The new version will use the new XHTML-based storage format, so that you can edit the page in the Confluence rich text editor.

In addition, if you are **upgrading to Confluence 4.3 or later** from an older version then as part of the upgrade an automatic migration of your page templates will take place. See [Migration of Templates from Wiki Markup to XHTML-Based Storage Format](#).

Note: Even though the process is non-destructive, you must be sure to perform a backup of your database and home directory prior to starting the new version of Confluence, as we recommend for any [Confluence upgrade](#).

Migration process

Depending on the size of your Confluence installation, the migration from wiki markup to the new XHTML-based storage format could prove time consuming. The duration of the migration is difficult to estimate; this is due to a number of site specific factors. As a rough guide, a test dataset we migrated was 130,000 pages, totalling approximately 700Mb, which took six minutes.

On this page:

- [Migration process](#)
- [Watching the migration logs during the upgrade](#)
- [Re-running the migration – for content that completely failed the migration](#)
- [Re-attempting the migration – for content in 'unmigrated-wiki-markup' macro](#)
- [Notes](#)

Related pages:

- [Migration of Templates from Wiki Markup to XHTML-Based Storage Format](#)
- [Upgrading Confluence](#)

The following properties that can be modified to allow finer control over the migration process:

Property	Purpose	Default

confluence.wiki.migration.threads	The number of concurrent worker threads migrating content	4
confluence.wiki.migration.batch.size	The number of items migrated in each batch of work	500
confluence.wiki.migration.versioncomment	The comment associated with the newly migrated version of each piece of content	"Migrated to Confluence 4.0"

(For instructions on setting Confluence system properties see [this document](#).)

Again, due to the large variability in Confluence installations it is hard to give specific recommendations for the above settings. One point to note though that both increasing batch size and the number of threads (or both) will increase the peak memory required for migration. If memory is an issue then as you increase one of these settings consider decreasing the other.

Another factor to be aware of if modifying these defaults is that of the cache settings employed in your site. The migration will quickly populate certain Confluence caches so be sure that if you have customised caches as [described here](#) that there is enough memory on the server for these caches should they reach maximum capacity.

Watching the migration logs during the upgrade

To monitor the progress of a site migration you should watch the output in the [application log](#).

Typical logging progress will be shown by multiple log entries at the INFO level of the following format:

```
WikiToXhtmlMigrationThread-n - Migrated 2500 of 158432 pages, this batch
migrated 500/500 without error
```

There may be a wide array of messages logged from each individual page but any errors are also collected for display in a single migration report once all content has been processed. Here is a typical example of such a report:

```
Wiki to XHTML Exception Report:
Summary:
  0 settings values failed.
  0 PageTemplates failed.
  2 ContentEntityObjects failed.
Content Exceptions:
  1) Type: page, Id: 332, Title: Release Notes 1.0b3, Space: DOC -
Confluence 4.0 Beta. Cause:
com.atlassian.confluence.content.render.xhtml.migration.exceptions.Unkno
wnMacroMigrationException: The macro link is unknown.. Message: The
macro link is unknown.
  2) Type: comment, Id: 6919, Title: null, Global Scope. Cause:
com.atlassian.confluence.content.render.xhtml.migration.exceptions.Unkno
wnMacroMigrationException: The macro mymacro is unknown.. Message: The
macro mymacro is unknown.
```

Each entry in the report will identify the content that caused migration exceptions as well as displaying the exceptions themselves.

In almost all cases any content reported as errored will have been migrated to the new XHTML-based storage format, but will actually consist of wiki markup content wrapped within an XML 'unmigrated-wiki-markup' macro. This content will still be viewable in Confluence and editable within the new Confluence Editor.

However, in some cases a batch of content may actually have completely failed to migrate. This is most typically due to an unhandled exception causing a database transaction rollback. This would be reported in the log with a message like this:

```
Unable to start up Confluence. Fatal error during startup sequence:  
confluence.lifecycle.core:pluginframeworkdependentupgrades (Run all the  
upgrades that require the plugin framework to be available) -  
com.atlassian.confluence.content.render.xhtml.migration.exceptions.Migra  
tionException: java.util.concurrent.ExecutionException:  
org.springframework.transaction.UnexpectedRollbackException: Transaction  
rolled back because it has been marked as rollback-only
```

Confluence provides no further report about this scenario and will also allow Confluence to restart as normal without retrying a migration. If a user tries to view any such unmigrated content they will see an exception similar to this:

```
java.lang.UnsupportedOperationException: The body of this  
ContentEntityObject ('Page Title') was 'WIKI' but was expected to be  
'XHTML'
```

The solution is to ensure you manually re-run the site migration after the restart.

Re-running the migration – for content that completely failed the migration

A Confluence Administrator can restart the site migration if there was any content that failed migration (see previous section). Only the content that is still formatted in wiki markup will be migrated, so typically a re-migration will take less time than the original migration.

To manually re-run migration:

1. Open this URL in your browser: <Confluence Address>/admin/force-upgrade.action
2. Select **wikiToXhtmlMigrationUpgradeTask** in the **Upgrade task to run** dropdown list.
3. Choose **Force Upgrade**.

Force Upgrade

This will run the selected upgrade task. You should only run any of these tasks if you were directed here by support documentation.

Upgrade task to run

Migrate all wiki formatted content in this Confluence instance to XHTML

Re-attempting the migration – for content in 'unmigrated-wiki-markup' macro

The previous section was about dealing with the exceptional circumstance where certain content was left completely unmigrated. The most common migration problem is that the content was migrated but remains formatted as wiki markup on the page, within the body of an 'unmigrated-wiki-markup' macro. Any content which is referenced in the migration report will be found in this state. This content is still viewable and editable but since it is wiki markup it cannot be edited using the full feature set of the rich text editor.

The most common reason for content to be in this state is that the page contains an unknown macro, or a macro that is not compatible with Confluence 4.x.

There are two possible fixes for this situation:

1. Install a version of the macro that is compatible with Confluence 4.x. See [Plugin Development Upgrade FAQ for 4.0](#).

2. Edit the page and remove the problematic macro.

Regardless of the solution you choose, you can then force a re-migration of all the content (including content in templates) that was left wrapped in an 'unmigrated-wiki-markup' macro. This feature is found at <Confluence Address>/admin/unmigratedcontent.action

Update content with incompatible macros

Confluence has detected that there are **0 pages** with macros that are not yet Confluence 4+ compatible. To ensure backwards compatibility, these macros are still being rendered as wiki markup when editing your pages.

If you have recently updated plugins, you should update your content to ensure that any macros that are not Confluence 4 compatible become compatible. You may have to run the update several times as you update incompatible macros.

Update Check

Update not required

You have not installed any new plugins since your last content upgrade. You do not need to run this upgrade unless you have been advised to by Atlassian Support staff.

Note: Once an upgrade has commenced you will not be able to pause or undo the upgrade. An update can severely affect the performance of your instance, we recommend you conduct this update during a quiet time. Users editing a page as it is updated may receive notice of a conflicting edit.

[Update Content](#)

Notes

We refer to the Confluence storage format as 'XHTML-based'. To be correct, we should call it XML, because the Confluence storage format does not comply with the XHTML definition. In particular, Confluence includes custom elements for macros and more. We're using the term 'XHTML-based' to indicate that there is a large proportion of HTML in the storage format.

Migration of Templates from Wiki Markup to XHTML-Based Storage Format

If you are **upgrading to Confluence 4.3 or later from an older version** (from Confluence 4.2.x or earlier) then as part of the upgrade an automatic migration of your page templates will take place. This is a non-destructive process. Your existing content is not overwritten. Instead, the migration process will create a new version of each space template and each global template on your Confluence site. The new version will use the new XHTML-based storage format, so that you can edit the template in the Confluence rich text editor.

Note: Nevertheless, you must be sure to perform a backup of your database and home directory prior to starting the new version of Confluence, as we recommend for any [Confluence upgrade](#).

Watching the migration logs during the upgrade

To monitor the progress of a site migration you should watch the output in the [application log](#).

A typical logging progress will be shown by multiple log entries at the INFO level of the following format:

```
WikiToXhtmlMigrationThread-n - Migrated 22 of 29 PageTemplates.
```

On this page:

- [Watching the migration logs during the upgrade](#)
- [Re-running the migration](#)
- [Notes](#)

Related pages:

- [Migration from Wiki Markup to XHTML-Based Storage Format](#)
- [Working with Templates](#)
- [Upgrading Confluence](#)

There may be a wide array of messages logged from each individual template, but any errors are also collected for display in a single migration report once all content has been processed. Here is a typical example of such a report:

Wiki to XHTML Exception Report:

Summary:

- 0 settings values failed.
- 2 PageTemplates failed.
- 0 ContentEntityObjects failed.

Content Exceptions:

- 1) Type: page, Id: 332, Title: Release Notes 1.0b3, Space: DOC - Confluence 4.0 Beta. Cause: `com.atlassian.confluence.content.render.xhtml.migration.exceptions.UnknownMacroMigrationException: The macro link is unknown.. Message: The macro link is unknown.`
- 2) Type: comment, Id: 6919, Title: null, Global Scope. Cause: `com.atlassian.confluence.content.render.xhtml.migration.exceptions.UnknownMacroMigrationException: The macro mymacro is unknown.. Message: The macro mymacro is unknown.`

Each entry in the report will identify the content that caused migration exceptions as well as displaying the exceptions themselves.

In almost all cases any content reported as errored will have been migrated to the new XHTML-based storage format, but will actually consist of wiki markup content wrapped within an XML 'unmigrated-wiki-markup' macro. This content will still be viewable in Confluence and editable within the Confluence rich text editor.

