A ATLASSIAN

Documentation for Confluence 7.16

Contents

Confluence administrator's guide	
Getting Started as Confluence Administrator	
Manage Users	 . 12
Add and Invite Users	 . 14
Delete or Disable Users	 . 19
Restore Passwords To Recover Admin User Rights	
Edit User Details	
Change a Username	
Managing Site-Wide Permissions and Groups	
Confluence Groups for Administrators	
Adding or Removing Users in Groups	
Global Permissions Overview	
Setting Up Public Access	
Revoke access for unlicensed users from Jira Service Management	
Configuring User Directories	
Configuring the Internal Directory	
Connecting to an LDAP Directory	
Configuring the LDAP Connection Pool	
Configuring an SSL Connection to Active Directory	 . 55
Connecting to an Internal Directory with LDAP Authentication	 . 66
Connecting to Crowd or Jira for User Management	
Reverting from Crowd or Jira applications to Internal User Management	
Managing Multiple Directories	
Managing Nested Groups	
Synchronizing 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	
Single sign-on for Confluence Data Center	
Managing System and Marketplace Apps	
Writing User Macros	
User Macro Template Syntax	
Customizing your Confluence Site	
Changing the Look and Feel of Confluence	 124
Customizing the Confluence Dashboard	 125
Changing the Site Logo	
Customizing Color Schemes	
Styling Confluence with CSS	
Basic Styling Tutorial	
Styling Fonts in Confluence	
Working with Themes	
Applying a Theme to a Site	
Creating a Theme	
Customizing Site and Space Layouts	
Upgrading Customized Site and Space Layouts	
Working With Decorator Macros	
Custom Decorator Templates	
Customizing a Specific Page	
Customizing the Login Page	
Modify Confluence Interface Text	
Customizing Email Templates	
Changing the Default Behavior and Content in Confluence	 149
Administering Site Templates	
Changing the Site Title	 151
Choosing a Default Language	
Configuring the Administrator Contact Page	
Configuring the Site Home Page	
- J. J	

Customizing Default Space Content	
Editing the Site Welcome Message	. 159
Integrating Confluence with Other Applications	. 161
Linking to Another Application	. 162
Configuring Workbox Notifications	
Integrating Jira and Confluence	
Registering External Gadgets	
Configuring the Office Connector	
Managing Webhooks	
Managing your Confluence License	120
Managing Confluence Data	
Database Configuration	
Database JDBC Drivers	
Database Setup for Oracle	
Database Setup for PostgreSQL	
Database Setup for SQL Server	. 196
Database Setup For MySQL	. 199
Embedded H2 Database	. 204
Migrating to Another Database	. 206
Configuring Database Character Encoding	
Configuring database query timeout	
Surviving Database Connection Closures	
Configuring a datasource connection	
Configuring Confluence Data Center to work with Amazon Aurora	
Site Backup and Restore	
Production Backup Strategy	
Configuring Backups	
User Submitted Backup & Restore Scripts	
Manually Backing Up the Site	
Restoring a Site	. 234
Restoring a Space	. 237
Restoring a Test Instance from Production	. 240
Restoring Data from other Backups	
Retrieving File Attachments from a Backup	
Troubleshooting failed XML site backups	
Troubleshooting XML backups that fail on restore	
Import a space from Confluence Cloud	
Attachment Storage Configuration	
Configuring Attachment Size	
Hierarchical File System Attachment Storage	
Confluence Data Model	
Finding Unused Spaces or Pages	
Data Import and Export	
Import a Text File	
Auditing in Confluence	. 271
Audit Log Events in Confluence	. 275
Audit Log Integrations in Confluence	. 282
Set retention rules to delete unwanted data	
Data pipeline	
Data pipeline export schema	
Configuring Confluence	
Viewing System Information	
Live Monitoring Using the JMX Interface	
Tracking Customizations Made to your Confluence Installation	
View Space Activity	
Viewing Site Statistics	
Viewing System Properties	
Configuring the Server Base URL	
Configuring the Confluence Search and Index	
Configuring Indexing Language	
Configuring Search	
Content Index Administration	
Enabling OpenSearch	
Rebuilding the Ancestor Table	
Trooping the fillester radio in the first transfer and transfer	

Setting Up Confluence to Index External Sites	
Setting Up an External Search Tool to Index Confluence	
Configuring Mail	
Configuring a Server for Outgoing Mail	. 334
Setting Up a Mail Session for the Confluence Distribution	. 336
Configuring the Recommended Updates Email Notification	. 337
The Mail Queue	. 338
Configuring Character Encoding	. 339
Troubleshooting Character Encodings	
"" Euro character not displaying properly	
MySQL 3.x Character Encoding Problems	
Other Settings	
Configuring a WebDAV client for Confluence	
Configuring HTTP Timeout Settings	
Configuring Number Formats	
Configuring Shortcut Links	
Configuring Time and Date Formats	
Enabling the Remote API	
Enabling Threaded Comments	
Installing a Language Pack	
Installing Patched Class Files	
Configuring System Properties	
Recognized System Properties	
Working with Confluence Logs	
Configuring Logging	
log4j Logging Levels	
Troubleshooting SQL Exceptions	
Configure access logs	
Scheduled Jobs	
Configuring the Allowlist	
Configuring the Time Interval at which Drafts are Saved	. 403
Configuring Confluence Security	. 404
Confluence Security Overview and Advisories	. 405
Proxy and HTTPS setup for Confluence	. 408
Configuring Web Proxy Support for Confluence	. 410
Connecting to LDAP or Jira applications or Other Services via SSL	. 412
Using Apache with mod_proxy	
Running Confluence behind NGINX with SSL	. 420
Running Confluence Over SSL or HTTPS	
Using Apache to limit access to the Confluence administration interface	. 430
Using Apache with mod_jk	
Using mod_rewrite to Modify Confluence URLs	. 433
Configuring Secure Administrator Sessions	
Confluence Cookies	
Using Fail2Ban to limit login attempts	
Securing Confluence with Apache	
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	
Configuring RSS Feeds	
Preventing and Cleaning Up Spam	
Configuring a Confluence Environment	
Confluence Home and other important directories	
Application Server Configuration	. 460
Managing Application Server Memory Settings	
Starting Confluence Automatically on System Startup	
Start Confluence Automatically on Linux	
Start Confluence Automatically on Windows as a Service	
Performance Tuning	. 470

Cache Performance Tuning	474
Cache Statistics	
Memory Usage and Requirements	479
Requesting Performance Support	481
Compressing an HTTP Response within Confluence	485
Garbage Collector Performance Issues	
Troubleshooting Slow Performance Using Page Request Profiling	487
Identifying Slow Performing Macros	
Confluence Diagnostics	
Faster permissions service	
Data Collection Policy	
Administering Collaborative Editing	
Possible Confluence and Synchrony Configurations	
Configuring Synchrony	
Set up a Synchrony cluster for Confluence Data Center	
Migrate from a standalone Synchrony cluster to managed Synchrony	
Troubleshooting Collaborative Editing	
Using read-only mode for site maintenance	
Administering the Atlassian Companion App	
Notifications from Atlassian	
Administer analytics	
Confluence installation and upgrade guide	
System Requirements	
Server Hardware Requirements Guide	538
Running Confluence in a Virtualized Environment	542
Confluence Installation Guide	543
Installing Confluence	544
Installing a Confluence trial	545
Installing Confluence on Windows	547
Installing Confluence on Windows from Zip File	
Uninstalling Confluence from Windows	
Installing Confluence on Linux	
Installing Confluence on Linux from Archive File	
Uninstalling Confluence from Linux	
Unattended installation	
Change listen port for Confluence	
Start and Stop Confluence	
Installing Confluence Data Center	
Upgrading Confluence Data Center	
Adding and Removing Data Center Nodes	
Change Node Discovery from Multicast to TCP/IP or AWS	
Running Confluence Data Center in AWS	
Getting started with Confluence Data Center on Azure	
Administering Confluence Data Center on Azure	
Running Confluence Data Center on a Kubernetes cluster	
Installing Java for Confluence	
Setting the JAVA_HOME Variable in Windows	
Change the Java vendor or version Confluence uses	
Creating a Dedicated User Account on the Operating System to Run Confluence	
Confluence Setup Guide	
Configuring Jira Integration in the Setup Wizard	
Upgrading Confluence	635
Upgrading Beyond Current Licensed Period	643
Confluence Post-Upgrade Checks	645
Migration from Wiki Markup to XHTML-Based Storage Format	646
Migration of Templates from Wiki Markup to XHTML-Based Storage Format	
Upgrading Confluence Manually	
Create a staging environment for upgrading Confluence	
Upgrade Confluence without downtime	
Upgrade a Confluence cluster manually without downtime	
Upgrade a Confluence cluster on AWS without downtime	
Upgrade a Confluence cluster through the API without downtime	
Roll back a rolling upgrade	
Upgrade task troubleshooting	
- pyrone men nemeroeming	200

Supported Platforms	 682
End of Support Announcements for Confluence	 685
Bundled Tomcat and Java versions	
Supported Platforms FAQ	 720
Migrate your Confluence site	
Migrate from Server to Data Center	
Migrate from Confluence Cloud to Server	 724
Migrating Confluence Between Servers	
Move to a non-clustered installation	
From Confluence Evaluation through to Production Installation	 732
Cloud Migration Assistant for Confluence	 735
Confluence Data Center documentation	
Getting Started with Confluence Data Center	 747
Confluence Server and Data Center feature comparison	 749
Clustering with Confluence Data Center	
External Process Pool for Confluence Data Center	 760
Document conversion for Confluence Data Center	 762
PDF export in Confluence Data Center	
Restricted Functions in Confluence Data Center	 766
Set up a Confluence Data Center cluster	
Confluence Data Center Performance	
Confluence Data Center disaster recovery	 780
Data Center Troubleshooting	 785
Troubleshooting a Data Center cluster outage	 789
Use a CDN with Atlassian Data Center applications	
Configure your CDN for Confluence Data Center	
Improving instance stability with rate limiting	 799
Adjusting your code for rate limiting	
Running Confluence Data Center on a single node	 809

Confluence administrator's guide

About the Confluence administrator's guide

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

For information on creating and administering spaces, SeeSpaces.

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

- Getting Started as Confluence Administrator
- Manage Users
 - Add and Invite Users
 - Delete or Disable Users
 - Restore Passwords To Recover Admin User Rights
 - Edit User Details
 - Change a Username
 - Managing Site-Wide Permissions and Groups
 - Configuring User Directories
 - Single sign-on for Confluence Data Center
- Managing System and Marketplace Apps
- Writing User Macros
 - User Macro Template Syntax
- Customizing your Confluence Site
 - Changing the Look and Feel of Confluence
 - Changing the Default Behavior and Content in Confluence
- Integrating Confluence with Other Applications
 - Linking to Another Application
 - Configuring Workbox Notifications
 - Integrating Jira and Confluence
 - Registering External Gadgets
 - Configuring the Office Connector
 - Managing Webhooks
- Managing your Confluence License
- Managing Confluence Data
 - Database Configuration
 - Site Backup and Restore
 - Attachment Storage Configuration
 - Confluence Data Model
 - Finding Unused Spaces or Pages
 - Data Import and Export
 - Import a Text File
 - Auditing in Confluence
 - Set retention rules to delete unwanted data
 - Data pipeline
- Configuring Confluence
 - Viewing System Information
 - Configuring the Server Base URL
 - Configuring the Confluence Search and Index
 - Configuring Mail
 - Configuring Character Encoding
 - Other Settings
 - Configuring System Properties

Downloads



Download the Confluence documentation in PDF format.

Other resources

Confluence installation and upgrade guide

Confluence Knowledge Base

Atlassian Answers

- Working with Confluence Logs
- Scheduled Jobs
- Configuring the Allowlist
- Configuring the Time Interval at which Drafts are Saved
- Configuring Confluence Security
 - Confluence Security Overview and Advisories
 - Proxy and HTTPS setup for Confluence
 - Configuring Secure Administrator Sessions
 - Confluence Cookies
 - Using Fail2Ban to limit login attempts
 - Securing Confluence with Apache
 - 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
 - Configuring RSS Feeds
 - Preventing and Cleaning Up Spam
- Configuring a Confluence Environment
 - Confluence Home and other important directories
 - Application Server Configuration
 - Starting Confluence Automatically on System Startup
- Performance Tuning
 - Cache Performance Tuning
 - Memory Usage and Requirements
 - Requesting Performance Support
 - Compressing an HTTP Response within Confluence
 - Garbage Collector Performance Issues
 - Troubleshooting Slow Performance Using Page Request Profiling
 - Confluence Diagnostics
 - Faster permissions service
- Data Collection Policy
- Administering Collaborative Editing
 - Possible Confluence and Synchrony Configurations
 - Configuring Synchrony
 - Set up a Synchrony cluster for Confluence Data Center
 - Migrate from a standalone Synchrony cluster to managed Synchrony
 - Troubleshooting Collaborative Editing
- Using read-only mode for site maintenance
- Administering the Atlassian Companion App
- Notifications from Atlassian
- Administer analytics

Getting Started as Confluence Administrator

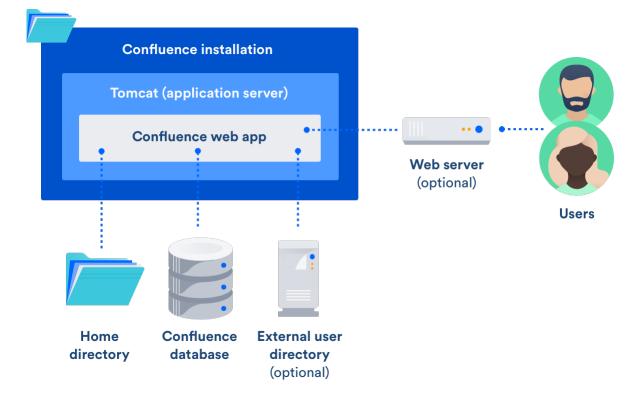
If you're just starting out as Confluence administrator, this page is for you. You'll find this page useful if your Confluence site is brand new, or if you're learning to administer an existing site.

Confluence is a Java-based web application. For the supported environments, there's 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 theConfluence Installation Guide for details.

Diagram: a Confluence installation

On this page:

- Quick access to admin functions via search
- How to administer and configure Confluence
- Getting started on a new Confluence site
- Getting to know an existing Confluence site



Quick access to admin functions via search

Start typing what you want to do into the Confluence search box at top right of the screen. The matching admin functions will appearwith a cog icon at the top of the search results.

Screenshot: searching for admin options



Even faster via /: Press / on your keyboard, then continue typing the action you want.

Notes about finding admin functions via search:

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

How to administer and configure Confluence

After installing Confluence, you will perform the initial configuration via a web interface called the Confluence Setup Wizard.Introducing the Confluence Administration Console: From this point onwards, many of the admin functions are available from the Confluence Administration Console, which is part of the Confluence web interface. If you have administrative permissions, you'll have access to the Confluence Administration Console via your web browser, using the standard Confluence URL for your site.

To access the Confluence Administration Console:

- 1. Open your Confluence URL in your web browser.
- 2. Choose General Configuration in the header.

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

Getting started on a new Confluence site

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

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

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

Getting to know an existing Confluence site

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

- Understand the Confluence permission scheme. SeePermissions and restrictions.
- Get to know the power of Marketplace apps (also known as add-ons or plugins), for extending and customizing your Confluence site. See Managing System and Marketplace Apps.
- Investigate more ways of customizing Confluence. See Customizing your Confluence Site.

Manage Users

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



Managing 500+ users across Atlassian products?

Find out how easy, scalable and effective it can be with Crowd!

Seecentralized user management.

Confluence user management

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

- Add and Invite Users
- Delete or Disable Users
- Managing Site-Wide Permissions and Groups

On this page:

- Confluence user management
- Authentication
 - Seraph
 - XML-RPC and SOAP authentication
 - Password authentication
 - SAML single sign-on

Related pages:

Configuring Confluence Security

By default, Confluence stores its users and groups in the Confluence database. This is called the internal directory. You can choose to connect Confluence to an external userbase instead, such as Microsoft Active Directory or another LDAP server. You can also use AtlassianCrowdandJiraapplicationsas 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. SeeConfiguring User Directories.

Authentication

Seraph

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

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

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

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

XML-RPC and SOAP authentication

Normally, requests for Confluence's XML-RPC and SOAP APIs (deprecated) 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.

SAML single sign-on

If you have a Confluence Data Center license you can connect Confluence to your SAML 2.0 identity provider for authentication and single sign-on.

SeeSingle sign-on for Confluence Data Center for more information.

Add and Invite Users

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

- By user signup: Ifusersignupis enabled on your Confluence site, people can add themselves as users of the site.
- Via an invitation link: You can invite people to sign up by sending them an invitation link. You can copy and paste the link, orprompt Confluence to send the link in an email message.
- By adding users manually: If you have Administrator or System Administratorpermission, youcan manually add new users.
- Via an external user directory: SeeConfiguring User Directories.

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

Allow user signup

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

On this page:

- Allow user signup
- Manageuser signupnotifications
- Invite people to sign up
- Reset the invitation link
- Add users manually
- Notes

Related pages:

- Manage Users
- Setting Up Public Access
- Configuring a Server for Outgoing Mail

You can restrict the signup to people whose email addresses are within a given domain or domains. This is useful if you want to ensure that only people within your organization 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 Administratorpermissions to change the signup options.