However, in some cases a batch of content may actually have completely failed to migrate. This is most typically due to an unhandled exception causing a database transaction rollback. This would be reported in the log with a message like this:

```
Unable to start up Confluence. Fatal error during startup sequence:
confluence.lifecycle.core:pluginframeworkdependentupgrades (Run all the
upgrades that require the plugin framework to be available) -
com.atlassian.confluence.content.render.xhtml.migration.exceptions.MigrationException: java.util.concurrent.ExecutionException:
org.springframework.transaction.UnexpectedRollbackException: Transaction
rolled back because it has been marked as rollback-only
```

Confluence provides no further report about this scenario and will also allow Confluence to restart as normal without retrying a migration. If a user tries to view or edit an unmigrated template, the wiki template editor will be used.

The solution is to manually re-run the site migration after the restart, as described below.


Re-running the migration

A Confluence administrator can restart the template migration if any templates have failed the migration (see previous section). Only the templates that are still formatted in wiki markup will be migrated again. Typically, a re-migration will take less time than the original migration.

To manually re-run the migration:

1. Open this URL in your browser: <Confluence Address>/admin/force-upgrade.action
2. Select **pageTemplateWikiToXhtmlMigrationUpgradeTask** in the **Upgrade task to run** dropdown list.
3. Choose **Force Upgrade**.

Screenshot: The 'Force Upgrade' screen in the Confluence administration console



Notes

We refer to the Confluence storage format as 'XHTML-based'. To be correct, we should call it XML, because the Confluence storage format does not comply with the XHTML definition. In particular, Confluence includes custom elements for macros and more. We're using the term 'XHTML-based' to indicate that there is a large proportion of HTML in the storage format.

Upgrading Confluence Manually

This document tells you how to upgrade from one version of Confluence to a later version. This document refers to the Confluence distribution that includes Apache Tomcat as the bundled application server. If you want to upgrade an **EAR/WAR distribution** deployed on your own existing application server, please refer to [Upgrading Confluence EAR-WAR Distribution](#) instead.

Please also check the following before you start using this guide:

- The **version of Confluence** that you will be upgrading to. Refer to the [documentation home page](#) to verify the latest Confluence version and to find documentation for older versions.
- The **supported platforms** for the version that you will be upgrading to. Please see the [Supported Platforms](#) page for the version of Confluence that you will be upgrading to, as well as the [End of Support Announcements for Confluence](#).
- If you are running **Confluence on a cluster**, please see [Upgrading a Confluence Cluster](#) instead of this document.

Upgrading to Confluence 5.4?

If so, please review the [Confluence 5.4 Release Notes](#) for important information about this version of Confluence. Ensure that you have read the [Confluence 5.3 Known Issues](#) in the Confluence Knowledge Base.

Also, we strongly recommend that you check the upgrade notes for every major version of Confluence that you are skipping, since there might be specific changes between Confluence versions that could affect your Confluence installation. The upgrade notes for recent major versions of Confluence are accessible from the [Upgrade Notes Overview](#) page.

Finally, please check the [Supported Platforms](#) page to ensure that your Java version, operating system, application server, database and browser are supported for this release of Confluence. The [End of Support Announcements for Confluence](#) page has important information regarding supported platforms.

On this page:

- [Before you Start](#)
- [Backing Up](#)
- [Testing the Upgrade in a Test Environment](#)
- [Performing the Upgrade](#)
- [Reapplying Customisations to your New Confluence](#)
- [Checking for Known Issues and Troubleshooting the Confluence Upgrade](#)
- [Useful Plugins](#)

Before you Start

1. If you are planning to change to a different database, we recommend that you complete the Confluence upgrade first. Then follow the instructions on [migrating to a different database](#).
2. Note that you need current [software maintenance](#) to perform the upgrade.
3. Confirm that your [license support period](#) is still valid before you try to upgrade.
4. If your current license has expired but you have a new license with you, please [update your license](#) in Confluence before performing the upgrade.
 If you forget to do this and your license has expired, you will receive errors during the upgrade process. Refer to the instructions on [upgrading beyond current license period](#).
5. Check the [release notes](#) for the new version of Confluence you are installing, plus the upgrade notes for any major versions you are skipping. It is important to read these upgrade notes as there might be specific changes between Confluence versions that could affect your Confluence instance. The upgrade notes pages for recent major versions of Confluence are accessible from the [Upgrade Notes Overview](#) page. (Each upgrade notes page is a 'child' of its respective release notes page.)
6. Make sure that your environment (e.g. the database system, the operating system, the application server and so on) still complies with the Confluence [System Requirements](#). A newer version of Confluence may have different requirements than the previous version.
7. If you are using Confluence EAR-WAR edition, check [Installing the Confluence EAR-WAR Edition](#) to see if there is anything extra you will need to do to get Confluence running.
8. 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 account) has sufficient permissions to modify the database schema.
9. Note which plugins (add-ons) are installed and enabled on your current Confluence site. Please verify whether a compatible version of the plugin is available in the version of Confluence you are upgrading to. This information is available via the **Confluence Upgrade Check** in the plugin administration section of Confluence. See the documentation: [Checking Add-on Compatibility with Application Updates](#). You can also check the respective home pages for these plugins on the [Atlassian Plugin Exchange](#). Once you have confirmed the availability of compatible versions, you should upgrade your plugins after successfully upgrading Confluence. Please test these first by applying them to the latest Confluence version in a test environment.
10. If you have made any customisations to Confluence, please verify their compatibility in the latest version. For example, if you have modified any layouts or are using your own custom theme, please test these first by applying them to the latest Confluence version in a test environment. You can [see the customisations applied to your Confluence installation](#).
11. Some anti-virus or other Internet security tools may interfere with the Confluence upgrade process and prevent the process from completing successfully. If you experience or anticipate experiencing such an issue with your anti-virus/Internet security tool, disable this tool first before proceeding with the Confluence upgrade.
12. After upgrading, Confluence may need to rebuild its indexes. If this happens, there may be some extra load placed on the server following the upgrade. Make sure to schedule any upgrade of production Confluence outside of hours where people need to use it.

Backing Up

Before you begin the Confluence upgrade, you must back up the following:

1. **Back up your Confluence Home directory.** The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you are using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.
 **Tip:** Another term for 'Home directory' would be 'data 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.  The Confluence installer will automatically prompt you to run a backup, storing the files in a .zip archive at the same level as your Confluence Home directory.

2. **Back up your database.** Perform a manual backup of your external database before proceeding with the upgrade, and double check that the backup was actually created properly. If you are not a database expert, or unfamiliar with the backup-restore facilities of your database, simply restore the backup to a different system to ensure the backup worked before proceeding. This recommendation is generally a good best practice. Surprisingly, many companies get in trouble for broken database backups because they skip this basic but vital "smoke test" of the operation.

 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 Confluence home directory. You should not be using this database for production systems at all, so if you happen to be using HSQLDB in a production system, please migrate to a proper database **before** the upgrade. Read about [the various shortcomings of HSQLDB](#).

3. **Back up your Confluence Installation directory or your Confluence webapp** (if you are using Confluence EAR-WAR edition).  The Confluence installer will automatically back up these files, storing the files in a .zip archive at the same level as your Confluence installation directory. The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.

Testing the Upgrade in a Test Environment

 **Be sure to test the upgrade in a test environment before proceeding on your production server.**

1. Create a snapshot of your current production Confluence environment on a test server, as described in the page on [Moving Confluence Between Servers](#).

XML imports

 Importing an old XML backup file to a new major version (for example, Confluence 3.5 to Confluence 4.0) **is not recommended**. Please recreate your production instance in a test environment first.

2. Perform the upgrade on your cloned environment.
3. Test all your unsupported plugins (add-ons) and any customisations with the new version before proceeding on your production server. You can read more about [supported and unsupported plugins](#).

Performing the Upgrade

 If you are [migrating servers](#) or [migrating databases](#), perform those operations in separate steps.

To install Confluence, unzip the new Confluence installation zip file into a directory of your choice and then edit the configuration files to point your new installation to your existing data files. Follow these instructions:

1. Shut down your existing Confluence instance.
2. Download the [Confluence zip file](#).
3. **If you are on Windows**, please check your unzip program before extracting the downloaded zip file. Some archive-extract programs cause errors when unzipping the Confluence zip file. You should use a third-party unzip program like 7Zip or Winzip. If you do not have one, please download and install one before continuing:
 - [7Zip](#) — Recommended. If in doubt, download the '32-bit.exe' version
 - [Winzip](#)
4. Use your unzip program to unzip the installation file. You should now have a new directory called `confluence-<version>`, e.g. `confluence-4.0.0-std`.
 - In the rest of this document, we will refer to this as the `<Installation-Directory>`.

- If you decide to change the location from the default, make sure that you choose a different location from your existing Confluence installation, because legacy files may cause problems if you install the new Confluence version into an existing directory.
 - Do not use spaces in your directory path.
 - You can read more about the [Confluence Installation directory](#).
5. Edit the `confluence-init.properties` file found at: `<Installation-Directory>\confluence\WEB-INF\classes\confluence-init.properties` and update `'confluence.home'` to point to your **existing** Confluence Home directory.
- You can read more about the [Confluence Home directory](#).
 - Make sure you have first backed up this directory, as instructed above.
 - Open the `confluence-init.properties` file in a text editor such as Notepad.
 - Scroll to the bottom and find this line:

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

- Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with `confluence.home`.
 - Update the directory name after the = sign, to point to your existing Confluence Home directory.
6. If you are running Confluence as a Windows service, use the command prompt and type `<Installation-Directory>\bin\service.bat remove Confluence`.