To set the user signup options:

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

Manageuser signupnotifications

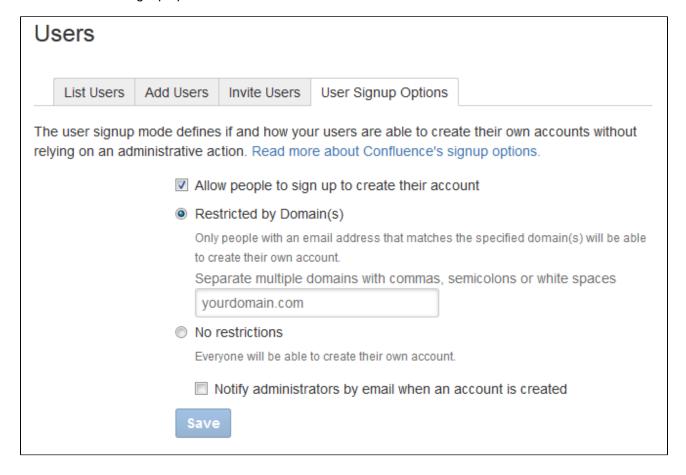
By default, Confluence will send an email notification to all Confluence administrators whenever someone signs up to your Confluence site. The administrators (people with Confluence Administrator or System Administrator permissions) will receive this message when someone signs up either by clicking the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

To disable this notification:

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

4. ChooseSave

Screenshot: User signup options



Invite people to sign up

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

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

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

To invite people to sign up:

- 1. Choose >User management
- 2. Select the Invite Users tab
- 3. Do either of the following:
 - Copy theInvitation Linkand paste it into an email message, or onto a page on your intranet, for example
 - Alternatively, prompt Confluence to send an email message for you:
 - a. Enter one or more email addresses in the field labeledEmail To Separate the addresses with commas. For example:john@example.com, sarah@example.com
 - b. Change the Message if you want to
 - c. ChooseSend

Reset the invitation link

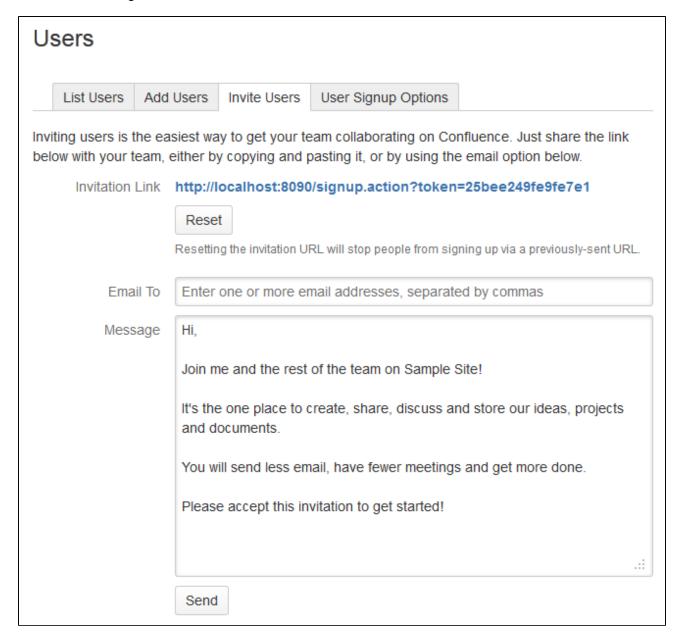
The invitation link includes a security token, like this:

 $\verb|http://confluence.example.com/signup.action?token=d513a04456312c47|$

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

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

Screenshot: Inviting users



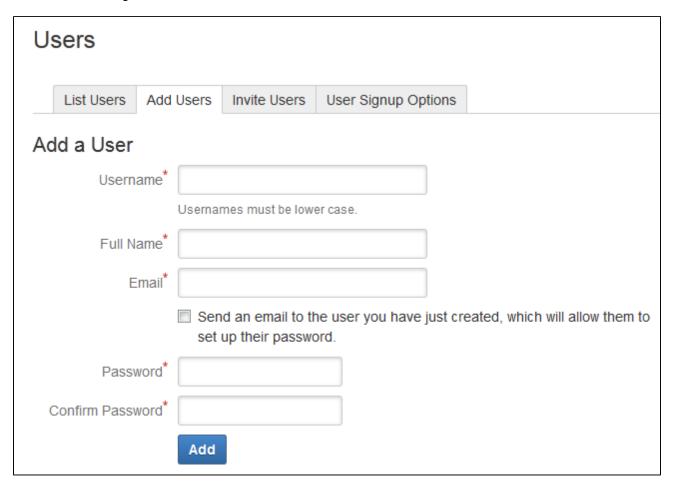
Add users manually

To add a new user:

- 1. Choose User management
- 2. Select the Add Userstab
- 3. Enter the user's details

- Choose whether Confluence should send anemailmessage informing the person of their new username
 - The email message will contain a link that the person can use to reset their password.
- 5. ChooseCreate

Screenshot: Adding users



Notes

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

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

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

SeeManaging Multiple Directories.

- Email server required for domain restricted signup and for invitations You need to set up amail serverfor 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 Configuration Security Configuration. See Disabling the Built-In User Management.

• Avoid hash, slash and question characters in usernames - there is a known issue where users with #, ? or / in their username cannot create spaces. See

CONFSERVER-43494 GATHERING IMPACT and CONFSERVER-13479 GATHERING IMPACT for more information.

Delete or Disable Users

When someone leaves your organisation, or no longer needs to use Confluence, you can either disable their user account, unsync it from any external directories, or delete it entirely.

On this page:

- Delete, disable, or unsync?
- Disable a user account
- Unsync a user account
- Delete a user account
 - Delete from an internal Confluence directory or read/write external directory
 - Delete from a read-only external directory, or multiple external directories
 - How deleted users appear to other people
- Only remove access to Confluence
- Limitations when deleting a user account
 - Free text is not anonymised
 - Data stored in Synchrony is not deleted immediately
 - Personal spaces are not deleted
 - Workbox notifications don't disappear immediately
 - Data stored by thirdparty apps is not deleted

Related pages:

- Manage Users
- Configuring User **Directories**

Managing 500+ users across Atlassian products? Find out how easy, scalable and effective it can be with Crowd! Seecentralized user management.

Delete, disable, or unsync?

It's useful to understand the difference between disabling a user account, unsyncing it from an external directory, and permanently deleting it from Confluence.

In most situations disabling or unsyncing a user account is the appropriate way to prevent a user from accessing Confluence, for example when someone leaves your organisation. However, if you do need to remove someone's name and personal details, you can permanently delete their user account.

When an user account is disabled:

- The user won't be able to log in.
- The user won't be included in your license count.
- · People won't be able to see the user in the People directory, mention them, or select their name /username as a search filter.
- Their full name will still appear on any spaces or content they created.
- They will be listed in User Management admin screens.
- Their account can be re-enabled (this will restore the connection to their content).
- Any content they created will be maintained.

When a user account is unsynced from all external directories:

- The user won't be able to log in.
- The user won't be included in your license count.
- People won't be able to see the user in the People directory, mention them, or select their name /username as a search filter.
- Their username will appear on any spaces or content they have created.
- They will only be listed on the Unsynced from Directory tab of the User Management admin screens.
- Their accountwill be restored if they are resynced with Confluence.
- Any content they created will be maintained.

When a user account is deleted:

- The user won't be able to log in.
- The user won't be included in your license count.
- People won't be able to see the user in the People directory, mention them, or select their name /username as a search filter.
- An anonymised alias will appear on any spaces or content they have created.
- They won't be listed in User Management admin screens.
- Their accountis deleted and anonymised permanently, and can't be restored.
- Any content they created will be maintained.

Disable a user account

How you disable a user account depends on whether you manage users in the internal Confluence directory, or in an external user directory (for example Jira, Crowd, Active Directory).

You need the Confluence Administrator global permission to do this.

To disable a user account:

- 1. Go to User management.
- 2. Search for the user you want to disable.
- 3. Choose Disable.

If there is no Disable option, it is likely that Confluence has a read-only connection to an external directory. If this happens, you'll need to remove the user's access to Confluence in your external directory. This might be done by disabling the user in that directory, or changing their group membership so they are not synced to Confluence. They will be treated as an unsynced user in Confluence after your next directory sync.

Unsync a user account

You unsync a user account by excluding it from the accounts to be synchronized with Confluence in your external directory.SeeSynchronizing Data from External Directoriesto learn more about how directory sync works.

To view users who have previously been synchronized with Confluence, but were not present in the last directory sync. go to User management>Unsynced from Directory.

It's important to note that user accounts can be unsynced intentionally, or because of a problem with your external directory. Don't assume all unsynced user accounts are to be deleted.

Delete a user account

Deleting a user is permanent, socannot be undone. If you're trying to reduce your license count, or only need to remove a someone's access to Confluence, you should disable their account instead.

How you delete a user account depends on whether you manage users in:

- an internal directory, or a single read/write external directory (such as Jira, Crowd, or Active Directory)
- multiple external directories, or a singleread-only external directory (such as Jira, Crowd, or Active Directory).

The delete process can take several minutes, depending on the amount of content the person had created. It can also flood your index queue, as it reindexes all pages the user contributed to, so you may want to perform this task at a time that won't impact other users.

You need the Confluence Administrator global permission to do this.



It's important to note that the person's content is not removed when you delete their account. Find out about limitations and what personal information may need to be removed manually.

Delete from an internal Confluence directory or read/write external directory

To permanently delete a user stored in the internal Confluence directory, or a single external directory that has a read/write connection to Confluence:

- 1. Go to >User management.
- 2. Search for the user you want to delete.
- 3. Choose Delete.
- 4. Wait for confirmation that the delete process is complete. This can take a few minutes.

The user account will be deleted from Confluence, and their name replaced with an anonymised alias. This can't be undone.

Delete from a read-only external directory, or multiple external directories

Deleting a user stored in a read-only external directory or in multiple external directories, is a two-step process. You need to remove them from all external directories and perform a directory resync before they can be deleted from Confluence.

To permanently delete a user stored in multiple external directories, or an external directory that has a readonly connection to Confluence:

- In your external directory, remove the user. If the user exists in multiple directories, remove them from each one.
- 2. In Confluence, go to User management> Unsynced from directory.
- 3. Search for the username of the person you want to delete. If the user doesn't appear, wait for Confluence to sync your external directory (or trigger a re-sync if you usually do this manually). SeeSynchronizing Data from External Directories.
- 4. Choose Delete.
- 5. Wait for confirmation that the delete process is complete. This can take a few minutes.

The user account will be deleted from Confluence, and their name replaced with an anonymised alias. This can't be undone.

How deleted users appear to other people

Once a user account has been deleted their identity will be anonymised throughout Confluence in places like the page byline, mentions, comments, and page history.

- full names be replaced with an alias like 'user-38782'
- usernames will be replaced with the user key (a long string of characters).
- their profile picture will be replaced with a default image.

The alias and user key stays the same throughout the site. This means people can see that pages and comments were made by the same person, but not know the identity of that person.

Only remove access to Confluence

If you want to remove someone's access to Confluence, but retain their user account (or you can't disable their account for some reason), you can do this by changing their group membership.

- 1. Create a group, for example no-confluence-access
- 2. Go to Configuration School Permissions.
- 3. Make sure theno-confluence-accessgroup doesn't have Can Use Confluence permission.
- 4. Change the user's group membership so they are only a member of theno-confluence-accessgro up.

If you don't manage groups in Confluence (for example group membership is always synced from your external directory), the same principles apply, but you'll need to change the user's membership in your external directory.

Remember that permissions are additive, so just being a member of a group without Confluence access is not enough. To ensure the user can't log in to Confluence they must not be a member of ANY group that has the Can Use Confluenceglobal permission (in any user directory).

Limitations when deleting a user account



The ability to delete and anonymize a user account was added in Confluence 6.13.

For earlier Confluence versions there's a workaround you can use to permanently delete a user account via the database. See Right to erasure in Confluence Server and Data Center.

You can also head to Confluence Server and Data Center GDPR support guidesto read more about Confluence and GDPR generally.

There are some situations where personal information may still be stored in Confluence after you have deleted a user account, andthe delete process does not remove any actual content, for example if someone has typed the user's name in plain text on a page, or if it is contained in an attached file.

Free text is not anonymised

Deleting a user does not delete any Confluence content (such as pages, files, or comments). This means that any references to a person's full name, user name, or other personal information that were entered as free text will remain after the user account is deleted. Text entered in the link text of a link or mention are also considered free text (for example if you mention someone on a page and change the mention link text to use just their first name, or a nickname).

Links to the deleted user's personal space (which contains their username in the URL) will also remain after their personal space has been deleted, if the links were inserted as a web link or free text.

We suggest searching for the deleted person's name and username to see if there is any residual content left behind.

There are also a couple of known issues that will require manual cleanup:

- When multiple people are mentioned on a task, only the first (the assignee) is replaced with the anonymised alias. This is due to an existing bug where subsequent mentions aren't indexed.
- If the user to be deleted is listed on the All Updates tab on the dashboard at the point they are deleted, their updated items will appear twice, once with their anonymised alias and once with their username. They will drop off the All Updates tab as new updates occur, but their username will still be listed in the search index. A full site reindex will resolve this issue.

Data stored in Synchrony is not deleted immediately

If you have collaborative editing enabled, every keystroke in the editor is stored by Synchrony in the Confluence database. This means that any references to a person's full name, user name, or other personal information typed in the editor will remain in the Synchrony tables in the database.

From Confluence 7.0 we provide two scheduled jobs for removing Synchrony data:

- Synchrony data eviction (soft)
- Synchrony data eviction (hard)

The soft eviction job runs regularly in the background. The hardeviction job is available for when you need to remove Synchrony data more aggressively, for example after you have deleted a user, and is disabled by default.

SeeHow to remove Synchrony datato learn more about how these jobs work.

Personal spaces are not deleted

When you delete a user, their personal space is not automatically deleted, as it may contain content owned by your organization. This means that:

- their username will still be visible in the space URL
- their name may still be visible in the space title or homepage title

We recommend moving any pages or blogs that you want to keep to a new space, and then deleting the personal space entirely. Any links to the personal space will be updated with the new space key automatically when the pages are moved, unless they have been added as a web link or free text.

If space permissions prevent you from accessing the user's personal space, a member of the confluenceadministrators super group will be able to access the space. They can then grant another user permission to administer the space, or delete it themselves.

Workbox notifications don't disappear immediately

The deleted user's full name will still appear in any existing workbox notifications. For example if the deleted user had shared a page with another user, the notification will still appear in that user's workbox for up to 28 days. SeeWorkbox Notifications for more information about how long a workbox notification is accessible before it is automatically deleted.

Data stored by third-party apps is not deleted

When you delete a user, we replace the person's full name and username with an anonymous alias in all the places we know about, such as mentions, page history, and in macros.

If you have installed apps from the Marketplace, there is a chance that these apps are storing data in their own tables in the Confluence database. Refer to the documentation for your app to find out the best way to remove this data.

Restore Passwords To Recover Admin User Rights

If you're unable to log in to Confluence as an administrator (for example, you've lost the administrator password) you can start Confluence in recovery mode to recover your admin user rights.

If you know the admin username, and it has a valid email address, you can reset the password using the forgot password link on the log in screen. We'll send a link to your admin email account to reset your password.

On this page:

 Use recovery mode to restore access

As an administrator, you may find yourself locked out of Confluence because:

- You've imported a site from Cloud, and it does not contain a system administrator account.
- You've forgotten the password to the administrator account, and don't have access to the email address associated with it.
- You're using an external directory or Jira for user management, have disabled the built in user management, and your external directory is not currently available.
- You need to make a change to the configuration of an external user directory in Confluence while that directory is not available.

In any of these situations you can use recovery mode to restore administrator access to Confluence.



(i) Using Confluence 6.5.0 or earlier? You'll need to use the database method to recover your admin user rights. Seethe earlier documentation.

Use recovery mode to restore access

Recovery mode works by creating a virtual user directory with a temporary admin account. You set the password for this admin account when applying the system property. Users can continue to log in and access Confluence while it is in recovery mode.

To recover administrator user rights:

- 1. Stop Confluence.
- 2. Add the following system property, replacing <your-password>with a unique, temporary password.

-Datlassian.recovery.password=<your-password>

The way you do this depends on how you run Confluence. SeeConfiguring System Properties for more information on how to apply system properties.

- 3. Start Confluenceusing your usual method.
- 4. Log in to Confluence with the usernamerecovery adminand the temporary password you specified in the system property.
- 5. Reset the password for your existing admin account, or create a new account and add it to the appropriate administrator group.
- 6. Confirm that you can successfully log in with your new account.
- 7. Stop Confluence.
- 8. Remove the system property you added earlier.
- 9. Restart Confluence using your usual method (manually or by starting the service).

Good to know:

- Remove the system property as soon as you have restored admin access.
- Don't leave Confluence in recovery mode, or use the recovery_admin account as a regular administrator account.
- Your temporary password should be a unique. Don't use an existing password or the one you intend to use for your admin account.

Edit User Details

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

Edit a user's details

- 1. Choose User management
- 2. Do either of the following:
 - ChooseShow all users to list everyone in the 'confluenceusers' or 'users' group
 - Enter ausername, full name or email address in the Find Userfi eld and hitSearch



If you're already viewing someone's profile, chooseAd minister User in the sidebar.

2. Select the user you want to manage

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

- View Profile View the user's profile.
- Edit GroupsAdd or removethis 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. SeeChang e a Usernamefor information.
- Delete Profile Picture remove current and all previous profile pictures uploaded by the user.
- Set Password Edit the user's password details.
- Disable You candisable(i.e. deactivate) access for a user who no longer needs access to Confluence.
- Delete You can permanently deletea user, and replace their full name and username with an anonymous alias.