It is vital that you stop and remove the existing service *prior to uninstalling* the old instance of Confluence! For more information on running Confluence as Windows service, please refer to the [Start Confluence Automatically on Windows as a Service](#) topic.



To remove the service installed by the [Confluence installer](#), you need to run the `<confluence auto installer installation folder>\UninstallService.bat`.

7. If you are using an external database (i.e. not the embedded HSQLDB database supplied for evaluation purposes), copy the jdbc driver jar file from your old Confluence installation to the new Confluence installation. The jdbc driver jar file in the old Confluence installation should be located in either the `<Install-Directory>/common/lib` or `<Installation-Directory>/confluence/WEB-INF/lib` directories. Once you have identified this file, copy it to either the `<Install-Directory>/lib` or `<Installation-Directory>/confluence/WEB-INF/lib` directories of your Confluence installation.
8. If you have delegated your user management to JIRA, LDAP, Crowd, or any other external user management system, copy the following files from your old Confluence installation to your new Confluence installation:
- `<Installation-Directory>/confluence/WEB-INF/classes/osuser.xml`.
 - `<Installation-Directory>/confluence/WEB-INF/classes/atlassian-user.xml` (if you are upgrading from Confluence 2.2 or later).



If you are upgrading from an earlier version of Confluence (2.5.5 and earlier) and are copying your existing `atlassian-user.xml` file from your previous instance, please ensure that the hibernate cache parameter in this file has been enabled, to avoid performance related issues. (NOTE: If you use Crowd for your user management, you do not need to do this.):

```
<hibernate name="Hibernate Repository"
key="hibernateRepository" description="Hibernate
Repository" cache="true" />
```

9. If you have delegated your user management to Crowd, you will also need to copy the Crowd configuration file from your old Confluence installation to your new Confluence installation: `<Installation-Directory>/confluence/WEB-INF/classes/crowd.properties`. If you need more information, please refer to the [Crowd documentation](#).

10. Consider any adjustments you need to make to customisations and special configurations, as described [below](#).
 Your new version of Confluence may not function correctly or could encounter problems or errors if these are not implemented.
11. Start your new version of Confluence.
 **Please note** that Confluence will need to re-index attachments and this can take 5-10 minutes. Please wait until Confluence has finished indexing the attachments before trying to access Confluence via your web browser.
12. During the startup process Confluence will create any missing database indexes. If you created any database indexes on your own, please check those afterwards and remove those that duplicate the indexes added by Confluence. Just in case you run into any errors which prevent Confluence from starting up, you can set the system property `hibernate.hbm2ddl.skip_creating_missing_indexes` to `true` to skip automatic index creation.
13. Visit Confluence in your web browser and log in using a username from your previous Confluence installation. You should be able to log in immediately, without seeing the Setup Wizard.
14. Take a quick look around your Confluence site to confirm that all your spaces and pages are present and everything looks normal. You should see the new Confluence version number in the page footer.

Reapplying Customisations to your New Confluence



Hint: The steps below are for advanced Confluence users, who have applied special settings to their Confluence server and/or Confluence look and feel

After upgrading your Confluence installation to a later version of Confluence, you need to consider any customisations you have applied to your system and other special configurations:

- If you had previously installed **Confluence/Tomcat as a Windows service**, uninstall the service (to ensure that the old Confluence cannot start automatically when the server restarts) and reinstall the new one. For details please see [Start Confluence Automatically on Windows as a Service](#).
- If you are using the Confluence distribution 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 had previously connected your Confluence installation to an external database via a JNDI datasource or you implemented [SSL](#), edit your new `web.xml` file and copy over any relevant modifications from your old `web.xml` file, which relate to these customisations.
- 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](#).
- If you had previously defined a **Tomcat datasource**, edit your new `<Installation-Directory>\conf\server.xml` and copy over the datasource definition from your old `server.xml`.
- If you were previously using any **plugins**, install the latest compatible version and disable any plugins that are incompatible with your new version of Confluence. The easiest way to do this is to use the [Plugin Repository](#) in the Confluence Administration Console.
- 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](#).
- If you had previously modified the Confluence **source code**, you will need to reapply your changes to the new version.
- 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 had previously modified the **memory flags** (`Xms` and `Xmx`) in either the `<Installation-Direct>\bin\setenv.sh` or the `<Installation-Directory>\bin\setenv.bat` file, you may want to make the modifications in your new installation. The parameters are specified in the `JAVA_OPTS` variable.
- If you had changed the [Confluence interface text](#), you will need to pull over the `ConfluenceActionSupport.properties` file.
- If you were using a custom SSO authenticator, change `seraph-config.xml` to the correct authenticator.

Checking for Known Issues and Troubleshooting the Confluence Upgrade

After you have completed the steps required to upgrade your Confluence installation, check all the items on the [Confluence post-upgrade checklist](#) to ensure that everything works as expected. If something is not working correctly, please check for known Confluence issues and try troubleshooting your upgrade as described below:

- **Check for known issues.** Sometimes we find out about a problem with the latest version of Confluence after we have released the software. In such cases we publish information about the known issues in the Confluence Knowledge Base. Please check the known issues for the relevant release on this page of the [Knowledge Base](#) and follow the instructions to solve the problem.
- **Check for answers from the community.** Other users may have encountered the same issue. You can check for answers from the community at [Atlassian Answers](#).
- **Did you encounter a problem during the Confluence upgrade?** Please refer to the guide to [troubleshooting upgrades](#) in the Confluence Knowledge Base.

Useful Plugins

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](#).

- Appfire's [Upgrade Assistant for Confluence \(UAC\)](#) is a commercial plugin that simplifies the upgrade process into an easy-to-use wizard.

RELATED TOPICS

[Upgrading Confluence](#)
[Upgrading Confluence EAR-WAR Distribution](#)
[Confluence Installation Guide](#)
[Important Directories and Files](#)
[Site Backup and Restore](#)
[Database Configuration](#)

Supported Platforms

This page describes the supported platforms for Confluence. Please review them before installing Confluence. The information on this page applies to **Confluence 5.5**.

Further information:

- End of support for various platforms and browsers when used with Confluence: [End of Support Announcements for Confluence](#).
- More information about these supported platforms and hardware requirements: [System Requirements](#).

Related pages:

- [Confluence Installation Guide](#)
- [Confluence Setup Guide](#)
- [Installing Confluence and JIRA Together](#)
- [Server Hardware Requirements Guide](#)
- [Supported Platforms FAQ](#)
- [Confluence Documentation Home](#)

Key:  = Supported.  = Not Supported

Java version	
Oracle JRE / JDK	 1.7
Operating systems for Confluence server installation	
Microsoft Windows (including 64-bit) ^(Supported Platforms v5.5#1)	 (Microsoft Supported Versions only)

Linux / Solaris (Supported Platforms v5.5#1 , Supported Platforms v5.5#2)	✓
Apple Mac OS X	✗ Not supported as server. ✓ Supported as client platform.
Application servers	
Apache Tomcat	✓ 6.0.x, 7.0.x
Databases	
PostgreSQL	✓ 8.4, 9.0, 9.1, 9.2, 9.3
MySQL (Supported Platforms v5.5#3)	✓ 5.1, 5.5, 5.6.16
Oracle	✓ 11.1, 11.2
Microsoft SQL Server	✓ 2008, 2008 R2, 2012
HSQLDB (Supported Platforms v5.5#4)	✓ (for evaluation purposes only)
Web browsers – desktop	
Microsoft Internet Explorer (Windows) (Supported Platforms v5.5#5 , Supported Platforms v5.5#6)	✓ 8, 9, 10, 11
Mozilla Firefox (all platforms)	✓ Latest stable version supported
Google Chrome (Windows and Mac) (Supported Platforms v5.5#7)	✓ Latest stable version supported
Safari (Mac)	✓ Latest stable version supported
Web browsers – mobile	
Mobile Safari (iOS) (Supported Platforms v5.5#8)	✓ Latest stable version supported
Android (Android) (Supported Platforms v5.5#9)	✓ The default browser on Android 4.0.3 (Ice Cream Sandwich)
Chrome (Android and iOS) (Supported Platforms v5.5#8 , Supported Platforms v5.5#9)	✓ Latest stable version supported

1. Confluence is a pure Java application and should run on this platform provided the JRE or JDK requirement is satisfied.
2. While some of our customers run Confluence on SPARC-based hardware, Atlassian only officially supports Confluence running on x86 hardware and 64-bit derivatives of x86 hardware.
3. Ensure that you configure your Confluence MySQL database to use the InnoDB storage engine as the MyISAM storage engine could lead to data corruption.
4. Confluence ships with a built-in HSQL database. While this database is fine for evaluation purposes, it is somewhat susceptible to data loss during system crashes. Hence, for production environments, we recommend that you configure Confluence to use an [external database](#).
5. Internet Explorer 8 and 9 do not support the drag-and-drop functionality of HTML5. As Confluence relies on this functionality, the drag-and-drop experience in Internet Explorer 8 and 9 is not complete. Internet Explorer 10 in 'desktop' mode **does** support the drag-and-drop functionality, and the implementation of drag-and-drop in Confluence works as expected with Internet Explorer 10 'desktop' mode. The 'modern' mode of Internet Explorer 10 does not support drag-and-drop.
6. Confluence is tested with these versions of Internet Explorer in standards-compliant rendering mode, not compatibility mode. Enabling compatibility mode may cause problems because it emulates older, unsupported rendering modes.

7. Chrome does not have WEBDAV / plugin support so features such as [Edit in Word](#) for attachments will not work. Please refer to [CONF-23322](#) for information on the progress of the issue.
8. Confluence does not support editing in Mobile Safari on iOS devices (such as iPhone and iPad). Please refer to [CONF-19523](#) for information on the progress of this issue.
9. Confluence does not support editing on Android devices.