View User: cassie « Back to Users View Profile Edit Groups Edit Details Delete Profile Picture Set Password Delete Disable User cassie Full Name Cassie Owens Email cassie@email.com Directory Confluence Internal Directory Created Apr 26, 2017 15:30 Last Updated Feb 01, 2018 17:10 Login Last Login: Sep 28, 2018 15:09 Last Failed Login: Jan 17, 2018 08:16 **Total Failed Login Count: 1 Current Failed Login Count: 0** Groups <a>a confluence-users Groups <a>a confluence-users

On this page:

- Edit a user's details
- Reset login count

Related pages:

- Delete or Disable Users
- Adding or Removing Users in Groups
- Add and Invite Users

Reset login count

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

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



Multiple user directories

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

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

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

SeeManaging Multiple Directories.

Change a Username

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

Each active users must have a unique username, so no two active users can have the same username. You can, however, assign the username of adisa bled userto another active user.

The procedure for changing a username depends on where you manage your users. SeeConfiguring User Directoriesfor more info.

On this page:

- Confluencemanaged users
- Users managed in an external directory
- Notes

Confluence-managed users

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

To change a username:

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

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

Users managed in an external directory

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

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

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

Notes

Some important things to note about changing usernames:

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

Managing Site-Wide Permissions and Groups

Permissions determine what people can do on your Confluence site. Confluence recognizes permissions at site level and at space level, as well as page-level restrictions.

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

You can also set the access levels for anonymous usersor deny access to unlicensed users from linked applications, such as Jira Service Management.



Managing 500+ users across Atlassian products?

Find out how easy, scalable and effective it can be with Crowd!

Seecentralized user management.

Related pages:

- Confluence Security Overview and Advisories
- Global Permissions Overview

Confluence Groups for Administrators

Grouping users in Confluence is a great way to cut down the work required when managing permissions and restrictions.

Groups can be used when setting:

- global permissions
- space permissions
- page restrictions.

If your site has a lot of users, using groups can really simplify your permissions management over time.

On this page:

- Default groups
- Create a new group
- Delete a group
- Confluenceadministrators super group
- About multiple user directories

Related pages:

- Manage Users
- Global Permissions Overview

Default groups

The two default groups in Confluence are:

- confluence-users- this is the default group into which all new users are usually assigned. In most sites this is the group that provides the permission to log in to Confluence.
- **confluence-administrators** this super group grants the highest level of administrator permissions. Members of this can view all pages, including restricted pages. While they can't edit existing pages, they can add, delete, comment, restore page history, and administer the space. They can also access the admin console and perform all administrative tasks.

Create a new group

To add a new group:

- 1. Go to O S General Configuration Groups.
- ChooseAdd Group.
- 3. Enter a name for your group and choose Save. Group names must be lower case.

You're now ready to startadding users to the group.

Delete a group

To delete a group:

- 1. Go to O Seneral Configuration Groups.
- 2. Choose **Delete**next to the group you want to remove.

Deleting a group removes all permission restrictions associated with it. This means that members of this group may loose access to spaces that use this group to grant their permissions, and pages / blogs that are only only restricted to this group will become available to all confluence users.

If you have Confluence Data Center, you can Inspect permissions to find out which spaces are using this group, before you delete it.

Confluence-administrators super group

The Confluence administrator global permission and the confluence-administrators group are not related. Going by the names, you would think they are the same thing, but they're not. Granting a user or a group Confluence administrator global permission allows access to a sub-set of administrative functions. Granting membership to the confluence-administratorsgroup grants the highest possible permissions, with complete access to all content and administration functions.

To find out more about what the various levels of administrator can do, seeGlobal Permissions Overview.

About multiple user directories

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

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

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

SeeManaging Multiple Directories.

Adding or Removing Users in Groups

Confluence Groupsarea great way to cut down the work required when managing permissions and restrictions.

You can edit group membership in two places:

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

You need Confluence Administrator or System Administrator global permission to do this.

On this page:

- Add people to a group
- Remove people from a group
- About multiple directories

Related pages:

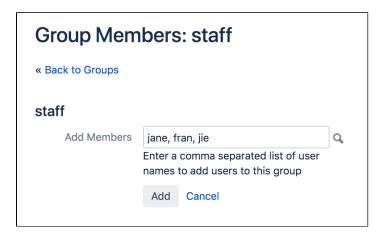
- Manage Users
- Confluence Groups
- Global Permissions Overview

Add people to a group

To add members to a group from the Groups screen:

- 1. Go to General Configuration > Groups.
- 2. Choose the group.
- 3. Choose Add Members.
- 4. Type the username of the person you want to add to the group. You can add multiple usernames, separated by a comma.
- 5. Choose **Add** to add members to the group.

Screenshot: Adding members



You can also change a user's group membership in the user management screen. Navigate to the user, then choose **Edit groups**, and select the groups the person should be a member of.



Remove people from a group

To remove members from a group:

- 1. Go to General Configuration > Groups.
- 2. Choose the group.
- 3. Choose the **Delete user from group**icon next to the user you want to remove.

You can also change a user's group membership in the user management screen. Navigate to the user, then choose**Edit groups**, and deselect the groups.

About multiple directories

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

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

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

SeeManaging Multiple Directories.



Managing 500+ users across Atlassian products? Find out how easy, scalable and effective it can be with Crowd! Seecentralized user management.

Global Permissions Overview

Global Permissions determine what a user can do at a site level, including whether they can log in, create spaces, or administer the site.



Unsure about the best way to set permissions in your site? Check out our Per missions best practices guide.

On this page:

- Overview of global permissions
- Grant global permissions
- Revoke global permissions
- System Administrator and Confluence Administrator permissions compared
- Confluence-administrators super group
- Troubleshooting

Overview of global permissions

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

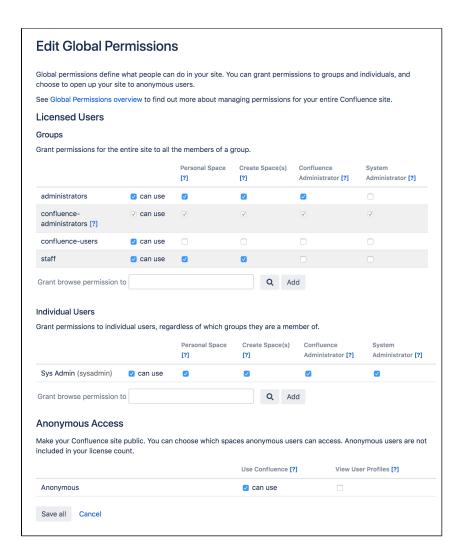
Global Permission	Description	
Can Use	This is the most basic permission that allows users to log in to this Confluence site. Users with this permission contribute to your licensed users count.	
Personal Space	Allows the user to create a personal space. The space key will be their username.	
Create Space(s)	Allows the user to create new spaces in your site. When a user creates a space they are automatically granted admin permissions for that space.	
Confluence Administrator	Allows the user to access the Confluence administration console, and perform basic administrative tasks such as adding users, changing group memberships, and changing the colour scheme of the site. See the detailed comparison of administrator permissions below.	
System Administrator	Allows the user to access the Confluence administration console and perform all administrative tasks. See the detailed comparison of administrator permissions below.	

Grant global permissions

To grant global permissions to a user or group:

- 1. Go to General Configuration Global Permissions
- 2. ChooseEdit Permissions.
- 3. Do one of the following:
 - a. Enter a group name in the **Grant browse permissions** field in the **Group** section
 - b. Enter a username in the Grant browse permissions field in the Individual Users section
- 4. ChooseAdd.
- 5. The user or group will appear in the list. Select the permissionsyou want to grant.
- 6. Choose Save all.

Screenshot: Editing global permissions



Revoke global permissions

To revoke the global permissions for a user or group:

- 1. Go to General Configuration Global Permissions
- 2. ChooseEdit Permissions.
- 3. Locate the user or group you want to edit, and deselect all checkboxes.
- 4. ChooseSave all.

If you are attempting to revoke permissions for an individual user, and they are not listed, you will need to check which groups they are a member of, and remove them from any groups that grant the global permission.

System Administrator and Confluence Administrator permissions compared

The table below lists the parts of the admin console that can be accessed by people with the Confluence Administrator and System Administrator global permissions.

Members of the confluence-administratorssuper group have System Administrator global permissions by default, as well as the ability to view all spaces and pages.

Admin console	Confluence administrator	System administrator
------------------	--------------------------	----------------------

Configuration	 General Configuration, except: Base URL Connection timeout Further configuration, except: Remote API Languages Shortcut Links Global Templates and Blueprints Recommended Updates Email Configure Code Macro, except add new languages WebDAV Configuration 	 General Configuration Further configuration Backup Administration Languages Shortcut Links External Gadgets Global Templates and Blueprints Recommended Updates Email Mail Servers User Macros In-app Notifications Spam Prevention PDF Export Language Support Configure Code Macro Office Connector WebDAV Configuration
Marketplace	 Find new apps Manage apps, except: Upload an app. 	Find new appsManage apps
Users & security	 Users Groups SSO 2.0 Security Configuration, except: External user management Append wildcards to user and group searches Enable Custom Stylesheets for Spaces Show system information on the 500 page RSS settings XSRF Protection Attachment download security Global Permissions Space Permissions Inspect Permissions (Data Center) 	 Users Groups SSO 2.0 Security Configuration Global Permissions Space Permissions Inspect Permissions (Data Center) User Directories Whitelist
Look and feel	 Themes Color Scheme Site Logo and Favicon PDF Layout PDF Stylesheet Sidebar, header and footer Default Space Logo 	 Themes Color Scheme Layouts Stylesheet Site Logo and Favicon PDF Layout PDF Stylesheet Sidebar, header and footer Default Space Logo Custom HTML
Upgrade	Latest upgrade report	Latest upgrade reportPlan your upgrade

Administration	 Macro Usage Audit Log Content Indexing License Details Application Links (OAuth only) Application Navigator 	 Mobile apps Collaborative Editing Maintenance (Data Center) System Information Macro Usage Audit Log Rate limiting (Data Center) Backup & Restore Content Delivery Network (Data Center) Content Indexing Mail Queue Scheduled Jobs Cache Management License Details Logging and Profiling Application Links Application Navigator Analytics Troubleshooting and support tools Clustering (Data Center)
Atlassian Cloud	Migration assistant	Migration assistant

Confluence-administrators super group

The Confluence administrator global permission and the confluence-administrators group are not related. Going by the names, you would think they are the same thing, but they're not. Granting a user or a group Confluence administrator global permission allows access to a sub-set of administrative functions. Granting membership to the confluence-administratorsgroup grants the highest possible permissions, with complete access to all content and administration functions.

When you install Confluence you'll be prompted to create a system administrator account. This user will be a member of theconfluence-administratorssuper group.

What can members of this group do?

This group provides the highest level of permission in your site, and these permissions can't be edited. People in this group can:

- perform all administrative tasks
- access all spaces
- access all pages, including pages with view restrictions.

Restricted pages and blog posts are not visible to members of theconfluence-administratorsgroup in the dashboard, blog roll, search and most macros, but are visible if the user has the page URL, or in the:

- page tree in the sidebar
- pages index page
- reorder pages screen
- page tree macro
- content by user macro

Members of this group can't edit pages by default. They need to grant themselves space permissions, or add themselves to the page restrictions in order to edit.

Should I use the confluence-administrators group?

Some organisations use the confluence-administrators group extensively, while others choose to limit its membership to just one special admin account, to limit the number of people who can see all content by default. System administrators can perform all the same administrative tasks, so membership of this group is not a requirement.

If you do decide not to use this group, be aware that the group can't be deleted, and that people with System Administrator global permissions can add themselves to this group.

Troubleshooting

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.
- If you're unable to log in to Confluence as an administrator (for example, you've lost the administrator password) you can start Confluence in recovery mode to recover your admin user rights. SeeRestore Passwords To Recover Admin User Rights.

Setting Up Public Access

If you use Confluence for documentation, as a knowledge base, you might want to make your site public. This means people don't need to log in to use Confluence.

On this page:

- Allow anonymous access to the site
- Disable anonymous access to the site
- Allow anonymous access to a space
- Alternatives to making your site public

Related pages:

- Configuring Captcha for Spam Prevention
- Add and Invite Users
- Global Permissions Overview

Allow anonymous access to the site

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

To enable anonymous access to your site:

- 1. Go to General Configuration Global permissions.
- 2. Choose Edit Permissions.
- 3. In the **Anonymous Access** section, select the **Can use** checkbox. You can also choose whether to allow anonymous users to see user profiles.
- 4. Choose Save All.

Disable 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.

Any spaces that granted permissions to anonymous users will still be available to all logged in users, until you remove these permissions from each space.

Allow anonymous access to a space

Allowing anonymous access to your site does not automatically allow people who are not logged in to see all the spaces in your site.

Space administrators must grant anonymous users permissions on a space by space basis. SeeMake a Space Public to find out how to do this.

Alternatives to making your site public

You can allow people to sign up for usernames themselves, and choose other options for user signup and invitations. See Add and Invite Users.

Revoke access for unlicensed users from Jira Service Management

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

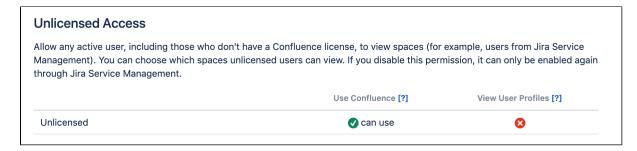
To revoke access for unlicensed users:

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

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

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

Screenshot: Unlicensed access section of the Global Permissions page.



This section only appears on the Global Permissions page in Confluence if you have linked a space to your service project (as a Knowledge base), and chosen to allow all active users and customers to access without a Confluence license. SeeSet up a knowledge base for self-servicein the Service Management Server and Data Center documentation for more info.

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 Jiraapplications as directory managers.

On this page:

- Configuring User Directories in Confluence
- Connecting to a Directory
- Updating Directories

Related pages:

- Add and Invite Users
- Managing Site-Wide Permissions and Groups



Managing 500+ users across Atlassian products? Find out how easy, scalable and effective it can be with Crowd! Seecentralized user management.

Configuring User Directories in Confluence

To configure your Confluence user directories:

- 1. Choose the cog icon , then choose General Configuration
- 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 or Jira 4.3 or later. See Connecting to Crowd or Jira for User Management.

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

Updating Directories

Limitations when Editing Directories

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

This limitation applies to all directory types. For example:

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

In some situations, reordering the directories will change the directory that the current user comes from, if a user with the same username happens to exist in both. This behavior 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

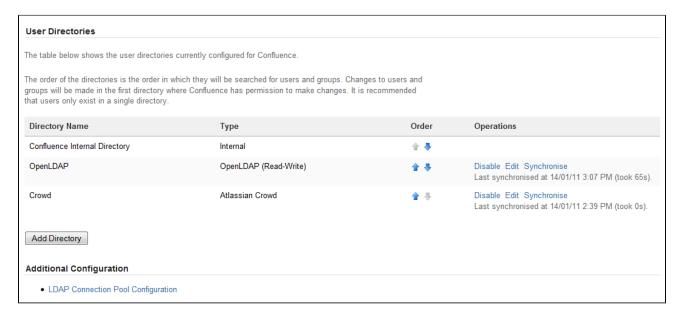
The recommended way to edit directory configurations is to log in as an internal user when making changes to external directory configuration.

⚠ 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 recognize 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.



Screenshot above: Configuring user directories

Configuring the Internal Directory

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

Overview

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

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

Diagram of Possible Configuration

Authentication, updates and queries Confluence database (internal directory)

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

On this page:

- Overview
- Diagram of Possible Configuration

Related pages:

- Configuring User Directories
- 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.



Managing 500+ users across Atlassian products?

Find out how easy, scalable and effective it can be with Crowd!

Seecentralized user 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
- 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
- 2. Click **User Directories** in the left-hand panel.
- 3. **Add** a directory and select one of these types:
 - Microsoft Active Directory This option provides a quick way to select AD, because it is the most popular LDAP directory type.
 - LDAP You will be able to choose a specific LDAP directory type on the next screen.
- 4. Enter the values for the settings, as described below.
- 5. Save the directory settings.
- 6. Define the directory order by clicking the blue up- and down-arrows next to each directory on the 'User Directories' screen. Here is a summary of how the directory order affects the processing:
 - Changes to users and groups will be made only in the first directory where the application has permission to make changes.

On this page:

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

Related pages:

Configuring User Directories

• The order of the directories is the order in which they will be searched for users and groups (by default Confluence aggregates group membership from all directories, so the order does not impact membership itself).

For details see Managing Multiple Directories.

Server Settings

_	Description
Name	Enter a meaningful name to help you identify the LDAP directory server. Examples: • Example Company Staff Directory • Example Company Corporate LDAP
y 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.
Hostna me	The host name of your directory server. Examples: • ad.example.com • ldap.example.com • opends.example.com
Port	The port on which your directory server is listening. Examples: • 389 • 10389 • 636 (for example, for 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.
	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
	① By default, all users can read the uSNChanged attribute; however, only administrators or users with relevant permissions can access the Deleted Objects container. The specific privileges required by the user to connect to LDAP are "Bind" and "Read" (user info, group info, group membership, update sequence number, deleted objects), which the user can obtain by being a member of the Active Directory's built-in administrators group.
	Note that the incremental sync will fail silently if the Active Directory is accessed by a user without these privileges. This has been reported asCWD-3093.

Password The password of the user specified above.