End of Support Announcements for Confluence

This page contains announcements of the end of support for various platforms and browsers when used with Confluence. This is summarised in the table below. Please see the sections following for the full announcements.

End of Support Matrix for Confluence

The table below summarises information regarding the end of support announcements for **upcoming** Confluence releases. If a platform (version) has already reached its end of support date, it is **not** listed in the table.

Platform	Confluence End of Support
Apache Tomcat 6	After Confluence 5.5.x (announcement)
PostgreSQL 8.3	With release of Confluence 5.5 (announcement)
Internet Explorer 8	After Confluence 5.5.x (announcement)



Why is Atlassian ending support for these platforms?

Atlassian is committed to delivering improvements and bug fixes as fast as possible. We are also committed to providing world class support for all the platforms our customers run our software on. However, as the complexity of our applications grows, the cost of supporting multiple platforms increases exponentially. Each new feature has to be tested on several combinations of application servers, databases, web browsers, etc, with setup and ongoing maintenance of automated tests. Moving forward, we want to reduce the time spent there to increase Confluence development speed significantly.

On this page (most recent announcements first):

- [Deprecated Tomcat platform for Confluence \(22 April 2014\)](#)
- [Deprecated Databases for Confluence \(2 December 2013\)](#)
- [Deprecated Web Browsers for Confluence \(24 September 2013\)](#)
- [Deprecated Databases for Confluence \(13 August 2013\)](#)
- [Deprecated Tomcat platform for Confluence \(29 August 2012\)](#)
- [Deprecated Java platform for Confluence \(6 August 2012\)](#)
- [Deprecated Databases for Confluence \(1 May 2012\)](#)
- [Deprecated Databases for Confluence \(13 March 2012\)](#)
- [Deprecated Operating Systems for Confluence \(21 July 2011\)](#)
- [Deprecated Databases for Confluence \(7 January 2011\)](#)
- [Deprecated Web Browsers for Confluence \(7 January 2011\)](#)
- [Deprecated Databases for Confluence \(12 October 2010\)](#)
- [Deprecated Web Browsers for Confluence \(12 October 2010\)](#)
- [Deprecated Databases for Confluence \(6 July 2010\)](#)
- [Deprecated Web Browsers for Confluence \(6 July 2010\)](#)
- [Deprecated Databases for Confluence \(24 March 2010\)](#)
- [Deprecated Application Servers for Confluence \(27 January 2010\)](#)
- [Deprecated Java Platforms for Confluence \(27 January 2010\)](#)
- [Deprecated Web Browsers for Confluence \(14 December 2009\)](#)

Deprecated Tomcat platform for Confluence (22 April 2014)

This section announces the end of Atlassian support for Tomcat 6.0.x for Confluence.

End of support means that Atlassian will not fix bugs related to the specified version of Tomcat, past the support end date for your version of Confluence. The details are below. Please refer to the list of [supported platforms](#) for details of platform support for Confluence. If you have questions or concerns regarding this

announcement, please email [eol-announcement at atlassian dot com](mailto:eol-announcement@atlassian.com).

End of Life Announcement for Tomcat 6.0.x Support

Platform	Support End Date
Tomcat 6.0.x	When Confluence 5.6 is released, due in mid 2014

Tomcat 6.0.x notes:

- Confluence 5.5 is the last major version that will support Tomcat 6.0.x. The Confluence 5.5.x bug-fix releases will also continue to support Tomcat 6.0.x.
- Confluence 5.5.x and previously-released versions will continue to work with Tomcat 6.0.x. However, we will not fix bugs affecting Tomcat 6.0.x after the [end-of-life date](#) for your version of Confluence.
- Confluence 5.6 will not be tested with Tomcat 6.0.x.

Deprecated Databases for Confluence (2 December 2013)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of [supported platforms](#) for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email [eol-announcement at atlassian dot com](mailto:eol-announcement@atlassian.com).

End of Life Announcement for Database Support

Database	Support End Date
PostgreSQL 8.3	When Confluence 5.5 is released, due in early 2014

PostgreSQL 8.3 notes:

- Confluence 5.4 is the last version that will support PostgreSQL 8.3.
- Confluence 5.4 and previously-released versions will continue to work with PostgreSQL 8.3. However, we will not fix bugs affecting PostgreSQL 8.3 after the [end-of-life date](#) for your version of Confluence.
- Confluence 5.5 will not be tested with PostgreSQL 8.3.

Deprecated Web Browsers for Confluence (24 September 2013)

To allow us to dedicate resources to providing the best experience on modern browsers, Confluence 5.5 will be the **last release that supports Internet Explorer 8 (IE8)**. The reasons behind this decision are to enable us to provide the best user experience to our customers, accelerate our pace of innovation and give us the ability to utilise modern browser technologies.

End of support means that Atlassian will not perform any maintenance on Confluence related to IE8 after the final release of Confluence 5.5.x, except for security related issues. In order to minimise the impact on you and the way your company uses Confluence, we have provided this announcement as early as possible, and hope that the subsequent 6 month period will give you adequate time to prepare for this change without disruption.

Atlassian will continue to support Internet Explorer 9 (IE9) and Internet Explorer 10 (IE10) as well as the latest versions of Chrome, Firefox and Safari. For further information, please refer to the [Supported Platforms](#) page. If you have questions or concerns regarding this announcement, please email [eol-announcement at atlassian dot com](mailto:eol-announcement@atlassian.com).

Deprecated Databases for Confluence (13 August 2013)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of [supported platforms](#) for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email eol-announcement@atlassian.com.

End of Life Announcement for Database Support

Database	Support End Date
MS SQL 2005	When Confluence 5.3 is released, due in late 2013

MS SQL 2005 notes:

- Confluence 5.2 is the last version that will support MS SQL 2005.
- Confluence 5.2 and previously-released versions will continue to work with MS SQL 2005. However, we will not fix bugs affecting MS SQL 2005 after the [end-of-life date](#) for your version of Confluence.
- Confluence 5.3 will not be tested with MS SQL 2005.

Deprecated Tomcat platform for Confluence (29 August 2012)

This section announces the end of Atlassian support for Tomcat 5.5.x for Confluence. Please note: Apache have announced that support for Apache Tomcat 5.5.x will end on 30 September 2012: [End of life for Apache Tomcat 5.5.x](#).

End of support means that Atlassian will not fix bugs related to the specified version of Tomcat, past the support end date for your version of Confluence. The details are below. Please refer to the list of [supported platforms](#) for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email eol-announcement@atlassian.com.

End of Life Announcement for Tomcat 5.5.x Support

Platform	Support End Date
Tomcat 5.5.x	When Confluence 5.0 is released, due in early 2013

Tomcat 5.5.x notes:

- Confluence 4.3 is the last major version that will support Tomcat 5.5.x. The Confluence 4.3.x bug-fix releases will also continue to support Tomcat 5.5.x.
- Tomcat 6.0.x will still be supported in Confluence 5.0.
- Confluence 4.3.x and previously-released versions will continue to work with Tomcat 5.5.x. However, we will not fix bugs affecting Tomcat 5.5.x after the [end-of-life date](#) for your version of Confluence.
- Confluence 5.0 will not be tested with Tomcat 5.5.x.

Deprecated Java platform for Confluence (6 August 2012)

This section announces the end of Atlassian support for Java 6 for Confluence. Please note that Oracle has announced the end of public updates for Java 6: [Java SE 6 End of Public Updates Notice](#).

End of support means that Atlassian will not fix bugs related to the specified version of Java, past the support end date for your version of Confluence. The details are below. Please refer to the list of [supported platforms](#) for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email eol-announcement@atlassian.com.

End of Life Announcement for Java 6 Support

Platform	Support End Date
Java 6 (JRE and JDK 1.6)	When Confluence 5.0 is released, due in early 2013

Java 6 notes:

- Confluence 4.3 is the last major version that will support Java 6. The Confluence 4.3.x bug-fix releases will also continue to support Java 6.
- Java 7 (JRE and JDK 1.7) will still be supported in Confluence 5.0.
- Confluence 4.3.x and previously-released versions will continue to work with Java 6. However, we will not fix bugs affecting Java 6 after the [end-of-life date](#) for your version of Confluence.
- Confluence 5.0 will not be tested with Java 6.

Deprecated Databases for Confluence (1 May 2012)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of [supported platforms](#) for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Database Support

Database	Support End Date
PostgreSQL 8.2	When Confluence 4.3 is released, due in mid 2012

PostgreSQL 8.2 notes:

- Confluence 4.2 is the last version that will support version 8.2 of PostgreSQL.
- Versions 8.3, 8.4 and 9.0 will still be supported in Confluence 4.3.
- Confluence 4.2 and previously-released versions will continue to work with PostgreSQL 8.2. However, we will not fix bugs affecting PostgreSQL 8.2 after the [end-of-life date](#) for your version of Confluence.
- Confluence 4.3 will not be tested with PostgreSQL 8.2.

Deprecated Databases for Confluence (13 March 2012)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of [supported platforms](#) for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Database Support

Database	Support End Date
DB2	When Confluence 4.3 is released, due in mid 2012

DB2 notes:

- Confluence 4.2 is the last version that will support DB2.
- From Confluence 4.3, no versions of DB2 will be supported.
- Confluence 4.2 and previously-released versions will continue to work with DB2. However, we will not fix bugs affecting DB2 after the [end-of-life date](#) for your version of Confluence.
- Confluence 4.3 will not be tested with DB2.
- For help with moving from DB2 to a supported database, please refer to the list of [supported databases](#) and the guide to [migrating to another database](#).

Deprecated Operating Systems for Confluence (21 July 2011)

This section announces the end of Atlassian support for certain operating systems for Confluence. End of support means that Atlassian will not fix bugs related to running Confluence server on that operating system past the support end date.

We will stop supporting the following operating systems from Confluence 4.0, due in late 2011:

- Mac OS X (as a Confluence server platform).

The details are below. Please refer to the list of [supported platforms](#) for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Operating System Support

Operating System	Support End Date
Mac OS X (as a Confluence server platform)	When Confluence 4.0 releases, due in late 2011

- **Mac OS X Notes:**
 - Atlassian intends to end support for Mac OS X (as a server platform) in Confluence 4.0 (due for release in late 2011). Confluence 3.5 is the last version that will support Mac OS X.
 - The Sun/Oracle JDK/JRE 1.6 is the only JDK platform officially supported by Atlassian. This means that Apple Mac OS X is not a supported operating system for the Confluence server, as the Sun/Oracle JDK does not run on Mac OS X.
 - Accessing Confluence as a user from Mac OS X via a compatible web browser will still be supported for the foreseeable future.

Deprecated Databases for Confluence (7 January 2011)

This section announces the end of Atlassian support for certain database versions for Confluence. End of support means that Atlassian will not fix bugs related to certain database versions past the support end date.

We will **stop supporting the following database versions** from Confluence 4.0, due in late 2011:

- MySQL 5.0.

The details are below. Please refer to the list of [supported platforms](#) for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Database Support

Database	Support End Date
MySQL (version 5.0 only)	When Confluence 4.0 releases, due in late 2011

- **MySQL Notes:**
 - Atlassian intends to end support for MySQL 5.0 in Confluence 4.0 (due for release in the middle of 2011). Confluence 3.5 is the last version that will support MySQL 5.0.
 - MySQL 5.1 will still be supported.
 - 'Support End Date' means that Confluence 3.5 and previously released versions will continue to work with MySQL 5.0. However, we will not fix bugs affecting MySQL 5.0 past the support end date.
 - Confluence 4.0 will not be tested with MySQL 5.0.

Deprecated Web Browsers for Confluence (7 January 2011)

This section announces the end of Atlassian support for certain web browser versions for Confluence. End of support means that Atlassian will not fix bugs related to certain web browser versions past the support end

date.

We will **stop supporting the following web browser versions** from Confluence 4.0, late middle of 2011:

- Microsoft Internet Explorer 7 (IE7).
- Safari 4.
- Firefox 3.5.

The details are below. Please refer to the list of [supported platforms](#) for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Web Browser Support

Web Browser	Support End Date
Microsoft Internet Explorer (version 7 only)	When Confluence 4.0 releases, late the middle of 2011
Safari (version 4 only)	When Confluence 4.0 releases, due in late of 2011
Firefox (version 3.5 only)	When Confluence 4.0 releases, due in late of 2011

- **Internet Explorer Notes:**
 - Atlassian intends to end support for IE7 in Confluence 4.0 (due for release in the middle of 2011). Confluence 3.5 is the last version that will support IE7.
 - IE8 will still be supported.
 - 'Support End Date' means that Confluence 3.5 and previously released versions will continue to work with IE7. However, we will not fix bugs affecting IE7 past the support end date.
 - Confluence 4.0 will not be tested with IE7.
- **Safari Notes:**
 - Atlassian will introduce support for Safari 5 in Confluence 3.5.
 - We intend to end support for Safari 4 in Confluence 4.0 (due for release in the middle of 2011). Confluence 3.5 is the last version that will support Safari 4.
 - 'Support End Date' means that Confluence 3.5 and previously released versions will continue to work with Safari 4. However, we will not fix bugs affecting Safari 4 past the support end date.
 - Confluence 4.0 will not be tested with Safari 4.
- **Firefox Notes:**
 - Atlassian will end support for Firefox 3.0 in Confluence 3.5, as [previously announced](#).
 - We intend to end support for Firefox 3.5 in Confluence 4.0 (due for release in the middle of 2011). Confluence 3.5 is the last version that will support Firefox 3.5.
 - Firefox 3.6 will still be supported.
 - 'Support End Date' means that Confluence 3.5 and previously released versions will continue to work with Firefox 3.5. However, we will not fix bugs affecting Firefox 3.5 past the support end date.
 - Confluence 4.0 will not be tested with Firefox 3.5.

Deprecated Databases for Confluence (12 October 2010)

This section announces the end of Atlassian support for certain database versions for Confluence. End of support means that Atlassian will not fix bugs related to certain database versions past the support end date.

We will **stop supporting the following database versions**:

- From Confluence 3.5, due in the first half of 2011, Confluence will no longer support PostgreSQL 8.1. *Note, PostgreSQL 8.2 and PostgreSQL 8.4 will still be supported.*

The details are below. Please refer to the [Supported Platforms](#) for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Database Support

Database	Support End Date
PostgreSQL (version 8.1 only)	When Confluence 3.5 releases, due in the first half of 2011

- **PostgreSQL (version 8.1 only) End of Support Notes:**

- Atlassian intends to end support for PostgreSQL 8.1 in Confluence 3.5 (due to release in the first half of 2011), with the final support for these platforms in Confluence 3.4. PostgreSQL 8.2 and PostgreSQL 8.4 will still be supported.
- 'Support End Date' means that Confluence 3.4 and previous released versions will continue to work with the PostgreSQL 8.1. However, we will not fix bugs affecting PostgreSQL 8.1 past the support end date.
- Confluence 3.5 (due to release in the first half of 2011) will not be tested with PostgreSQL 8.1.

Deprecated Web Browsers for Confluence (12 October 2010)

This section announces the end of Atlassian support for certain web browser versions for Confluence. End of support means that Atlassian will not fix bugs related to certain web browser versions past the support end date.

We will **stop supporting the following web browser versions:**

- From Confluence 3.5, due in the first half of 2011, Confluence will no longer support Firefox 3.0.
Note, Firefox 3.5 and Firefox 3.6 will still be supported.

The details are below. Please refer to the [Supported Platforms](#) for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Web Browser Support

Web Browser	Support End Date
Firefox (version 3.0 only)	When Confluence 3.5 releases, due in the first half of 2011

- **Firefox (version 3.0 only) End of Support Notes:**

- Atlassian intends to end support for Firefox 3.0 in Confluence 3.5 (due to release in the first half of 2011), with the final support for these platforms in Confluence 3.4. Firefox 3.5 and Firefox 3.6 will still be supported.
- 'Support End Date' means that Confluence 3.4 and previous released versions will continue to work with Firefox 3.0. However, we will not fix bugs affecting Firefox 3.0 past the support end date.
- Confluence 3.5 (due to release in the first half of 2011) will not be tested with Firefox 3.0.

Deprecated Databases for Confluence (6 July 2010)

This section announces the end of Atlassian support for certain database versions for Confluence. End of support means that Atlassian will not fix bugs related to certain database versions past the support end date.

We will **stop supporting the following database versions:**

- From Confluence 3.4, due in the second half of 2010, Confluence will no longer support Oracle 10g (i.e. Oracle 10.1 and Oracle 10.2).
Note, Oracle 11g (i.e. Oracle 11.1 and Oracle 11.2) will still be supported.

We have made these decisions in line with Oracle's decision to stop support for Oracle 10g, as per the "Oracle Database (RDBMS) Releases Support Status Summary [ID 161818.1]" article on the [Oracle Support site](#) (note, you will need an Oracle Support account to find and view the article). This also will reduce the testing time required for each release and help us speed up our ability to deliver market-driven features. We are committed to helping our customers understand this decision and assist them in upgrading to Oracle 11g if needed.

The details are below. Please refer to the [Supported Platforms](#) for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Database Support

Database	Support End Date
Oracle (version 10.1 and 10.2 only)	When Confluence 3.4 releases, due in the second half of 2010

- **Oracle (version 10.1 and 10.2 only) End of Support Notes:**
 - Atlassian intends to end support for Oracle 10.1 and Oracle 10.2 in Confluence 3.4 (due to release in the second half of 2010), with the final support for these platforms in Confluence 3.3. Oracle 11.1 and Oracle 11.2 will still be supported.
 - 'Support End Date' means that Confluence 3.3 and previous released versions will continue to work with the Oracle 10.1 and Oracle 10.2. However, we will not fix bugs affecting Oracle 10.1 or Oracle 10.2 past the support end date.
 - Confluence 3.4 (due to release in the second half of 2010) will not be tested with Oracle 10.1 and Oracle 10.2.

Deprecated Web Browsers for Confluence (6 July 2010)

This section announces the end of Atlassian support for certain web browser versions for Confluence. End of support means that Atlassian will not fix bugs related to certain web browser versions past the support end date.

We will **stop supporting the following web browser versions**:

- From Confluence 3.4, due in the second half of 2010, Confluence will no longer support Safari 3 or Safari 3.1.
Note, Safari 4 will still be supported.

The details are below. Please refer to the [Supported Platforms](#) for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Web Browser Support

Web Browser	Support End Date
Safari (version 3 and 3.1 only)	When Confluence 3.4 releases, due in the second half of 2010

- **Safari (version 3 and 3.1 only) End of Support Notes:**
 - Atlassian intends to end support for Safari 3 and Safari 3.1 in Confluence 3.4 (due to release in the second half of 2010), with the final support for these platforms in Confluence 3.3. Safari 4 will still be supported.
 - 'Support End Date' means that Confluence 3.3 and previous released versions will continue to work with the Safari 3 and Safari 3.1. However, we will not fix bugs affecting Safari 3 and Safari 3.1 past the support end date.
 - Confluence 3.4 (due to release in the second half of 2010) will not be tested with Safari 3 and Safari 3.1.