Note: Connecting to an LDAP server requires that this application log in to the server with the username and password configured here. As a result, this password cannot be one-way hashed - it must be recoverable in the context of this application. The password is currently stored in the database in plain text without obfuscation. To guarantee its security, you need to ensure that other processes do not have OS-level read permissions for this application's database or configuration files.

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 ldp.exe which is useful for finding out and configuring the the LDAP structure of your server.
Addition al 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
Addition al 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



 \bigwedge If no value is supplied for **Additional User DN** or **Additional Group DN** this will cause the subtree search to start from the base DN and, in case of huge directory structure, could cause performance issues for login and operations that rely on login to be performed.

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.
	1 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
	CONFSERVER 28621 - User Loses all Local Group Memberships If LDAP Sync is Unable to find the User, but the User appears again in subsequent syncs CLOSED
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 Member ships	Option available in Confluence 3.5 and later, and JIRA 4.3.3 and later. 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. In Confluence 3.5 to Confluence 3.5.1: 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. In Confluence 3.5.2 and later, and JIRA 4.3.3 and later: 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 not be added automatically to any groups. This change in behavior 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, authorization 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 • confluence-users, jira-administrators, jira-core-users

Advanced Settings

Setting	Description
Enable Nested Groups	Enable or disable support for nested groups. Some directory servers allow you to define a group as a member of another group. Groups in such a structure are called <i>nested groups</i> . Nested groups simplify permissions by allowing sub-groups to inherit permissions from a parent group.
Manage User Status Locally	If true, you can activate and deactivate users in Crowd independent of their status in the directory server.

If true, user accounts marked as expired in Active Directory will be automatically removed. For cached directories, the removal of a user will occur during the first synchronization after the account's expiration date. Note: This is available in Embedded Crowd 2.0.0 and above, but not available in the 2.0.0 m04 release. Follow Referral applies and the search will retrieve sets of data rather than all of the 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 at once. Enter the desired page size that is, the maximum number of search results are once. Enter the desired page size that is, the maximum number of search results are once. Enter the desired page size that is, the maximum number of search results are once. Enter the desired page size that is, the maximum number of search results are once. Enter the desired page size that is, the maximum number of search results at once. Enter the desired page size that is, the maximum number of search results at once. Enter the desired page size that is, the maximum number of search results at once. Enter the desired page size that is, the maximum number of search results at once. Enter the desired page size that is, the maximum number of search results at once. Enter the desired page size that is, the maximum number of search results at once and the search results at once and t		
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 Referrals Referrals Referrals In the 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. Raive 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. If this checkbox is selected, the application will compare DNs to determine if they are equal. If this checkbox is not selected, the application will parse the DN and then check the parsed version. Enable Incremental synchronization if you only want changes since the last synchronization to be queried when synchronization if you only want changes since the last synchronization to be queried when synchronizing a directory. A Please be aware that when using this option, the user account configured for synchronization must have read access to: The usinchanged attribute of all users and groups in the directory that need to be synchronized. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchron Canada attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory server. The application will send a request to your directory server every x min	out expired	cached directories, the removal of a user will occur during the first synchronization after the
Paged 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 java.naming.referral) configuration setting. It is generally needed for Active Directory servers configured without proper DNS, to prevent a javax.naming.PartialResultException: Unprocessed Continuation Reference(s)* error. Naive DN Matching 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 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 Synchron nization Enable enable incremental synchronization if you only want changes since the last synchronization to be queried when synchronizing a directory. • Please be aware that when using this option, the user account configured for synchronization must have read access to: • Theusinchanged attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchron ization is the process by which the application updates its internal store of user data to agree with the data on the directory		
Referrals uses the node referral (JNDI lookup java.naming.referral) configuration setting. It is generally needed for Active Directory servers configured without proper DNS, to prevent a javax.naming.PartialResultException: Unprocessed Continuation Reference(s) error. Naive DN History 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 Increme hall synchronization if you only want changes since the last synchronization to be queried when synchronizing a directory. Please be aware that when using this option, the user account configured for synchronization must have read access to: TheusnChanged attribute of all users and groups in the directory that need to be synchronized. The objects and attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchron ization 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. The time, in seconds, to wait for a response from a search operation. A value of 0	Paged	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
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 Increme hall Synchronization if you only want changes since the last synchronization to be queried when synchronizing a directory. Please be aware that when using this option, the user account configured for synchronization must have read access to: TheusnChanged attribute of all users and groups in the directory that need to be synchronized. The objects and attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchronization 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 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 60 seconds.		uses the node referral (JNDI lookup java.naming.referral) configuration setting. It is generally needed for Active Directory servers configured without proper DNS, to prevent a
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 synchronization if you only want changes since the last synchronization to be queried when synchronizing a directory. Please be aware that when using this option, the user account configured for synchronization must have read access to: Theusnchanged attribute of all users and groups in the directory that need to be synchronized. The objects and attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchronization 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 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 60 seconds.	DN	enable naive DN matching. Using naive DN matching will result in a significant performance
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 Synchronization if you only want changes since the last synchronization to be queried when synchronizing a directory. Please be aware that when using this option, the user account configured for synchronization must have read access to: TheusNchanged attribute of all users and groups in the directory that need to be synchronized. The objects and attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchronization 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. 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. 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.		This setting determines how your application will compare DNs to determine if they are equal.
Enable Incremental Synchronization if you only want changes since the last synchronization to be queried when synchronizing a directory. A Please be aware that when using this option, the user account configured for synchronization must have read access to: • Theusnchanged attribute of all users and groups in the directory that need to be synchronized. • The objects and attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchronization 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 (second s) 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 60 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.		comparison. This is the default and recommended setting for Active Directory, because Active Directory guarantees the format of DNs.
Increme ntal Synchro nization Please be aware that when using this option, the user account configured for synchronization must have read access to: TheuSNChanged attribute of all users and groups in the directory that need to be synchronized. The objects and attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchro nization Interval (minutes) Synchronization 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 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 60 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.		
 Synchro nization ⚠ Please be aware that when using this option, the user account configured for synchronization must have read access to: TheuSNChanged attribute of all users and groups in the directory that need to be synchronized. The objects and attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchronization 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 (second s) 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 60 seconds. Search Timeout (second 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. 	Increme	
synchronized. The objects and attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchronization 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 (second Important the process of the	Synchro	
If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchro nization Interval (minutes) Synchronization 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 (second s) 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. 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.		
deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory". Synchronization 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 (second s) 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. 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.		The objects and attributes in the Active Directory deleted objects container.
Synchronization 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 (second s) Search Timeout (second Timeout (second)		
nization Interval (minutes) Read Timeout (second s) Search Timeout (second directory server) 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. 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.		This setting is only available if the directory type is set to "Microsoft Active Directory".
Timeout (second s) Search Timeout (second s) S	nization Interval	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
Timeout means there is no limit. The default value is 60 seconds.	Timeout (second	specified time period, the read attempt will be aborted. A value of 0 (zero) means there is no
	Timeout (second	

Connect	This setting affects two actions. The default value is 0.
Timeout (second s)	 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 in our knowledge base. See How to write LDAP search filters.
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 Passwor d 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 toRFC 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

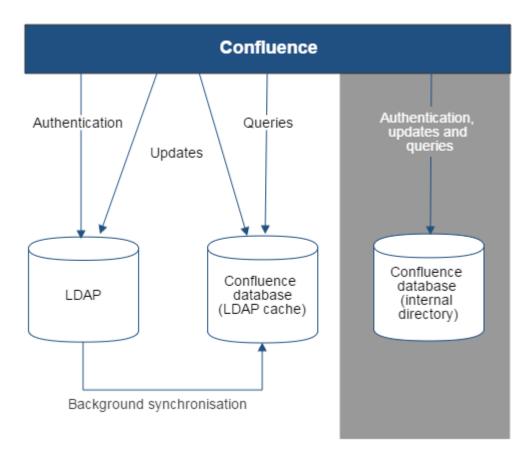


Diagram above: Confluence connecting to an LDAP directory.

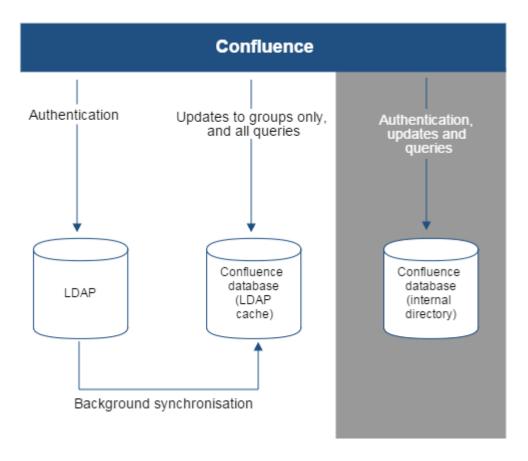


Diagram above: Confluence connecting to an LDAP directory with permissions set to read only and local groups.

Configuring the LDAP Connection Pool

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

Related pages:

- Connecting to an LDAP Directory
- Configuring User Directories

To configure your LDAP connection pool:

- 1. Choose the cog icon , then choose General Configuration
- 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
Maximu m 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 Authentic ation	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.

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

Related pages:

- Connecting to an LDAP Directory
- Configuring User Directories

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'skeystore.

Prerequisites

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

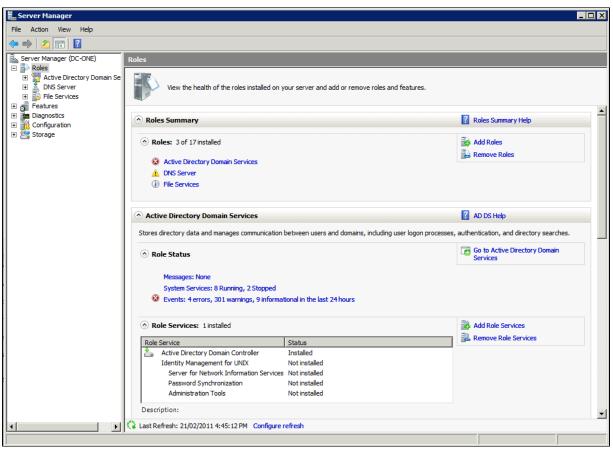
Required Component	Description	
Internet Information Services (IIS)	This is required before you can install Windows Certificate Services.	
Windows Certificate Services	This installs a certification authority (CA) which is used to issue certificates. Step 1, below, explains this process.	
Windows 2000 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

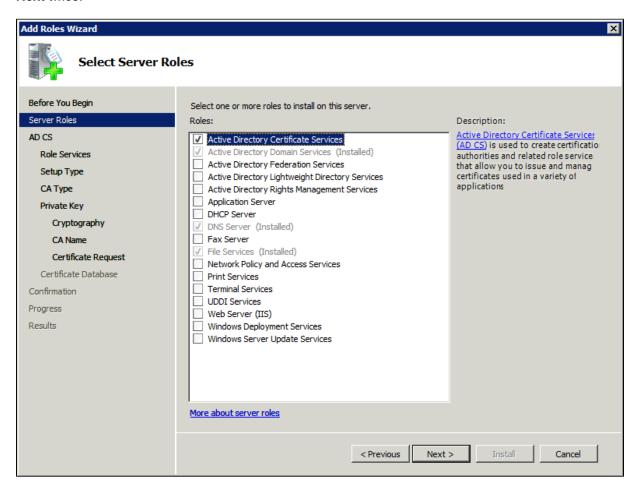
If Certificate Services are already installed, skip to step 2, below. The screenshots below are from Server 2008, but the process is similar for Server 2000 and 2003.

- 1. Log in to your Active Directory server as an administrator.
- 2. Click Start, point to Administrative Tools, and then click Server Manager.

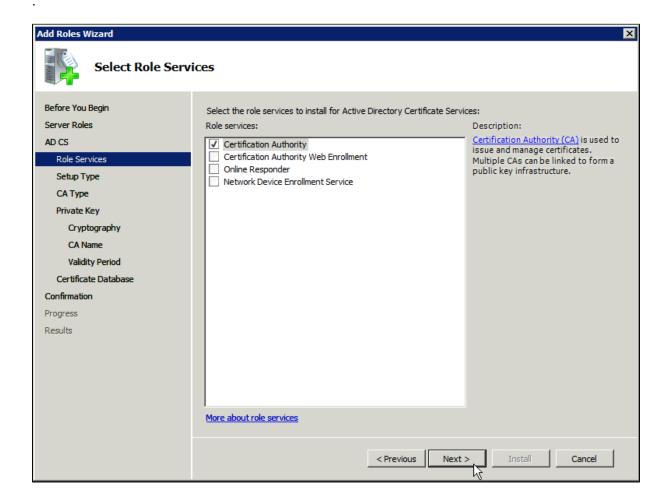
3. In the Roles Summary section, click Add Roles.



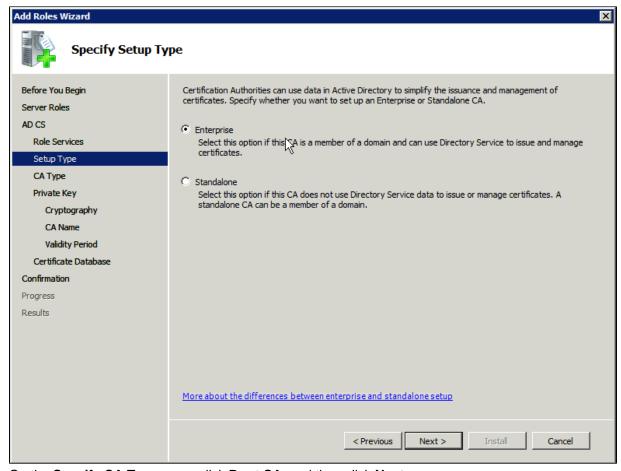
 On the Select Server Roles page, select the Active Directory Certificate Services check box. Click Next twice.



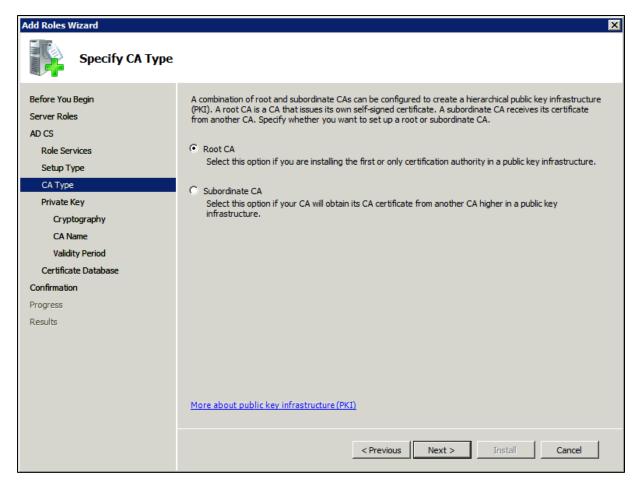
5. On the Select Role Services page, select the Certification Authority check box, and then click Next



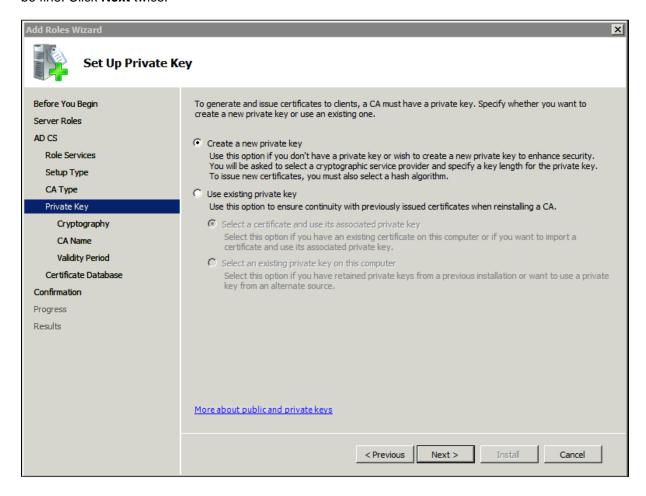
6. On the Specify Setup Type page, click Enterprise, and then click Next.



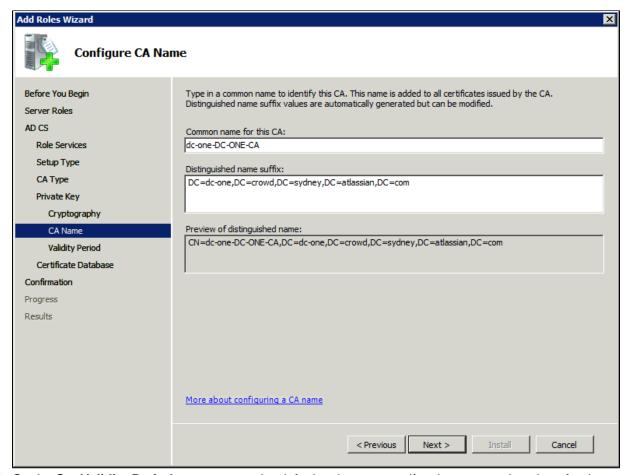
7. On the Specify CA Type page, click Root CA, and then click Next.