Deprecated Databases for Confluence (24 March 2010)

This section announces the end of Atlassian support for certain database versions for Confluence. End of support means that Atlassian will not fix bugs related to certain database versions past the support end date.

We will **stop supporting the following database versions**:

- From Confluence 3.3, due in Q3 2010, Confluence will no longer support DB2 8.2.

Note, DB2 9.7 will still be supported.

We are reducing our database support to reduce the amount of testing time and help us speed up our ability to deliver market-driven features. We are committed to helping our customers understand this decision and assist them in upgrading to DB2 9.7 if needed.

The details are below. Please refer to the [Supported Platforms](#) for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Database Support

Database	Support End Date
DB2 (version 8.2 only)	When Confluence 3.3 releases, due Q3 2010

- **DB2 (version 8.2 only) End of Support Notes:**
 - Atlassian intends to end support for DB2 8.2 in Q3 2010, with the final support for these platforms in Confluence 3.2. DB2 9.7 will still be supported.
 - 'Support End Date' means that Confluence 3.2 and previous released versions will continue to work with the DB2 8.2. However, we will not fix bugs affecting DB2 8.2 past the support end date.
 - Confluence 3.3 (due to release in Q3 2010) will not be tested with DB2 8.2.

Deprecated Application Servers for Confluence (27 January 2010)

This section announces the end of Atlassian support for certain application servers for Confluence. End of support means that Atlassian will not fix bugs related to certain application servers past the support end date.

We will **stop supporting the following application servers**:

- From Confluence 3.2, due late Q1 2010, Confluence will no longer support JBoss application servers.
- From Confluence 3.3, due in Q3 2010, Confluence will no longer support Oracle WebLogic, IBM WebSphere or Caucho Resin.

We are reducing our application server platform support to reduce the amount of testing time and help us speed up our ability to deliver market-driven features. We are committed to helping our customers understand this decision and assist them in migrating to Tomcat, our supported application server.

The details are below. Please refer to the [Supported Platforms](#) for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Application Server Support

Application Servers	Support End Date
JBoss 4.2.2	When Confluence 3.2 releases, due late Q1 2010
Oracle WebLogic 9.2	When Confluence 3.3 releases, due Q3 2010
IBM WebSphere 6.1	When Confluence 3.3 releases, due Q3 2010
Caucho Resin 3.0, 3.1.6, 3.1.7	When Confluence 3.3 releases, due Q3 2010

- **JBoss End of Support Notes:**
 - 'Support End Date' means that Confluence 3.1 and previous released versions will continue to work with stated application servers. However, we will not fix bugs affecting JBoss application servers.
 - Confluence 3.2 will not support JBoss application servers.
- **WebLogic, WebSphere and Resin End of Support Notes:**
 - Atlassian intends to end support for Oracle WebLogic, IBM WebSphere, and Caucho Resin in Q3 2010, with the final support for these platforms in Confluence 3.2.
 - 'Support End Date' means that Confluence 3.2 and previous released versions will continue to work with the stated application servers. However, we will not fix bugs affecting Oracle WebLogic,

- IBM WebSphere, and Caucho Resin application servers past the support end date.
- Confluence 3.3 (due to release in Q3 2010) will only be tested with and support Tomcat 5.5.20+ and 6.0.
- If you have concerns with this end of support announcement, please email eo1-announcement@atlassian.com.

Why is Atlassian doing this?

We have chosen to standardise on Tomcat, because it is the most widely used application server in our user population. It is fast, robust, secure, well-documented, easy to operate, open source, and has a huge community driving improvements. It is the de facto industry standard, with several companies available that specialise in providing enterprise grade support contracts for it, ranging from customisations to 24/7 support.

Deprecated Java Platforms for Confluence (27 January 2010)

This section announces the end of Atlassian support for certain Java Platforms for Confluence.

We will **stop supporting the following Java Platforms**:

- From Confluence 3.3, due Q3 2010, support for Java Platform 5 (JDK/JRE 1.5) will end.

We are ending support for Java Platform 5, in line with the [Java SE Support Roadmap](#) (i.e. "End of Service Life" for Java Platform 5 dated October 30, 2009). We are committed to helping our customers understand this decision and assist them in updating to Java Platform 6, our supported Java Platform.

The details are below. Please refer to the [Supported Platforms](#) for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Java Platform Support

Java Platform	Support End Date
Java Platform 5 (JDK/JRE 1.5)	When Confluence 3.3 releases, due Q3 2010

- Java Platform 5 End of Support Notes:**
 - Atlassian intends to end support for Java Platform 5 in Q3 2010.
 - 'Support End Date' means that Confluence 3.2.x and previous released versions will continue to work with Java Platform 5 (JDK/JRE 1.5), however we will not fix bugs related to Java Platform 5 past the support end date.
 - Confluence 3.3 will only be tested with and support Java Platform 6 (JDK/JRE 1.6).
 - If you have concerns with this end of support announcement, please email eo1-announcement@atlassian.com.

Deprecated Web Browsers for Confluence (14 December 2009)

This section announces the end of Atlassian support for certain web browsers for Confluence.

We will **stop supporting older versions of web browsers** as follows:

- From Confluence 3.2, due late Q1 2010, support for Firefox 2 and Safari 2 will end.
- From 13 July 2010, in line with Microsoft's Support Lifecycle policy, support for IE6 will end.

The details are below. Please refer to the [Supported Platforms](#) for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email eo1-announcement@atlassian.com.

End of Life Announcement for Web Browser Support

Web Browsers	Support End Date
Firefox 2	When Confluence 3.2 releases, late Q1 2010

Safari 2	When Confluence 3.2 releases, late Q1 2010
Internet Explorer 6	When Confluence 3.3 releases (target Q3 2010) or 13 July 2010, whichever is sooner

- **Firefox 2 and Safari 2 Notes:**

- Confluence 3.1 is the last version to officially support Firefox 2 and Safari 2.
- You may be able to use these older browser for the most common use cases like viewing and editing content, but official support for these browsers will end once you upgrade to Confluence 3.2.
- Confluence 3.2 is currently targeted to release late Q1 2010 and will not be tested with Firefox 2 and Safari 2. After the Confluence 3.2 release, Atlassian will not provide fixes in older versions of Confluence for bugs affecting Firefox 2 and Safari 2.

- **Internet Explorer 6 Notes:**

- Confluence 3.2 (due late Q1 2010) will be the last version to officially support Internet Explorer 6.
- Confluence 3.3 is currently targeted to release Q3 2010 and will **not** support IE6.
- Atlassian will support IE6 in Confluence until the 13th of July 2010, in line with Microsoft's [Support Lifecycle policy](#). Beyond that date, released versions of Confluence will continue working with IE6 just as they did before, but we will not fix bugs affecting Internet Explorer 6.
- You may be able to use Internet Explorer 6 for the most common use cases like viewing and editing content, but official support for this browser will end once you upgrade to Confluence 3.3.

Supported Platforms FAQ

Q: How does Atlassian choose which JRE versions, application servers and databases to support?

For application servers and databases, we try to pick a good cross-section of open source options and popular commercial platforms. We then choose which JRE versions to support based on the recommended environments for these servers.

Q: What is a supported platform?

A supported platform is one that:

- Confluence is regularly tested on during the development cycle
- One that is available within Atlassian for support technicians and developers to reproduce problems
- Bugs raised against it will be given a high priority

Supporting a platform means we know how to get Confluence running in that environment and can troubleshoot Confluence issues within it. It does not mean we have any particular expertise beyond that. As such, we may not be able to provide assistance with customising or tuning that application server or database. (Atlassian support is not a substitute for a good database administrator.)

Q: Can I get assistance with running Confluence on a platform that is *not* supported?

If you are running Confluence on an unsupported platform, then we can not guarantee providing any support for it. Furthermore, we will recommend that you switch to a platform which is supported.

Q: If you write your application to standards like J2EE, JDBC and SQL, doesn't that mean it should run on any compliant server?

Confluence is a complicated application and we commonly encounter interesting edge-cases where different servers have interpreted the specifications differently. Then again, each server has its own different collection of bugs.

Q: How can I get Atlassian to support Confluence on a new platform?

Supporting a new platform involves a significant investment of time by Atlassian, both up-front costs to set up new testing environments and fix any issues we might encounter and the ongoing costs involved in maintaining the application against this new environment in the future. As such, supporting a new platform is not something we will do unless we know there is significant demand for it.

Please be aware that your interest alone will not be enough for us to add support for your application server or database. We would need to see a significant number of votes on the issue raised in our public JIRA site or a significant level of interest in our forums, before considering supporting that platform.

Q: My organisation has standardised on an operating environment that Confluence does not support.

What can I do?

In this situation, you have the following two options:

1. Run Confluence in the unsupported environment, with the caveats mentioned above.
2. Make an exception to your standardised operating environment and set up Confluence based on its supported platforms.

Migrating Confluence Between Servers

This page describes how to move Confluence between physical servers. It is distinct from other functions. It does not cover database migration, application server migration, or upgrading. Atlassian suggests doing each of these steps separately. See also:

- [Upgrading Confluence](#)
- [Migrating to Another Database](#)
- [Switching to Apache Tomcat](#)

How to Create a Test or Development Site

Administrators may need to move a Confluence site from one server to another for upgrades or downtime. This page tells you how to copy a Confluence site from one server to another. For example, you may want to transfer your current production snapshot to a test server as [permitted in the licence agreement](#).



- **Avoid upgrades while migrating.** If you are planning to switch databases, application servers or Confluence versions, firstly perform the application transfer in isolation, and test that it was successful before making other changes.
- Development licenses are available for any Commercial or Academic license. [Create one](#) or [contact Atlassian](#) for help.