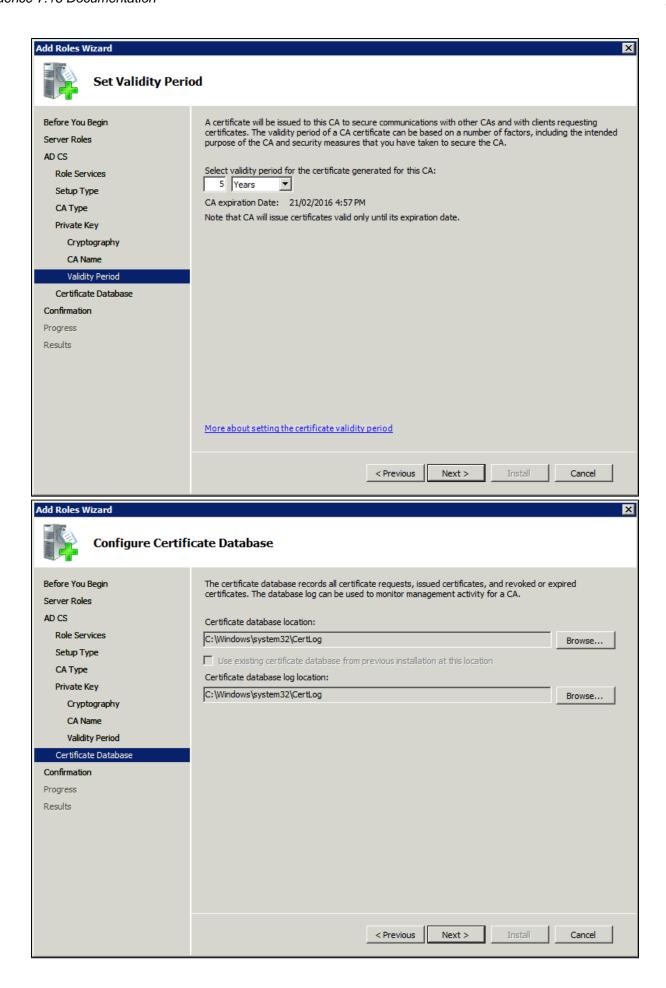
8. On the Set Up Private Key and Configure Cryptography for CA pages, you can configure optional configuration settings, including cryptographic service providers. However, the default values should be fine. Click Next twice.



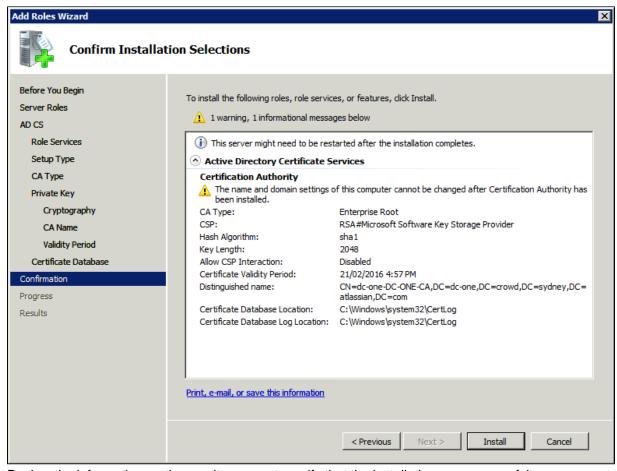
9. In the Common name for this CA box, type the common name of the CA, and then click Next.



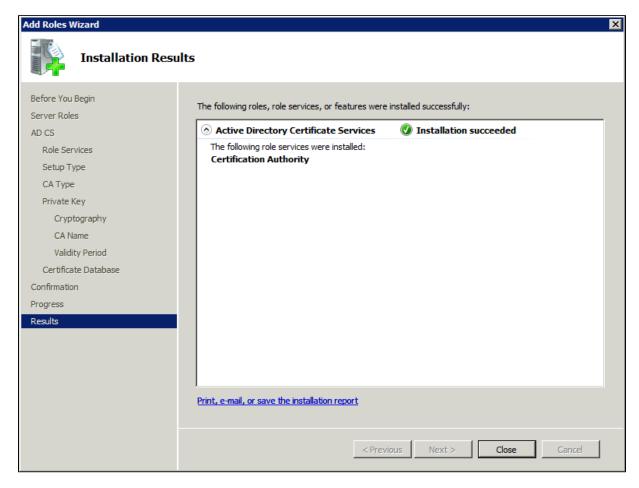
10. On the Set Validity Period page, accept the default values or specify other storage locations for the certificate database and the certificate database log, and then click Next.



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



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



Step 2. Obtain the Server Certificate

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

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

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

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

You might still fail to be authenticated using the certificate file above. In this case, Microsoft'sLDAP over SSL (LDAPS) Certificatepage might help. Note that you need to:

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

Step 3. Import the Server Certificate

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

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

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

Windows

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

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

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

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

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

5. Restart the application to take up the cacerts changes.

6. You may nowchange 'URL' to use LDAP over SSL (i.e.Idaps://<HOSTNAME>:636/) anduse the 'Secure SSL' optionwhen connecting your application to your directory server.

UNIX

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

```
cd $JAVA_HOME
```

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

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

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

- 5. Restart the application to take up the cacerts changes.
- 6. You may nowchange 'URL' to use LDAP over SSL (i.e.ldaps://<HOSTNAME>:636/) anduse the 'Secu re SSL' optionwhen connecting your application to your directory server.

Mac OS X

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

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

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

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

- 3. keytool will prompt you for a password. The default keystore password is changeit.
- 4. When prompted Trust this certificate? [no]: enter yesto 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
```

- 5. Restart the application to take up the cacerts changes.
- 6. You may nowchange 'URL' to use LDAP over SSL (i.e.Idaps://<HOSTNAME>:636/) anduse the 'Secu re SSL' optionwhen connecting your application to your directory server.

Connecting to an Internal Directory with LDAP Authentication

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

Overview

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

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

On this page:

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

Related pages:

Configuring User Directories

Connecting Confluence to an Internal Directory with LDAP Authentication

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

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

For details see Managing Multiple Directories.

7. Add your users and groups in Confluence. See Add and Invite Users and Managing Site-Wide Permissions and Groups .

Server Settings

Setting Description

Name	A descriptive name that will help you to identify the directory. Examples:
	 Internal directory with LDAP Authentication Corporate LDAP for Authentication Only
Director y 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.
Hostna me	The host name of your directory server. Examples: ad.example.com ldap.example.com opends.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.
Userna me	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 synchronized on each subsequent log in. If this box is not checked, the user's login will fail if the user wasn't already manually created in the directory.
	If you check this box the following additional fields will appear on the screen, which are described in more detail below:
	 Default Group Memberships Synchronize Group Memberships User Schema Settings (described in a separate section below)

Update User attribute s on Login

Whenever your users authenticate to the application, their attributes will be automatically updated from the LDAP server into the application. After you select this option, you won't be able to modify or delete your users directly in the application.

- If you need to modify a user, do it on the LDAP server; it will be updated in the application after authenticating.
- If you need to delete a user, do it on the LDAP server, but also in the application. If you delete the user only on the LDAP server, it will be rejected from logging in to the application, but it won't be set as inactive, which will affect your license. You'll need to disable the **Update User attributes on Login** option to delete the user, and then enable it again.

Default Group Member ships

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, authorization 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-administrators, jira-core-users

Synchro nize Group Member ships

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 synchronized 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 ldp.exe which is useful for finding out and configuring the the LDAP structure of your server.

User Name	The attribute field to use when loading the username. Examples:
Attribute	• 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 <i>nested groups</i> . Nested groups simplify permissions by allowing sub-groups to inherit permissions from a parent group.
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
Additiona I 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 **Synchronize 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 **Synchronize 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

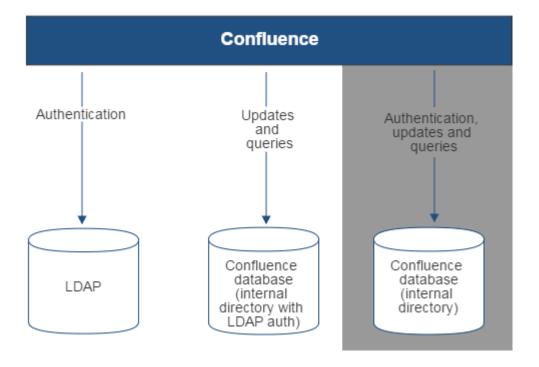


Diagram above: Confluence connecting to an LDAP directory for authentication only.

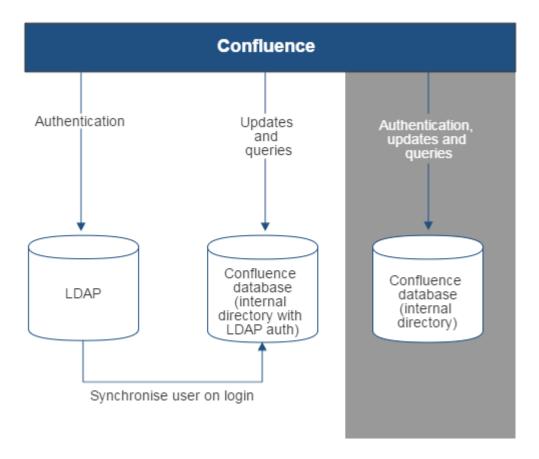


Diagram above: Confluence connecting to an LDAP directory for authentication only, with each user synchronized with the internal directory that is using LDAP authentication when they log in to Confluence.

Connecting to Crowd or Jira for User Management

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

You can't use Jira Cloud for user management.

Connecting Confluence to Crowd for User Management

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

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

On this page:

- Connecting Confluence to Crowd for User Management
- Connecting Confluence to Jira applications for User Management
- Diagrams of Some Possible Configurations
- Troubleshooting

Related pages:

Configuring User Directories



Managing 500+ users across Atlassian products? Find out how easy, scalable and effective it can be with Crowd! Seecentralized user management.

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. Go to General Configuration User directories.
- 3. Add a directory and select type 'Atlassian Crowd'. Enter the settings as described below.
- 4. Save the directory settings.
- 5. Define the **directory order** by clicking the blue up- and down-arrows next to each directory on the '**Us er Directories**' screen. Here is a summary of how the directory order affects the processing:
 - Changes to users and groups will be made only in the first directory where the application has permission to make changes.
 - The order of the directories is the order in which they will be searched for users and groups (by default Confluence aggregates group membership from all directories, so the order does not impact membership itself).

For details see Managing Multiple Directories.

6. 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
Applicati on Name	The name of your application, as recognized by your Crowd server. Note that you will need to define the application in Crowd too, using the Crowd administration Console. See the Crowd documentation on adding an application.
Applicati on 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.

① Note: There is a known issue where the password is not saved in some instances

CONFSERVER-33979 - New JIRA/Crowd password not saved after test GATHERING IMPACT when configuring Confluence to use Jira/Crowd as a external user directory.

Crowd Permissions

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

Advanced Crowd Settings

Setting	Description
Enable Nested Groups	Enable or disable support for nested groups. Before enabling nested groups, please check to see if the user directory or directories in Crowd support nested groups. When nested groups are enabled, you can define a group as a member of another group. If you are using groups to manage permissions, you can create nested groups to allow inheritance of permissions from one group to its sub-groups.
Enable Increme ntal Synchro nization	Enable or disable incremental synchronization. Only changes since the last synchronization will be retrieved when synchronizing a directory. Note that full synchronization is always executed when restarting the application.

Synchro nization Interval (minutes) Synchronization is the process by which the application updates its internal store of user data to agree with the data on the directory server. The application will send a request to your directory server every x minutes, where 'x' is the number specified here. The default value is 60 minutes.

Connecting Confluence to Jira applications for User Management



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

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

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

To connect Confluence to a Jira Server or Data Center application:

- 1. In your Jira application go to User Management > Jira User Server.
 - (For Jira 6.4 and earlier go to your Jira administration screen then **Users > Jira User Server**)
 - Click Add Application.
 - Enter the application name and password that Confluence will use when accessing Jira.
 - Enter the **IP address** or addresses of your Confluence server. Valid values are:
 - A full IP address, e.g. 192.168.10.12.
 - A wildcard IP range, using CIDR notation, e.g. 192.168.10.1/16. For more information, see the introduction to CIDR notation on Wikipedia and RFC 4632.
 - Save the new application.
- 2. Set up the Jira user directory in Confluence:
 - Go to O Seneral Configuration User directories.
 - 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.
 - ADon't change the directory order until you have done the next step or you may accidentally lock yourself out of the Confluence admin console.
- 3. In order to use Confluence, users must be a member of the confluence-users group or have Confluence 'can use' permission. Follow these steps to configure your Confluence groups in your JIRA application:
 - a. Add the confluence-users and confluence-administrators groups in your JIRA application.
 - b. Add your own username as a member of both of the above groups.
 - c. Choose one of the following methods to give your existing JIRA users access to Confluence:
 - Option 1: In your JIRA application, find the groups that the relevant users belong to. Add the groups as members of one or both of the above Confluence groups.
 - Option 2: Log in to Confluence using your JIRA account and go to the Confluence Administration Console. Click 'Global Permissions' and assign the 'can use' permission to the relevant JIRA groups.
- 4. In Confluence you can now 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 seeManaging Multiple Directories.



⚠ 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 in the list of directory servers. Examples:
	• Jira Software • My Company Jira
Server URL	The web address of your Jira server. Examples: • http://www.example.com:8080 • http://jira.example.com
Applicati on 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.
Applicati on 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.
Synchron ization Interval (minutes)	Synchronization 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

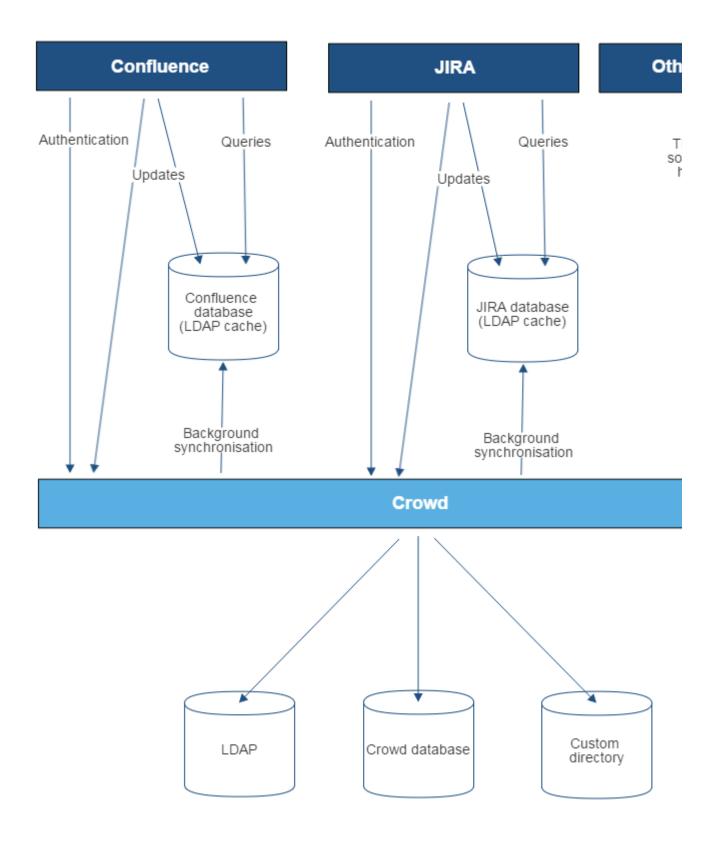


Diagram above: Confluence, JIRA and other applications connecting to Crowd for user management.

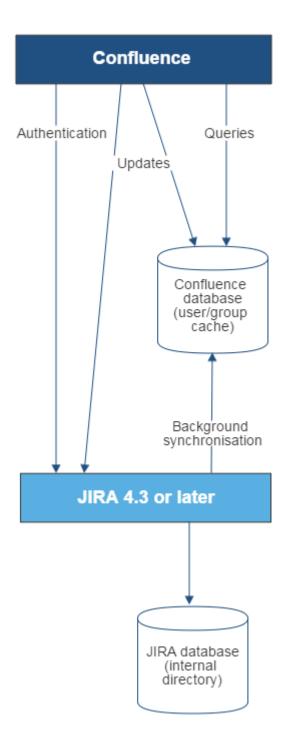


Diagram above: Confluence connecting to JIRA for user management.

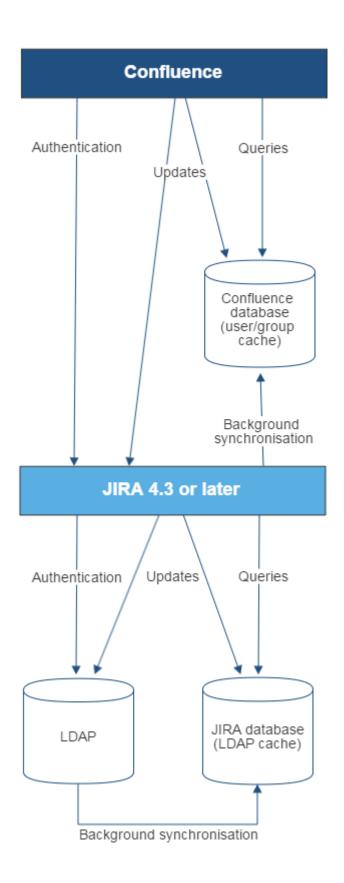


Diagram above: Confluence connecting to JIRA for user management, with JIRA in turn connecting to LDAP.

Troubleshooting

Below are some error messages you may encounter. If you run into problems, you should turnon WARN logging for the relevant class. SeeConfiguring Logging.

Error	Message	Cause
error.jirabaseurl. connection. refused	Connection refused. Check if an instance of Jira is running on the given url	 This may be because: 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 a Jira application.	A runtime exception has occured. Refer to the Confluence log files for further troubleshooting information.

Reverting from Crowd or Jira applications to Internal **User Management**

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



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 application 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 theconfluenceusers 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
- 5. Disable the 'Atlassian Crowd' directory.
- 6. Manually add the required users and groups in Confluence. They will be added to the internal directory, because you have moved it to the top of the list of directories.
 - If you have assigned Confluence permissions to a group which exists in your Jira application, you must create a group in Confluence with the same name.
 - If a user who exists in your Jira application has created content or has had permissions assigned to them in Confluence, you must also create that user in Confluence.
- 7. Add the users to the required groups.

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



This method is not officially supported. The Atlassian Support team won't be able to assist you with this process.

We strongly recommend trying this in a test environment, and then making a full backup of your database before deciding to deploy the change in your production environment.

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

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

Step 1. Create Backups

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

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

Step 2. Replace Confluence User Management

Use the SQL below to move groups and users from your External Application to Confluence by transferring table content. The SQL provided is specific to MySQL and must be modifed 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>';
```

Find and remove duplicate users who belong to the same group in multiple directories

To make sure you don't introduce duplicates in the next step, when you move groups to Confluence, use the following SQL query to locate any users that belong to a group with the same name in both your external directory and internal Confluence directory.

1. Run the following command to find any users with the same name, that belong to the same group across different directories:

```
SELECT count(*), a.user_name, c.group_name from cwd_user a
join cwd_membership b on b.child_user_id = a.id
join cwd_group c on c.id = b.parent_id group by 2,3 having count(*)>1
```

Make a note of each of the usernames and groups returned. You'll need this in the next step.

- In your external directory, remove the users from their respective groups. Their membership will still be retained in the Confluence internal directory.
- 3. Run the SQL query above again. Once it returns no results, you can move to the next step.

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.

```
\label{local_problem} \begin{tabular}{ll} $\tt update \ cwd\_group \ set \ directory\_id=<External \ Application \ ID>; \end{tabular}
```

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

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

Managing Multiple Directories

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

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



Managing 500+ users across Atlassian products?

Find out how easy, scalable and effective it can be with Crowd!

Seecentralized user management.

On this page:

- Overview
- Configuring the Directory Order
- Effect of Directory Order
 - Login
 - Permissions
 - Updating Users and groups

Overview

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

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

Configuring the Directory Order

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

Directory Name	Туре	Order
Confluence Internal Directory	Internal	₩.
OpenLDAP	OpenLDAP (Read-Write)	1

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.
- Before you move an external directory above Confluence's internal directory, make sure you (and your admin users) are members of a group called confluence-administrators your external directory or you may accidentally lock yourself out of the Confluence admin console.

Effect of Directory Order

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

Login

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

Permissions

Aggregating membership (default)

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

Example:

- You have connected two directories: The Customers directory and the Partners directory.
- The Customers directory is first in the directory order.
- A usernamejsmithexists in both the Customers directory and the Partners directory.
- The userjsmithis a member of groupglin the Customers directory and groupglin the Partners directory.
- The userjsmithwill have permissions based on membership of bothglandg2regardless of the directory order.

For administrators upgrading to Confluence 5.7 or later:

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



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

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

Confluence User Directories



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



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

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

To Confluence his group membership looks like this:



Confluence group membership confluence-users developers sydney

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

Non-aggregating membership

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

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

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

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

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

Example:

- You have connected two directories: The Customers directory and the Partners directory.
- The Customers directory is first in the directory order.
- A username jsmithexists in both the Customers directory and the Partners directory.
- The userjsmithis a member of groupglin the Customers directory and groupglin the Partners directory.
- The userjsmithwill have permissions based on membership ofg1only, notg2.

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.

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*. Nested groups simplify permissions by allowing sub-groups to inherit permissions from a parent group.

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 'Use r 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.
- You can't edit the directory you are currently logged in via. This means that in most cases you need to log in with an administrator account stored in the internal directory.

On this page:

- Enabling Nested Groups
- Effect of Nested Groups
 - Login
 - Permissions
 - Viewing lists of group members
 - Adding and updating group membership
- Examples
 - Example 1: User is member of subgroup
 - Example 2: Sub-groups as members of the jira-developers group
- Notes

Related pages:

Configuring User Directories

Effect of Nested Groups

This section explains how nested groups affect logging in, permissions, and viewing and updating users and groups.

Login

When a user logs in, they can access the application if they belong to an authorized group or any of its subgroups.

Permissions

The user can access 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 can't view or edit the nested groups themselves, or see that one group is a member of another group.

Adding and updating group membership

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 of groups in the flattened list, the user is removed from the top group.
- Otherwise, you 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

Imagine 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 the groups **engineering-group**and **techwriters-group**. We want to grant both groups developer-level access to the JIRA. We will have a group called **jira-developers** that has developer-level access.

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

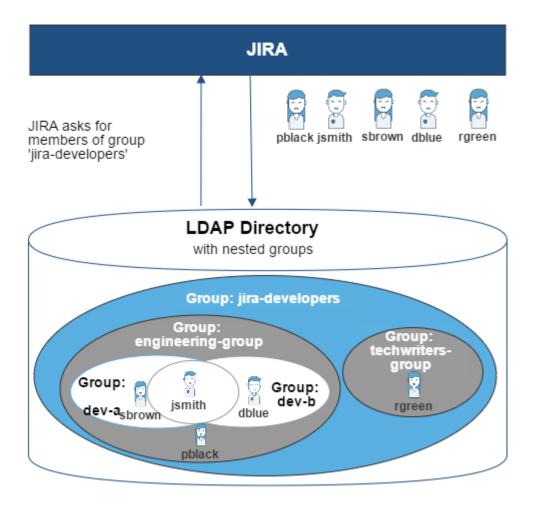
Group memberships are now:

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

When the JIRA application requests a list of users in the jira-developers group, it receives the following list:

- pblack
- jsmith
- sbrown
- dblue
- rgreen

Diagram: Sub-groups as members of the jira-developers 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 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 Onemight 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

Synchronizing 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 synchronization 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 Synchronize
- Manually Synchronizing the Cache
- Configuring the Synchronization Interval
- Unsynced users

Related pages:

Configuring User Directories

Affected Directory Types

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

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

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

- 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 synchronization 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 synchronization task has copied that user's details into the cache.
- A periodic synchronization task will run to update the database with any changes made to the
 external directory. The default synchronization interval, or polling interval, is one hour (60 minutes).
 You can change the synchronization interval on the directory configuration screen.
 - Note for Confluence Data Center: The sync will take place on a single node of the cluster to update the database. This may make it seem like automatic synchronization will not be happening, but the task is assigned to one of the nodes.
- You can manually synchronize 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.

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

Manually Synchronizing the Cache

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

Screen snippet: User directories, showing information about synchronization

OpenLDAP	OpenLDAP (Read-Write)	4 4	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 Synchronization Interval

Note: The option to configure the synchronization 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 '**Synchronization Interval**' on the directory configuration screen. The synchronization interval is the period of time to wait between requests for updates from the directory server.

The length you choose for your synchronization 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 synchronize more frequently, then your data will be more up to date. The downside of synchronizing 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.

Unsynced users

To view users who have previously been synchronized with Confluence, but were not present in the last directory sync, go to User management Unsynced from Directory.

Users may appear in the Unsynced from Directory tab be due to a problem with your last sync, or because the user has been intentionally removed from the external directory (for example because they've left your organisation).

If a user who has created content is removed from an external directory, and a new account is created with the same username, that username will be associated with the original user's content. This is intentional, to ensure that if a directory sync problem occurs, users are correctly re-associated with their own content.

If the user was intentionally unsynced, administrators can choose to:

- Leave the unsynced account as it is. The person's username will appear on any content or comments they've created.
- Delete the account from the Unsynced from Directory tab, which then replaces the username with an anonymous alias. This final deletion step is usually only required if you've received a formal erasure request.

SeeDelete or Disable Users for more information. Don't assume that because a user appears in the unsynced users list, that they are to be deleted from Confluence.

You may see a user in the Unsynced from Directory tab with the username 'exporter'. This account is used when creating the demonstration space when you first install Confluence, and can be included when importing a Cloud site. You can safely ignore this unsynced account.

Diagrams of Possible Configurations for User Management

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

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

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

On this page:

- Confluence Internal Directory
- Confluence with Read/Write Connection to LDAP
- Confluence with Read-Only Connection to LDAP, with Local Groups
- Confluence Internal Directory with LDAP Authentication
- Confluence with LDAP Authentication, Copy Users on First Login
- Confluence Connecting to Jira
- Confluence Connecting to Jira and Jira Connecting to LDAP
- Confluence and Jira Connecting to Crowd

Related pages:

• Configuring User Directories

Confluence Internal Directory

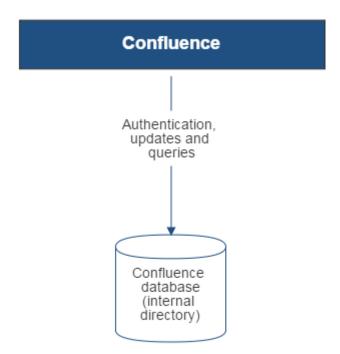


Diagram above: Confluence using its internal directory for user management.

Confluence with Read/Write Connection to LDAP

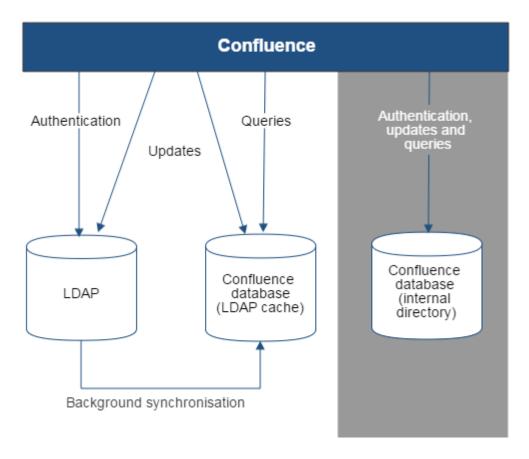


Diagram above: Confluence connecting to an LDAP directory.

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

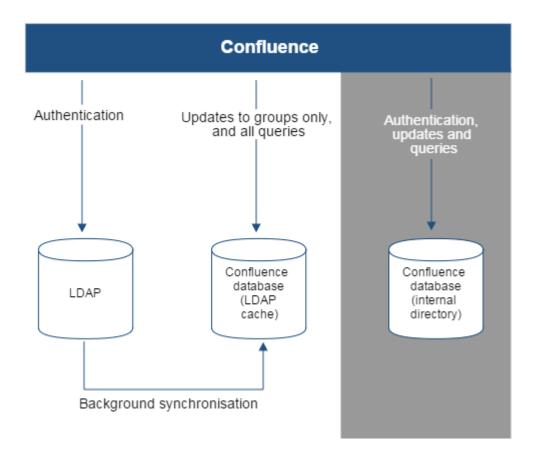


Diagram above: Confluence connecting to an LDAP directory with permissions set to read only and local groups.

Confluence Internal Directory with LDAP Authentication

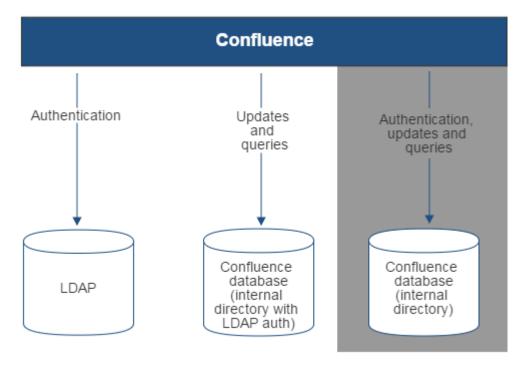


Diagram above: Confluence connecting to an LDAP directory for authentication only.

Confluence with LDAP Authentication, Copy Users on First Login

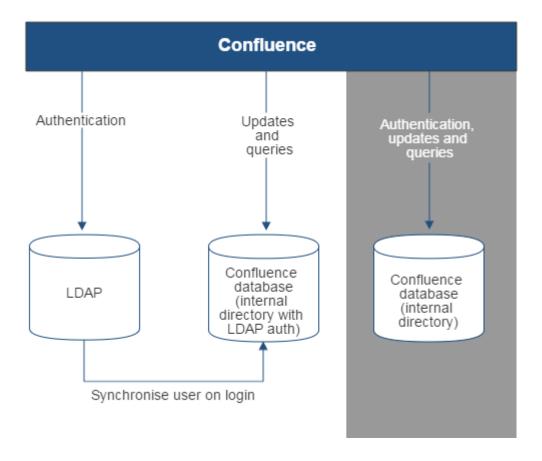


Diagram above: Confluence connecting to an LDAP directory for authentication only, with each user synchronized with the internal directory that is using LDAP authentication when they log in to Confluence.

Confluence Connecting to Jira

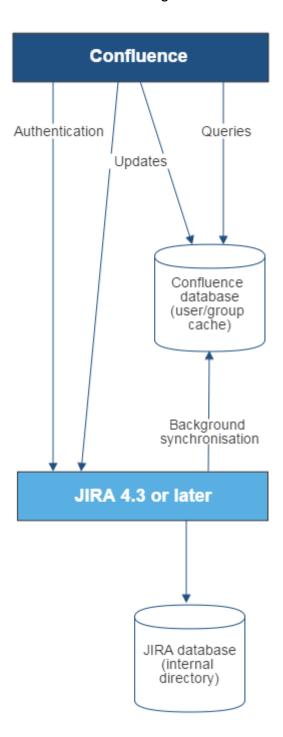


Diagram above: Confluence connecting to JIRA for user management.

Confluence Connecting to Jira and Jira Connecting to LDAP

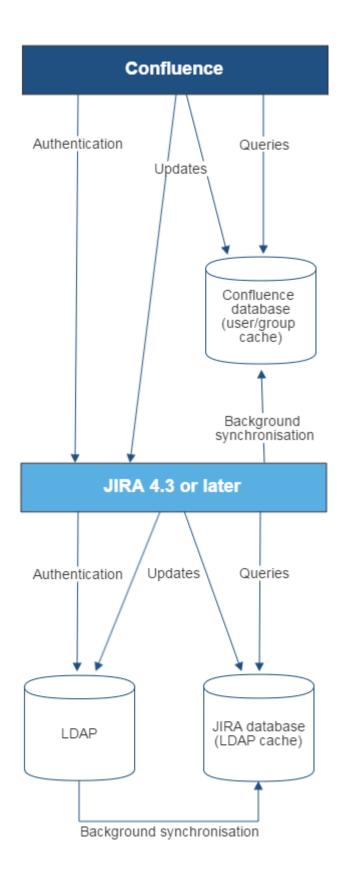


Diagram above: Confluence connecting to JIRA for user management, with JIRA in turn connecting to LDAP.

Confluence and Jira Connecting to Crowd

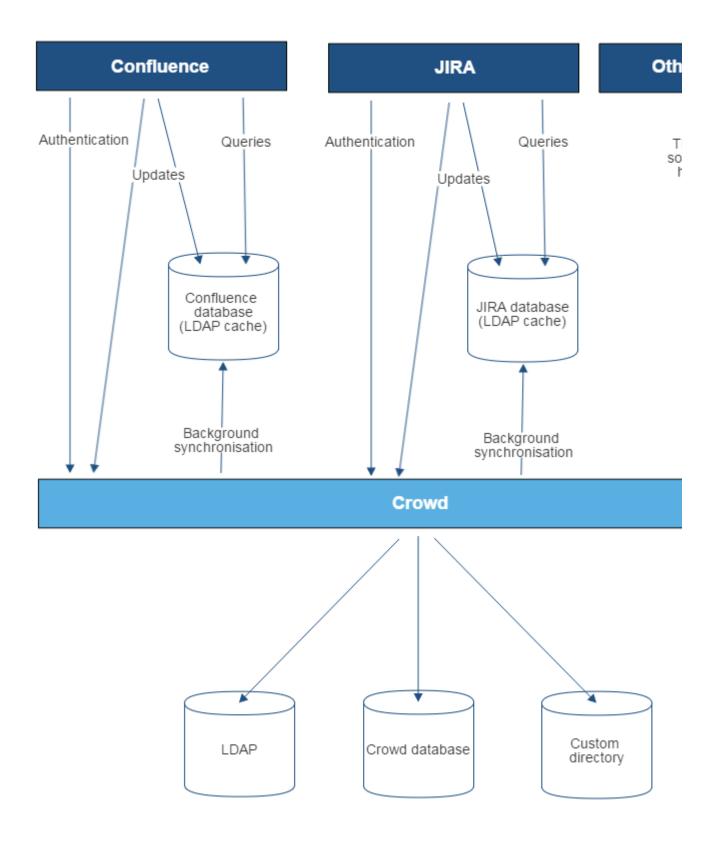


Diagram above: Confluence, JIRA and other applications connecting to Crowd for user management.

User Management Limitations and Recommendations

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

General Recommendations

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

Be careful when deleting users in remote directories.

If you are connecting to an LDAP directory, a Crowd directory or a Jira directory, please take care when deleting users from the remote directory. If you delete a user that is associated with data in Confluence, this will cause problems in Confluence.

If a user who has created content is deleted from an external directory, and an account is then re-created with the same username, it will automatically be reassociated with that content. This is intentional, so that if a directory sync problem occurs, users are correctly re-associated with their content.

Avoid hash, slash and question characters in usernames

There is a known issue where users with #, ? or / in their username cannot create spaces. See

CONFSERVER-43494 GATHERING IMPACT

and CONFSERVER-13479 GATHERING IMPACT

for more information.

On this page:

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

Related pages:

- Connecting to an LDAP Directory
- Connecting to Crowd or Jira for User Management
- Configuring User Directories

Recommendations for Connecting to LDAP

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

Optimal Number of Users and Groups in your LDAP Directory

The connection to your LDAP directory provides powerful and flexible support for connecting to, configuring and managing LDAP directory servers. To achieve optimal performance, a background synchronization 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 synchronization 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 synchronization task.	