On this page:

- [How to Create a Test or Development Site](#)
- [Transferring Confluence To Another Server Using The Same Operating System](#)
- [Transferring Confluence To Another Server Using a Different Operating System](#)
- [Ensuring no contact with production systems](#)
- [Migrating from HTTPS to HTTP](#)
- [Notes](#)

Transferring Confluence To Another Server Using The Same Operating System

If the operating systems on both servers are the same, then the home and install folders can be copied straight into an identical external database and user management setup.

1. On the original server, create zips of the Confluence install and home directories. Copy the zips to the new server.
2. On the new server, unzip the install and home directories. Windows users should avoid unzipping with the Windows built-in extractor, instead use Winzip or the free [7Zip](#).



If you are changing the location of the home directory, open the Confluence `install\confluence\WEB-INF\classes` directory and edit `confluence-init.properties` by changing the line starting with `'confluence.home='`.

3. If you are using the EAR/WAR distribution, modify the location of your war file if need be. If using Tomcat, this is likely in `/Conf/Catalina/localhost`. You'll want to make sure the `docbase` attribute is pointing to the right location.
4. This next step is dependent on your database:
 - For users of the *internal database*, the database content is stored inside the home directory. You should [switch to an external database](#) after the transfer is successful. The internal database is for evaluation only and is not recommended for use in Production systems.
 - For *external databases stored on another server*: change the user account or datasource

permissions so that the new server has the same network access permissions as the original. Then confirm from the new server that the hostname can be resolved and is listening for database connections on the expected port.

- For *external databases hosted locally (ie. localhost)*: on the original server, create a manual database backup using a native db dump backup tool. Copy the database backup to the new server.
5. On the new server, install or upgrade the database version to match the original server.
 6. Import the database backup.
 7. Add a database user account with the same username and password as the original.
 8. Provide the database user with the full access to the imported database.
 9. Use a database administration tool to confirm that the user can login from the localhost.
 10. This step depends on your database connection:
 - a. If you use JDBC (the default option) to connect to the database, to modify any database connection information, go to the Confluence home directory and edit `confluence.cfg.xml`. The connection URL is set under `hibernate.connection.url`. **Ensure it does not point to your production database server.**
 - b. If you use a data source, follow the instructions for your database type and ensure the data source points to the new database: [PostgreSQL](#), [MySQL](#), [SQL Server](#) or [Oracle](#).
 11. If you are using internal user management, skip this step. For users who have JIRA or LDAP integration, provide the new server with network or local access to the same hosts as the original. If this is a true test site, set up a test of your JIRA site or LDAP server so as not to disrupt production systems and change the `server.xml` or `atlassian-user.xml` files (Confluence 3.4 and below), or modify the directory settings in Confluence Admin > User Directories (Confluence 3.5 and above) to point to the appropriate test servers. Note that it might be acceptable to use a production connection here, as users won't be logging on to the test system in high volume.
 12. If appropriate, make sure [no emails are sent out from the test system](#).
 13. If you have previously installed Confluence using the guided installer *and* plan on starting Confluence using the `startup` or `start-confluence` scripts in the `ConfluenceInstall/bin/` directory, check `setenv.sh` (Unix/Linux) or `setenv.bat` (Windows) in the same directory. If there is a `JRE_HOME` set, ensure that the path to the JRE is up to date in regards to the new environment.
 14. Start Confluence.
 15. Go to Administration > License Details and add your development license key. You can generate one at <http://my.atlassian.com>. There are more details in [How to get a Confluence developer license](#).
 16. If you configured Confluence as a Windows service, [repeat those instructions](#).
 17. Add your [development license key](#).
 18. Some customers have experienced problems with Confluence's search functions after performing a migration, or that the content of their `{recently-updated}` macro is not being updated correctly. Errors in the `atlassian-confluence.log` file corroborate such problems. Hence, to avoid these issues, it is strongly recommended that you perform a [rebuild of your content indices](#) after performing a migration.

Transferring Confluence To Another Server Using a Different Operating System

Migrating from Windows to Linux

You will need to replace the backslash with forward slash in the following lines in `confluence.cfg.xml`:

```
<property
name="attachments.dir">${confluenceHome}/attachments</property>
  <property name="lucene.index.dir">${confluenceHome}/index</property>
  <property
name="webwork.multipart.saveDir">${confluenceHome}/temp</property>
```

Using database tools (preferred option)

If you are using the [Production backup strategy](#), follow these steps:

1. Download the proper distribution (**the same one you have from your original site**) from the [Download Archive](#).
2. Copy your Confluence *home* (not *install*) directory from your original server (even if it was a different OS).
3. If you are changing the location of the home directory, open the Confluence

install\confluence\WEB-INF\classes directory and edit confluence-init.properties by changing the line starting with 'confluence.home='.

4. For external databases stored locally, on the original server, create a manual database backup using a native db dump backup tool.
5. Copy the database backup to the new server.
6. On the new server, install or upgrade the database version to match the original server.
7. Import the database backup.
8. Add a database user account with the same username and password as the original.
9. Provide the user with the full access to the imported database.
10. Use a database administration tool to confirm that the user can login from the localhost.
11. To modify any database connection information, go to the Confluence home directory and edit confluence.cfg.xml. The connection URL is set under hibernate.connection.url. **Ensure it does not point to your production database server.**
12. If you are using internal user management, skip this step. For users who have JIRA or LDAP integration, provide the new server with network or local access to the same hosts as the original.
13. Copy server.xml, atlassian-user.xml, osuser.xml, any patches, and any other customized files velocity or properties files. If this is a true test site, set up a test of your JIRA site or LDAP server so as not to disrupt production systems and change the server.xml or atlassian-user.xml files to point to the appropriate test servers. Note that it might be acceptable to use a production connection here, as users won't be logging on to the test system in high volume.
14. If appropriate, make sure [no emails are sent out from the test system](#).
15. If you have previously installed Confluence using the guided installer *and* plan on starting Confluence using the startup or start-confluence scripts in the Confluenceinstall/bin/ directory, check setenv.sh (Unix/Linux) or setenv.bat (Windows) in the same directory. If there is a JRE_HOME set, ensure that the path to the JRE is up to date in regards to the new environment.
16. Start Confluence.
17. Go to **Administration > License Details** and add your development license key. You can generate one at <http://my.atlassian.com>. There are more details in [How to get a Confluence developer license](#).
18. If you configured Confluence as a Windows service, [repeat those instructions](#).
19. Add your [development license key](#).
20. Some customers have experienced problems with Confluence's search functions after performing a migration, or that the content of their {recently-updated} macro is not being updated correctly. Errors in the atlassian-confluence.log file corroborate such problems. Hence, to avoid these issues, it is strongly recommended that you perform a [rebuild of your content indices](#) after performing a migration.

Using XML data backups (only for small to medium sized installations)

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](#).

If you're not yet using the [Production backup strategy](#), you can migrate Confluence to a different server machine by creating an XML data backup as usual, and then importing that to Confluence on the new server.

1. Create an XML data backup from Confluence as follows:
 - a. Choose the **cog icon** , then choose **General Configuration** under Confluence Administration.
 - b. Select **Backup & Restore**.
 - c. Check the **Backup Attachments** option and click **Backup**.
2. Identify the version of Confluence that you are currently using. This is displayed at the bottom of each Confluence page.
3. Download Confluence to the new server. Get the version of Confluence that you identified above, but for the operating system of the new server. You may be using either the [latest Confluence version](#), or an [older version](#).
4. Install Confluence on the new server.
5. Go to **Administration > License Details** and add your [development license key](#). You can generate a license at <http://my.atlassian.com>. You can find more details in [How to get a Confluence developer license](#).
6. Restore your XML data backup from **Administration > Backup and Restore**.
7. If appropriate, make sure that [no email contact](#) can be made with the test system.
8. Some customers have experienced problems with Confluence's search functions after performing a

migration, or that the content of their {recently-updated} macro is not being updated correctly. Errors in the atlassian-confluence.log file corroborate such problems. Hence, to avoid these issues, it is strongly recommended that you [rebuild your content indices](#) after performing a migration.

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 'SELECT * FROM' to 'DELETE FROM' in the above queries once you are sure you want to remove the specified accounts.

Once this is done, you can start your test site without any mails being sent or retrieved. Think carefully about other plugins which may access production systems (SQL macro, JIRA macro, etc.). If these write content, or create unwanted load on external systems, they should be disabled promptly after starting the test site.

Migrating from HTTPS to HTTP

You may want to migrate from a server secured by SSL to one which is not secured by SSL. For example, this may be useful if you are copying a Confluence site from a production to a test site.

To migrate from HTTPS to HTTP, undo the HTTPS-specific settings that are described on this page: [Adding SSL for Secure Logins and Page Security](#).

Notes

- Ricky Sheaves ([calebscreek](#)) has written an interesting blog post on [Moving Confluence from Windows to \(Ubuntu\) Linux](#).
- If you wish to **merge** two Confluence sites, you can consider using the [remote import plugin](#). This plugin is currently not supported. The supported method would be to [export a space](#) and then [import each space](#) one by one. The two Confluence sites must be running the same version of Confluence.

From Confluence Evaluation through to Production Installation

So, you want to try Confluence on an evaluation installation, then move to a production installation when you are ready? This page gives an overview of the steps to follow.

Assumptions:

- This page starts with telling you how to install an evaluation Confluence site. If you have already finished evaluating Confluence, you can safely skip steps 1 to 3.
- Your production installation will be an installed version of Confluence, **not** a Confluence OnDemand site.
- You will evaluate Confluence on an installed version too, not a Confluence OnDemand site.

If you are using Confluence OnDemand to evaluate Confluence, please refer to the following guide when you want to move to an installed version: [Migrating from Confluence OnDemand to a Confluence Installed Site](#).