Our Test Results

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

We found that the initial synchronization took about 5 minutes. Subsequent synchronizations 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 synchronization 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 synchronizations 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 synchronization process caches data in the database, the performance of your database will affect the performance of the synchronization.
- **JVM heap size.** If your heap size is too small for your userbase, you may experience heavy garbage collection during the synchronization process which could in turn slow down the synchronization.

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 synchronizes with Active Directory (AD), the synchronization task requests only the changes from the LDAP server rather than the entire user base. This optimizes the synchronization process and gives much faster performance on the second and subsequent requests.

On the other hand, this synchronization 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 synchronize the directory cache after you have made the changes to ensure cache consistency.

- 2. Synchronizing 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 synchronization. (You can of course define multiple different directories, each pointing to its own respective AD server.)
- 3. Synchronizing with AD servers behind a load balancer is not supported. As with synchronizing 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 synchronize 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 synchronization 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 synchronization task relies on the uSNChanged attribute to detect changes, and so must be in the appropriate AD security groups to see this attribute for all LDAP objects in the subtree.

Recommendations for Connecting to Jira for User Management

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

Single Sign-On Across Multiple Applications is Not Supported

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

Custom Application Connectors are Not Supported

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

Custom Directories are Not Supported

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

Load on your JIRA instance

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

JIRA Cloud applications not supported

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

Recommendations

Your environment	Recommendation
 If all the following are true: Your JIRA application is not under high load. You want to share user and group management across just a few applications, such as one JIRA Software server and one Confluence server, or two JIRA servers. 	Your environment meets the optimal requirements for using a JIRA application for user management.

- You do not need single sign-on (SSO) between your JIRA application and Confluence, or between two JIRA servers.
- You do not have custom application connectors. Or, if you do have them, you are happy to convert them to use the new REST API.
- You are happy to shut down all your servers when you need to upgrade your JIRA application.

If **one or more** of the following are true:

- If your JIRA application is already under high load.
- You want to share user and group management across more than 5 applications.
- You need single sign-on (SSO) across multiple applications.
- You have custom applications integrated via the Crowd SOAP API, and you cannot convert them to use the new REST API.
- You are not happy to shut down all your servers when you need to upgrade JIRA.

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.

Requesting Support for External User Management

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

The cause of such problems may be:

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

On this page:

- Troubleshooting the Connection to your External User Directory
- Problems During Initial Setup
- Complex Authentication or Performance Problems

Related pages:

- Troubleshooting Problems and Requesting Technical Support
- Configuring User Directories

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
- 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 knowedge 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 ser aph-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.atlassian.seraph=DEBUG
log4j.logger.com.opensymphony.user=DEBUG
```

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

Disabling the Built-In User Management

In some circumstances you may want to disable Confluence's built in user management, and delegate all user management to an external application, such as Jira Software or Jira Service Management. You can disable internal user management by turning on Confluence's **External User Management** setting. You'll need to be asy stem administratorto do this.

You might disable Confluence's internal user management:

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

To disable management of users and groups within Confluence:

- 2. Click Edit.
- 3. Select the External user managementcheckbox then Save your change.

Note: If you turn on External user management:

- You will not be able to add users or groups in Confluence.
- You will not be able to edit user details (full name and email) of users in Confluence Internal Directory
- You will not be able to use public signup in your site.
- The Forgot Password link will not appear on the Confluence login page.
- Users will not be able to reset their password in Confluence.

Single sign-on for Confluence Data Center

We provide the functionality for Confluence Data Center to connect to your preferred identity provider (IdP) so that you can provide your users with a single sign-on (SSO) experience.

On this page:		

This only handles authentication. Application access and any required authorizations, such as ensuring that users belong to the appropriate groups/roles and have the necessary permissions, should be configured in the user directory and/or the application itself.

The way you configure SSO depends on the protocol your IdP uses:

- For SAML based identity providers, seeSAML single sign-on for Atlassian Data Center applications
- For OpenID based identity providers, seeOpenID Connect for Atlassian Data Center applications
- For Atlassian Crowd, seeCrowd SSO 2.0
- If you need to configure multiple different IdPs seeUsing multiple identity providers

Looking for a cross-domain SSO solution?

Atlassian Crowd 3.4, with its Crowd SSO 2.0 feature, offers one solution for Server, Data Center, and Cloud applications and setting it up takes only minutes.

Are you are ready for the change? See Crowd SSO 2.0

Managing System and Marketplace Apps

An app is a separately installed component that extends the basic Confluence functionality.

Not to be confused with the Confluence mobile app that users install on their own device, these apps are installed by a Confluence admin, and act like an extension to Confluence. They are also known'plugins' or 'addons'.

There are two main types of apps:

- System apps these are bundled with Confluence and provide core functionality
- User installed apps these are usually downloaded from The Marketplace and may have been created by Atlassian or by a third party developer.

For information about developing your own apps for Confluence, see the Confluence Server and Data Center Developer documentation.

About the Universal Plugin Manager

System and Marketplace apps are managed via the Universal Plugin Manager (known as the UPM). The UPM can be found in most Atlassian applications, and provides a consistent experience for administering apps. To visit the UPM, go to Manage apps in the Confluence header.

The UPM allows you to:

- Discover and install new apps from the Atlassian Marketplace.
- Install or remove apps.
- Configure app settings.
- Enable or disable apps and their component modules.
- Confirm app compatibility before upgrading Confluence.

You'll need Confluence Administrator permissions to access the UPM.

SeeRequest Marketplace Appsfor information on how users can find and request add-ons.

SeetheUniversal Plugin Manager documentation for more information on using the UPM.

Disable and uninstall apps

Youcan disable or unsubscribe from user installedappsthat are no longer being used on your site. See Disabling and enabling apps to find out how to do this.

Once the app is disabled, its features are immediately unavailable. If the app included macros, pages that contained those macros will show an'unknown macro' error. To avoid this, you can check which macros are being used on your site before disabling an app by checking the macro usage statistics.

Go to O S General Configuration > Macro Usage.

Writing User Macros

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

User macros are created and managed within Confluence itself, you do not need to develop an app (plugin). You will need some coding skills though.

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

On this page:

- Create a User Macro
- Edit a user macro
- Delete a user macro
- Best practices
- Example user macros

Related Pages:

User Macro Module (Developer documentation)

Create a User Macro

To add a new user macro:

- 1. Go to O Seneral Configuration Suser Macros
- 2. ChooseCreate a User Macro
- 3. Enter the macro details (see table below)
- 4. ClickAdd

Macro details field	Description
Macro name	This is the name of the macro, as it appears in the code.
Visibility	 This controls who can see this macro in the macro browser or auto-complete. Options are: Visible to all users Visible only to system administrators Note that if you select Visible only to system administrators, users will still see the output of the macro on a page, and the macro placeholder will still be visible when a user edits a page. It is only hidden in the macro browser and autocomplete. All macro information is discoverable, including the macro title, description, parameter names and other metadata. Do not include confidential data anywhere in the definition of a user macro, even if it is marked as visible only to system administrators.
Macro Title	This is the title that will appear in the macro browser and auto-complete.
Descrip tion	This is the description that will appear in the macro browser. The macro browser's search will pick up matches in both the title and description.
Categor ies	Select one or more macro browser categories for your macro to appear in.
Icon URL	Enter an absolute URL (for example http://mysite.com/mypath/status.png) or path relative to the Confluence base URL (for example /images/icons/macrobrowser/status.png) if you want the macro browser to display an icon for your macro.
Docum entation URL	If you have documentation for your macro, enter the URL here.

Macro **Body Process** ing

Specify how Confluence should process the body before passing it to your macro.

The macro body is the content that is displayed on a Confluence page. If your macro has a body, any body content that the user enters will be available to the macro in the \$body variable.

Options for processing the macro body include:

No macro body

Select this option if your macro does not have a body.

Escaped

Confluence will add escape characters to the HTML markup in the macro body. Use this if you want to show actual HTML markup in the rendered page. For example, if the body is

t >Hello World it will render as Hello World.

Unrendered

HTML in the body will be processed within the template before being output. Ensure that HTML is ultimately output by the template.

Rendered

Confluence will recognize HTML in the macro body, and render it appropriately. For example, if the body is Hello World it will render as Hello World.

Template This is where you write the code that determines what the macro should do.

- Use HTML and Confluence-specific XML elements in the macro template.
- You can use the Velocity templating language. Here is more information on the Velocity
- If your macro has a body, your template can refer to the macro body text by specifying '\$bo
- Each parameter variable you use must have a matching metadata definition. Use@paramto 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 specifed in the@parammetadata definition.
- Use@noparamsif your macro does not accept parameters.

See User Macro Template Syntax for more information and examples.

① 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 Use r Macro plugin module. If you want to create more complex, programmatic macros in Confluence, you may need to write a Macro plugin.

Edit a user macro

To edit a user macro:

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

Delete a user macro

To delete a user macro:

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

Before deleting a user macro, you shouldsearchfor all occurrences of the macro in pages and blog posts. Users will see an 'unknown macro' error if you delete a user macro that is still in use on a page.

Best practices

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

Add a descriptive header to your macro template

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

```
## Macro title: My macro name
## Macro has a body: Y or N
## Body processing: Selected body processing option
## Output: Selected output option
##
##
## Developed by: My Name
## Date created: dd/mm/yyyy
## Confluence version: Version it was developed for
## Installed by: My Name
## Short description of what the macro does
```

Expose your parameters in the macro browser

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

Supply default values for macro parameters

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

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

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

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

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

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

Consider security implications

We recommend thoroughly testing your user macro with a number of permission scenarios, such as restricted pages and space permissions to avoid inadvertently displaying content that a user has no permission to see. SeeUser Macro Template Syntax for more information.

Example user macros

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

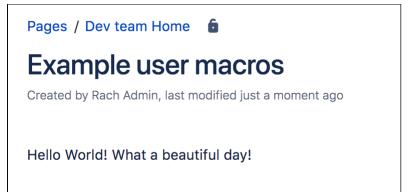
Field	Value		
Macro name	helloworld		
Visibility	Visible to all users in the Macro Browser		
Macro Title	Hello World		
Description	Displays "Hello World" and the macro body.		
Categories	Confluence Content		
Icon URL	You can leave this field blank		
Documentation URL	You can leave this field blank		
Macro body processing	Rendered		
Template	Enter the code below in the template field - this example will print the text straight onto the page.		
	## @noparams Hello World! \$body		
	If you wanted the text to appear in a panel you could include the relevant AUI message class as shown here.		
	<pre>## @noparams <div class="aui-message closeable"> Hello World! \$body </div></pre>		

Using the 'Hello World' macro on a page

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



The result is:

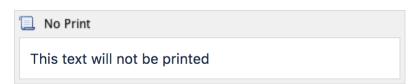


This example demonstrates how to create a user macro that can contain text that is visible when viewing a page, but does not print.

Field	Value	
Macro name	noprint	
Visibility	Visible to all users in the Macro Browser	
Macro Title	No Print	
Description	Hides text from printed output.	
Categories	Confluence Content	
Icon URL	You can leave this field blank	
Documentation URL	You can leave this field blank	
Macro body processing	Rendered	
Template	Enter the code below in the template field.	
	## @noparams <div class="noprint">\$body</div>	

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 recognize the NoPrint macro

SeeAdvanced PDF Export Customizations.

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

Field	Value
Macro name	stylish
Visibility	Visible to all users in the Macro Browser
Macro Title	Stylish
Descripti on	Applies colour and size to text.
Categories	Confluence Content
Icon URL	You can leave this field blank
Documen tation URL	You can leave this field blank
Macro body processi ng	Rendered
Template	Enter the code below in the template field. If your macro requires more than one parameter, you can use variables \$param0 to \$param9 to represent them.
	<pre>## @param 0:title=colour type=string ## @param 1:title=size type=string \$body</pre>
	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.</x></x>
	<pre>## @param Colour:title=colour type=string ## @param Size:title=size type=string \$body</pre>

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