On this page:

- [Step 1. Set up your evaluation Confluence site](#)
- [Step 2. Add users and content to your evaluation site](#)
- [Step 3. Look for interesting add-ons as part of your evaluation](#)
- [Step 4. Set up your production Confluence site](#)

Related pages:

- [Supported Platforms](#)
- [Embedded HSQLDB Database](#)
- [Adding and Inviting Users](#)
- [About Add-ons](#)
- [Getting Started as Confluence Administrator](#)
- [Confluence Installation and Upgrade Guide](#)

Step 1. Set up your evaluation Confluence site

If you have already set up an evaluation Confluence site, you can skip this step.

Below is a summary of the installation and setup procedure, focusing on the choice of database.

To install Confluence:

1. Download the automatic installer from the [Confluence download site](#).
Note: If you are using a Mac or another unsupported platform for your evaluation, you will need to install from a zip file. Details are in the full installation guide.
2. Run the installer and choose the express or custom installation. If you are not sure, choose **Express Install**.
 - The **express** option will install Confluence with default settings.
 - The **custom** option allows you to choose the Confluence installation directory, home (data) directory, ports and other options.
3. When prompted, choose the option to **open Confluence in your browser**, where you can complete the setup.

To set up Confluence, including the database:

1. Follow the prompts in the browser-based setup wizard, to get your Confluence license.
2. Choose the **evaluation** or **production** installation type. If you are not sure, choose **Evaluation Installation**.
 - The **evaluation** option will install Confluence with default settings, including the embedded database.
 - If you decide to do a **production** installation for increased flexibility during evaluation, the setup wizard will prompt you for various options. When choosing a database, we recommend the **embedded database** for evaluation purposes, because it is simpler and faster to set up.
 - When you move to a production installation, you will be able to move your Confluence data to a production-ready database, as described below.

Step 2. Add users and content to your evaluation site

If you have finished evaluating Confluence, you can skip this step.

Depending on your choices during the Confluence setup, your evaluation site may include sample content. The example pages, blog posts and attachments are in the 'Demonstration space'. This space is present if:

- You chose the 'Evaluation Installation' during setup.
- Or you chose the 'Production Installation', then chose to include the 'Example Site'.

You can update the sample content, and create more of your own. You can also invite people to join you on the site.

When you move to a production site, you can choose to copy the content and users to the new site.

To create content in your evaluation site:

- Choose **Spaces > Create Space** to add a space, which is like a library of pages.
- Choose **Create** to add pages and blog posts.

To add users: Choose the **cog** icon  then choose **User Management**.

For more tips about getting started, see [Confluence 101](#).

Step 3. Look for interesting add-ons as part of your evaluation

If you have finished evaluating Confluence, you can skip this step.

Add-ons, also called plugins, provide additional features that you can install into your Confluence site. Some of them are provided free of charge. Many of the commercial add-ons are available free for an evaluation period.

You can browse and download add-ons on the [Atlassian Marketplace](#). You can also find add-ons via the Confluence user interface, which interacts with the Atlassian Marketplace for you.

To find useful add-ons via the Confluence user interface:

1. Choose the **cog** icon  then choose **Add-ons**.
2. Choose **Find new add-ons**.

Step 4. Set up your production Confluence site

When you are ready to move from an evaluation site to a production site, you need to migrate to a production-ready database. This involves installing a new Confluence site with a new database, and instructing Confluence to copy the data from your evaluation site to the new site. You will also need to check some important configuration settings, and define your backup strategy. The instructions below lead you through all the steps required.

Migrating your data to a production database:

1. Choose a database carefully, with a focus on reliability and backups. See our [list of supported databases](#). If you are unsure which one to choose, we recommend PostgreSQL.
2. Install a new database and a new Confluence site, by following our guide to [migrating to another database](#). The guide will lead you through the following steps:
 - Setting up your database server.
 - Adding a Confluence database (schema) to your database server.
 - Installing a new, production-ready Confluence site.
 - Copying your Confluence data from your evaluation site to your new production site.

Setting important configuration options on your production site:

- Set the base URL. See [Configuring the Server Base URL](#).
- Make sure you have configured an email server. See [Configuring a Server for Outgoing Mail](#).
- Decide on proxy setup and other settings that determine where Confluence fits into your network. See [Web Server Configuration](#).
- Consider setting up a secure connection via SSL. See [Running Confluence Over SSL or HTTPS](#).
- Read our guidelines on security. See [Best Practices for Configuring Confluence Security](#).
- Decide whether you will manage your users in Confluence or connect to an external LDAP directory. See [Configuring User Directories](#).
- Decide whether you want to allow public (anonymous) access to your site. See [Setting Up Public Access](#).
- Set up your permission scheme. See [Giving People Access to Content](#).
- Connect Confluence to JIRA and other applications. See [Linking to Another Application](#).

Defining your backup strategy:

By default, Confluence will create daily XML backups of your content and user data. This is suitable when you are evaluating Confluence. When you move to a production site, you need more robust backup procedures and technologies. See [Production Backup Strategy](#).

Migrating from Confluence OnDemand to a Confluence Installed Site

This page is for people who are currently using a [Confluence OnDemand](#) site, and wish to move to a Confluence site that is hosted on their own servers.





For migrating an Installed Site to Confluence OnDemand, see [Importing wiki data](#).

Summary

You will need to download and install a special **OnDemand release** of Confluence (for example, 'Confluence 5.0-OD-1') and then move your data from your hosted Confluence OnDemand site into your newly installed site. You cannot move your data from Confluence OnDemand to a site installed from the standard Confluence download.

Instructions

Note: You must do the data export and the Confluence download (both described in the steps below) **on the same day**. This will ensure that your data and your Confluence installation are of the same version and are therefore compatible.

To migrate from Confluence OnDemand to a Confluence installed site:

1. Export the data from your Confluence OnDemand site, using the Confluence backup manager.
 - For instructions, see this page in the Confluence OnDemand documentation: [Exporting wiki data](#).
 - You now have a backup file, also called an XML export, of your Confluence OnDemand data.
2. Download the OnDemand release of Confluence. Go to the [Confluence OnDemand download page](#) and get the latest 'OD' release for your operating system. The latest downloads are at the top of the list. For example, get the following files, replacing 'x' with the latest number available:
 - **For Windows 64-bit:** Get '5.x-OD-x - Windows Installer (64 bit)'
 - **For Windows 32-bit:** Get '5.x-OD-x - Windows Installer (32 bit)'
 - **For Linux 64-bit:** Get '5.x-OD-x - Linux Installer (64 bit)'
 - **For Linux 32-bit:** Get '5.x-OD-x - Linux Installer (32 bit)'
 - To install Confluence from an archive on UNIX or Mac OS X: Get '5.x-OD-x - Standalone (TAR.GZ Archive)'
 - To install Confluence from an archive on Windows: Get '5.x-OD-x - Standalone (ZIP Archive)'
 - EAR/WAR archives are also available.
3. Install Confluence as described in the [Confluence Installation Guide](#).
4. Import the data from your backup file (XML export) into your new Confluence installation. See [Restoring a Site](#).
5. Log in to your new Confluence site, using the following credentials:
 - Username: sysadmin
 - Password: sysadmin
6. Change the password immediately after logging in.

On this page:

- [Summary](#)
- [Instructions](#)
- [Background](#)
- [Support, limitations, and recommendations](#)
 - [Upgrading to an official release](#)
 - [Compatibility of third-party plugins](#)
 - [PostgreSQL database recommended](#)
 - [Confluence license](#)

Related pages:

- [Confluence Installation Guide](#)
- [Confluence Administrator's Guide](#)

Background

Backups taken from Confluence OnDemand are only compatible with the current OnDemand release (for example, 'Confluence 5.0-OD-1'). The reason is that Confluence OnDemand is typically ahead of the downloadable version of Confluence, meaning that you will have new features in Confluence OnDemand that are not yet available in the downloadable version.

It is therefore not possible to migrate your data to a Confluence site installed from the standard Confluence download. You will need to download and install the special OnDemand release of Confluence (for example,

'Confluence 5.0-OD-1') as described above.

The advantage is that you will be able to keep the Confluence OnDemand features currently not available to other customers who are using the standard downloadable version of Confluence. However, there are a few major limitations as noted below.

Support, limitations, and recommendations

Please note the following points about your Confluence site installed from an OnDemand release.

Upgrading to an official release

 You must upgrade to an official release of Confluence as soon as one is available. For instance, if you use 5.2 OD-2, you must upgrade to 5.2 when it is formally released for public consumption. OD releases are supported, but only until the official release is available.

Compatibility of third-party plugins

Because Confluence OnDemand is typically ahead of the downloadable version of Confluence, most third-party plugins will not be compatible with the OnDemand release. You may have some problems with third-party plugins on your Confluence site, until you are able to upgrade to the full release. Note, however, that any third-party plugins that you were using in Confluence OnDemand should be compatible with your newly installed site too.

If you have any questions about the compatibility of third-party plugins with your OnDemand release, please contact the plugin vendors. Contact details are on the [Atlassian Marketplace](#).

PostgreSQL database recommended

If you are uncertain about which database to choose for your Confluence site, we recommend PostgreSQL. See [Database Setup for PostgreSQL](#). The Confluence OnDemand site runs on PostgreSQL, and we therefore know it to be compatible with your OnDemand release.

If you decide to choose another supported database and discover any problems with compatibility, please contact [Atlassian Support](#). For a list of supported databases, see [Supported Platforms](#).

Confluence license

Your Atlassian OnDemand license cannot be used in a site installed from the downloadable version of Confluence. Please get your new Confluence license at <https://my.atlassian.com>.