```
(Title)
```

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

Field	Value
Macro name	formpanel
Visibility	Visible to all users in the Macro Browser
Macro Title	Formatted Panel
Description	Creates a panel preformatted with specific colors

Categories	Formatting
Icon URL	You can leave this field blank
Documentation URL	You can leave this field blank
Macro body processing	Escaped
Template	Enter the code below in the template field. See below for a more detailed explanation of the code below. ## @param Title:title=Title type=string desc=Title <ac:structured-macro ac:name="panel"></ac:structured-macro>

Explanation of the code in the macro template

Below is a breakdown of the user macro template code.

Item	Description
## @param Title: title=Title type=st	@param defines the metadata for your macro parameters.
ring desc=Title	@param Title
	This parameter is called "Title".
	title=Title
	defines the parameter title that will appear in the macro browser as "Title".
	type=string
	defines the field type for the parameter as a text field.
	desc=Title
	defines the description of the parameter in the macro browser.
<ac:structured- macro ac:name="</ac:structured- 	This calls the Confluence Panel macro.
panel">	The easiest way to find out the code name of a Confluence macro by viewing the Storage Format of a page containing the macro. You'll need Confluence Administrator permissions to view the storage format.

<ac:parameter ac:<br="">name="titleBGColor" >#ccc<th>Sets the parameters for the macro: the background color, border style, border color, border width and title color.</th></ac:parameter>	Sets the parameters for the macro: the background color, border style, border color, border width and title color.
parameter> <ac:parameter ac:="" name="borderStyle">solid <ac:parameter ac:="" name="borderColor">#6699CC <ac:parameter ac:="" name="borderWidth">2</ac:parameter> <ac:parameter> <ac:parameter> <ac:parameter> cac:parameter ac: name="titleColor" >#000000</ac:parameter></ac:parameter></ac:parameter></ac:parameter></ac:parameter>	To discover the names of the parameters for a Confluence macro, view the storage format as described above.
<ac:parameter ac:<br="">name="title">\$! paramTitleparameter></ac:parameter>	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.
<ac:rich-text- body>\$bodyrich-text-body></ac:rich-text- 	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.
	This command marks the end of the macro.

O Do more with Confluence

Not keen to write your own macro? There are a ton of free and paid macros available in the Atlassian Marketplace. Here are some of our most popular:

- Numbered Headings: Automatically number headings for easy navigation and documentation
- HideElements for Confluence: Hide several Confluence page elements e.g. title, comments, buttons - with just one click
- Composition Tabs & Page Layout: Bring your content to life tabs, highlights, instant focus, menus and expandable sections

User Macro Template Syntax

SeeWriting User Macros for an introduction to writing a user macro.

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

Accessing your macro's body

Use the \$bodyobject within your user macro template to access the content passed to your macro in the macro body.

The \$bodyobject 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 calledhellow orld.

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 behavior on a Confluence page.

How your macro parameters are used on a Confluence page

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

Defining the parameters

A parameter definition in the template contains:

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

Format:

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

On this page:

- Accessing your macro's body
- Using parameters in your user macro
- Objects available to your macro
- Rendering HTML with variables
- Controlling parameter appearance in the editor placeholder

Related pages:

Writing User Macros

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:
	## @param colour:title=Colour type=enum enumValues=Grey,Red,Yellow,Green
	Note about i18n: Confluence does not support internationalization of the enum values. The value the user sees is the one passed to the macro as the parameter value, with the capitalization given. In this case 'Grey', 'Red', etc.
string	A text field. This is the default type. Example with a required field:
	## @param status:title=Status type=string required=true desc=Status to display
confluence-	Offers a control allowing the user to search for a page or blog post. Example:
Content	## @param page:title=Page type=confluence-content required=true desc=Select a page do use
username	Search for user.
	## @param user:title=Username type=username desc=Select username to display
spacekey	Offers a list of spaces for selection. Passes the space key to the macro. Example:
	## @param space:title=Space type=spacekey
date	Confluence accepts this type, but currently treats it in the same way as 'string'. Example:
	## @param fromDate:title=From Date type=date desc=Date to start from. Format: dd/mm /YYYY
	Note about dates: 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 treats it in the same way as 'string'. Example with a default value:
	## @param numPosts:title=Number of Posts type=int default=15 desc=Number of posts to display
percentage	Confluence accepts this type, but treats it in the same way as 'string'. Example:
	## @param pcent:title=Percentage type=percentage desc=Number of posts to display

Using the parameters in your macro code

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

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

Using no parameters

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

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

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

@noparams

Objects available to your macro

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

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

Macros can also access objects available in the default Velocity context, as described in the developer documentation.



Security consideration

When creating a User Macro you should avoid using \$content.getChildren() or \$content. getDescendants() as these methods will list all pages, regardless of page restrictions or space permissions. This may lead to page viewers seeing pages that they do not have permission to see.

We also recommend thoroughly testing your user macro with a number of permission scenarios, such as restricted pages and space permissions.

Rendering HTML with variables

If HTML is rendered with variable assignment, the variable name needs to end with "Html". This will render HTML instead of escaping it.

Example: \$outputHtmlinstead of \$output

Controlling parameter appearance in the editor placeholder

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

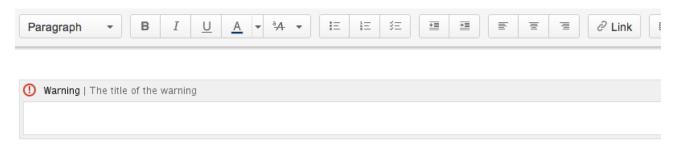
By default as many parameters as can fit will be displayed in the placeholder, as shown here:



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 behavior 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).

Customizing your Confluence Site

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

For guidelines on customizations at a personal and space level, see Your User Profile or Customize your Space.

We've documented the customizations under two broad headings:

- You can change the **appearance** of Confluence by customizing the dashboard, adjusting the colors, adding a site logo, and more. See Changing the Look and Feel of Confluence.
- You can determine the **default behavior** 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
 Behavior and Content in Confluence.

Related pages:

- Integrating Confluence with Other Applications
- Tracking Customizations Made to your Confluence Installation
- Confluence administrator's guide

Changing the Look and Feel of Confluence

You can change the appearance, or look and feel of Confluence for the whole site (globally) or for individual spaces.

Changes you make to the whole site will also apply to all spaces that are inheriting the global look and feel. Users with space administrator permissions can further customize the appearance of a spaceand override the global look and feel for that space. SeeCustomize your Spacefor more.

Related pages:

- Administering Site Templates
- Working With Decorator Macros
- Customizing a Specific Page
- Upgrading Customized Site and Space Layouts

Ways to customize the look and feel of your site:

- Add your own site logo. See Changing the Site Logo.
- Change the color scheme of the user interface. See Customizing Color Schemes.
- Usethemesfor advanced layout customization. SeeWorking with Themes.
- 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. SeeCustomizing Site and Space Layouts.

Customizing the Confluence Dashboard

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

Editing the site welcome message

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

SeeEditing the Site Welcome Messageto find out how to add announcements, useful links, images, macros and more.

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

Using a page as the site landing page

If you want more control, you can choose to use an ordinary Confluence page as your site landing page, instead of sending people to the dashboard. SeeConfiguring the Site Home Pageto find out more.

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

Advanced customizations

You can further customize the dashboard by editing the global layout file. SeeCustomizing Site and Space Layoutsfor more information on how to do this. You'll need some knowledge of Velocity to modify the layout files.

There are two locations that you can add content to:

- Web panels added toatl.dashboard.secondarywill appear below the site welcome message.
- Web items added tosystem.dashboard.buttonwill appear next to the Create space and Invite
 users button at the top right of the dashboard.

If you modify layouts in Confluence you will need to reapply your modifications each time you upgrade Confluence. The more dramatic your customizations are, the harder it may be to reapply the changes when upgrading. SeeUpgrading Customized Site and Space Layoutsto find out what will be involved before modifying the layouts.

On this page:

- Editing the site welcome message
- Using a page as the site landing page
- Advanced customizations

Related pages:

- Save for later
- Changing the Look and Feel of Confluence

Changing the Site Logo

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

You can change:

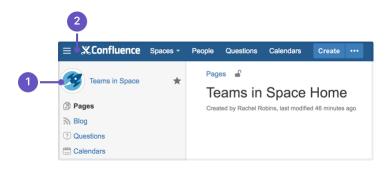
- the site logo
- the default space logo for all spaces
- the **space logo** for individual spaces.

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



- 1. **Space logo:**appears in the sidebar and on the dashboard.
- 2. Site logo: always visible, click the logo to go to the dashboard (or site homepage).

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 O, then choose General Configuration
- 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 colors in your new logo, and change the site color scheme to match.

If you would prefer to use the default color scheme with your custom logo go to General Configuration> Color Scheme > Edit and then choose Reset to revert back to the default scheme.



- 1. Site logo: auto look and feel has updated the header colours to complement the logo.
- 2. Site title: this is the name of your site.

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. Go to General Configuration>Site Logo and Favicon.
- 2. 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. The default space logo applies to all spaces that do not have a custom space logo applied - see Configure the Sidebar.

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

To change the default space logo:

- 1. Go to Some General Configuration Default Space Logo.
- 2. Choose Logo:ON
- 3. ChooseBrowseto upload a new logo
- 4. Choose Upload Logo
- 5. ChooseSave.

Changing a specific space logo

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

SeeseeConfigure the Sidebarto find out how to change the logo in a specific space.

Customizing Color Schemes

Confluence administrators can configure a new color scheme for the site. The default color scheme for the site will also become the default for all spaces within it.

To change the site's color scheme:

- 1. Choose thecog icon O, then chooseGeneral Configuration
- 2. Choose Color Scheme in the left-hand panel
- 3. Click**Edit**
- 4. Enter standard HTML/CSS2 color codes, or use the color-picker to choose a new color from the palette provided.
- 5 HitSave

Any changes you make will immediately be reflected across the Confluence site.

Reset your color scheme after uploading a site logo

When you upload a site logo, Confluence automatically detects the colors in your logo and customizes the color scheme for you.

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

To reset the color scheme:

- 1. Choose the cog icon , then choose General Configuration
- 2. Choose Color Scheme in the left-hand panel
- 3. ClickEdit
- 4. HitReset

On this page

 Reset your color scheme after uploading a site logo

Related pages:

 Changing the Look and Feel of Confluence

Styling Confluence with CSS

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

Introduction

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

Note: By default, only system administratorscan edit the CSS for a space or for the site. To allow any user with Space Admin permissions to edit the CSS for a space, go to Security Configuration and select Custom Stylesheets for Spaces.

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

In order to get you started, we have compiled this introduction, a basic styling tutorial.

Considerations for Using Custom CSS

CSS Knowledge is Required

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

Security

Custom CSS can be used to inject scripts into a page, opening the risk of cross-site scripting (XSS) attacks. With this feature enabled, space administrators could upload styles that steal other users' login credentials, trick their browsers into performing actions on the wiki without their knowledge, or even obtain global administration privileges. As such, this feature is disabled by default. Confluence administrators should only enable custom CSS if they are comfortable with the risks listed in this paragraph.

Scaling

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

Features Cannot Be Disabled

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

Features Should Not Be Disabled

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

On this page:

- Introduction
- Considerations for Using Custom CSS
- Getting Started
- CSS Resources

Related pages:

- Basic Styling Tutorial
- Styling Fonts in Confluence

Custom CSS does not apply to Admin screens

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

Confluence Version Compatibility

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

Test on Different Web Browsers

As a rule you should test your modifications on the various web browsers supported by Confluence.

CSS Customization is Not Supported

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

Getting Started

Editing the CSS

To edit a space's CSS style sheets:

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

To edit your global CSS stylesheet:

- 2. Choose Edit.
- 3. Paste your custom CSS into the text field.
- 4. Choose Save.

Note:

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

Follow the Tutorial

Follow the examples in the Basic Styling Tutorial to get started.

CSS Resources

- W3C CSS Standards
- Mozilla Developer Network
- W3resource.com

Basic Styling Tutorial

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

CSS Editing Quick-Start

To edit a space's CSS style sheets:

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

On this page:

- CSS Editing Quick-Start
- Tutorial: Changing the Header Background
- CSS Editing Tips
- Notes

Related pages:

 Styling Confluence with CSS

Tutorial: Changing the Header Background

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

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

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

- 5. Log in as the Space Administrator.
- 6. Open the Space Admin page.
- 7. Click Stylesheet.
- 8. Click **Edit** to change the code in the text field.
- 9. Paste your custom CSS into the text field.
- 10. Click **Save** and then reload the page (you may have to shift-reload). The background of the header will change.
- 11. The custom header will be visible on all content pages in the space. To revert your change, simple delete the custom code from the 'Stylesheet' page and click **Save**.

CSS Editing Tips

Begin With a Space Stylesheet

A space stylesheet is a good starting point for CSS customization, as it already includes all of the elements that can be changed. When you work on the space stylesheet it styles all content pages in the space. Build and test it at space-level, before considering applying the new stylesheet to your entire site. Once you are satisfied with your space design, test it thoroughly until you are confident that it has no problems. Then, you can look into advanced customization of the Confluence CSS such as adjusting the Search page, the Dashboard and other integral pages.

Use the Right Tools

As the Confluence CSS is reasonably sophisticated, web development applications will help you to understand how the page styles have been created. In particular, you will need to view the existing source for the pages you're starting to work on. If you don't already have some, tools such as the following free applications will allow you to do this.

1. Firebug

Firebug, a plugin for the Firefox web browser, allows you to take a look at the style of each element on your page. This is very useful to see what styles are currently applied, for example styles applied to the header only.

2. Web Developer

The Web Developer plugin for Firefox allows you to edit CSS inline and create new page designs.

3. CSS Edit

CSS Edit is a stand-alone CSS editor for Macintosh that extracts all existing styles from a given page and allows you to overwrite these.

Edit Simple Elements First

Begin by editing simple elements and checking that they work. By making changes, then checking that each one worked, you can easily isolate any CSS code that is causing problems. Be aware that some page elements are more suited to customization than others. For example, adding a gradient to the toolbar is less likely to 'break' the page than changing the page width. Editing reasonably static elements such as background graphics will render more predictably than designs which attempt to completely change the user interface or the Javascript-powered drop-down menus (which we don't recommend editing).

Notes

Note: By default, only system administratorscan edit the CSS for a space or for the site. To allow any user with Space Admin permissions to edit the CSS for a space, go to General Configuration Security Configuration and select Custom Stylesheets for Spaces.

Styling Fonts in Confluence

Confluence provides the ability to adjust its visual style via Cascading Style Sheets (CSS). This tutorial shows you to change the fonts and font sizes of a Confluence page, using a few lines of CSS.

Below is the code for the custom font. Copy and paste it into the Space Stylesheet form within the Space Administration section.

Related pages:

- Basic Styling Tutorial
- Styling Confluence with CSS

Changing the fonts

In order to customize the fonts in Confluence, you first need to set the body font to the font you want. Secondly, you may want to adjust the font size because different fonts have different relative sizes.

The relevant CSS is shown below. It changesConfluence's font from the default of Helvetica/Arial sans serif to Times/Times New Roman serif. To adjust for the fact that Times is a bit smaller than Helvetica, we increase the font size to 14 pixels. The many styles that 'wiki-content' in their definition are necessary to change the font size for all the tags in the wiki content.

To edit a space's stylesheet:

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

```
.wiki-content p,
.wiki-content table,
.wiki-content tr,
.wiki-content td,
.wiki-content td,
.wiki-content ol,
.wiki-content ol,
.wiki-content ul,
.wiki-content li {
  font-family: Times, "Times New Roman", serif;
  font-size: 14px;
}
```

Notes

Note: By default, only system administratorscan edit the CSS for a space or for the site. To allow any user with Space Admin permissions to edit the CSS for a space, go to General Configuration Security Configuration and select Custom Stylesheets for Spaces.

Working with Themes

Themes are used to change the appearance of your Confluence site or spaces.

Confluence comes with a single default theme installed, or you can download and install other themes from The Atlassian Marketplace.

Once a theme is installed it can be applied to the whole site or to individual spaces.

To see the themes installed in your site:

- 1. Go to O-> General Configuration > Themes.
- 2. You'll see a list of all the themes installed in your site.

When a new space is created, whichever theme is applied to the whole site will be applied by default to the new space. The space theme can then be changed by anyone with space administrator permissions for that space.

Note about the Documentation theme

The Documentation theme was available in Confluence 5.9 and earlier. Many of the Documentation theme features are now available in the Confluence default theme. Check out Develop Technical Documentation in Confluence for more information about using Confluence for documentation using the default theme.

Related pages:

- Apply a Theme to a Space
- Applying a Theme to a Site
- Creating a Theme

Applying a Theme to a Site

Themes are used to change the appearance of your Confluence site. SeeW orking with Themes for an overview of how themes apply to your whole site, and how you can add more themes. To apply a theme across the site:

- 1. Go to General Configuration > Themes.
- 2. The screen will display all available themes. Choose a theme.
- 3. Choose Confirm.

All spaces that have the **Global look and feel** applied as their space theme will inherit this theme and any customizations you make to it.

Related pages:

 Apply a Theme to a Space

Creating a Theme

If you want to create your own theme, you will need to write a Confluence plugin. Please refer to the following pages in our developer documentation:

- Get started with plugin development.Follow the developer's tutorial for writing a Confluence theme.
- Create a theme using the theme plugin module.

Related pages:

- Applying a Theme to a Site
- Apply a Theme to a Space

Customizing Site and Space Layouts

You can modify Confluence's look and feel by editing layout files (also known as decorators). Editing these files allows you to change the look and feel of the whole Confluence site, or just an individual space.

When you edit a site layout, you'll be modifying the default decorators in every space in your site, except for those that have already been edited in a space. SeeCustomize Space Layoutsfor more information on how to edit the decorators for a single space.

You'll need System Administrator permissions to edit site layouts.

On this page:

- Editing a site decorator file
- Using Velocity macros
- Advanced customizations

Related pages:

- Velocity Template Overview
- Basic Introduction to Velocity
- Customizing your Confluence Site

If you modify layouts in Confluence you will need to reapply your modifications each time you upgrade Confluence. The more dramatic your customizations are, the harder it may be to reapply the changes when upgrading. SeeUpgrading Customized Site and Space Layoutsto find out what will be involved before modifying the layouts.

Confluence is built on top of the open source SiteMesh library, a web-page layout system.

To edit the layout of Confluence, you will need to modify these decorator files. A decorator file is a . vmd file and is written in Velocity. You can learn more from the Velocity User Guide.

Once you are familiar with Velocity, you can edit the decorator files to personalize the appearance of Confluence.

The decorator files in Confluence are grouped into the following categories:

- Site layouts: These are used to define the controls that surround each page in the site. For example, the header, footer and dashboard.
- Content layouts: These control the appearance of content such as pages and blog posts. They do not change the way the pages themselves are displayed, but allow you to alter the way the surrounding comments or attachments are displayed.
- Export layouts: These control the appearance of spaces and pages when they are exported to HTML.

Editing a site decorator file

To edit a site decorator:

- 1. Go to General Configuration Layouts (under Look and Feel)
- 2. Click Create Customnext to the decorator . vmdfileyou want to modify.
- 3. Make your changes and click Update.

If something goes wrong: HitReset 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 customizations

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 texteditor 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 Customized 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 functionally is changed, the default layouts are modified to support these changes.

Related pages:

Customizing Site and Space Layouts



⚠ If you are using custom layouts based on defaults from a previous Confluence version, you run the risk of breaking functionality, or worse, missing out on great new features!

Take care on each new release of Confluence to reapply your changes to the new default templates.

To reapply your custom layouts, you need to:

- 1. Obtain the source of your custom layouts from your current version of Confluence.
- 2. Reapply your customizations to the new default layouts.

Step 1. Obtain your Custom Layouts

Ideally, you should keep a record of each customization you have applied to each of your Confluence site or space layouts.

If not, you should be able to find your customizations 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 customizations.



This method is handy to use if you have:

- Many spaces with space layout customizations, or
- Do not have an independent record of your site or space layout customizations.

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 Customizations

When you upgrade Confluence to another major release of Confluence, you will need to manually reapply any customizations you made to any site-wide or space-specific layouts. Unless otherwise stated, you should not need to reapply customizations 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 customizations:

- 1. Choose the cog icon O, then choose General Configuration
- Select Layouts in the left-hand navigation panel. The decorators are grouped under Site, Content and Export layouts.
- 3. Ensure you have all your customizations available (preferably in a form which can be copied and pasted).
- 4. Click **Reset Default** next to the layout whose customizations need to be reapplied.
- 5. Click **Create Custom** next to the same layout and reapply your customizations (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 customizations need to be reapplied.

If you have made space-specific layout customizations:

- 1. Go to the space and choose Space tools > Look and Feel from the bottom of the sidebar
- 2. ChooseLayout. The decorators are grouped under Site, Content and Export layouts.
- 3. Ensure you have all your customizations available (preferably in a form which can be copied and pasted).
- 4. Click Reset Default next to the layout whose customizations need to be reapplied.
- 5. Click **Create Custom** next to the same layout and reapply your customizations (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 customizations need to be reapplied.

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, wherex.x.x is the Confluence version number. The JAR file is located in the WEB-INF/lib directory. If you wish to make modification to the files in the JAR, we recommend the following steps:

- 1. Stop Confluence.
- 2. Make a backup copy of the JAR file.
- 3. Un-jar the file
- 4. Locate and edit the appropriate file that you wish to modify.
- 5. Re-jar the confluence-x.x.x. jar file.
- 6. Relocate the JAR file to the appropriate directory.
- 7. Restart Confluence.

Working With Decorator Macros

Decorator Macros are Velocity macros which are used to draw complex or variable parts of the page such as menus and breadcrumbs when editing Custom decorators. Decorator macros can be inserted anywhere in your templates.

The macro is called by inserting a string of the form: #macroName("argument1" "argument2" "argument3"). There are no commas between the arguments. Unless otherwise noted, these macros take no arguments.

NOTE: These macros will only work reliably when customizing main. vmd. They may not work in other Velocity decorators. Decorator macros will not work inside normal confluence pages.

Macro	Usage
#brea dcrum bs()	Draws the "You are here" breadcrumbs list, like the one found above the page name in the default template.
#incl udePa ge (page Title)	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.
#sear chbox ()	Inserts a search box into the page, like the one to the far right of the breadcrumbs in the default template.
#glob alnav bar (type)	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:
#glob alnav bar ("tab le")	Displays the navigation bar in its default mode: drawn as a table of links with colored backgrounds and mouse-over effects.
#glob alnav bar ("tex t")	Displays the navigation bar as series of text links separated by characters.
#user navba r()	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.
<pre>#help icon()</pre>	Draws the help icon, and link to the Confluence help page.
<pre>#prin table icon()</pre>	On pages where a printable version is available, draws the printable page icon, linking to the printable version of the page. Otherwise, draws nothing
<pre>#page title (clas s)</pre>	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 customized your Confluence installation's CSS file, you should call this with "spacenametitle" as the class: #pagetitle("spacenametitle")

<pre>#powe redby ()</pre>	Writes out the "Powered by Confluence" and Confluence version-number boilerplate found at the bottom of the default template.
#bott omsha dow()	Draws the fading shadow-effect found at the bottom of the content area in the default template.
#dash board link()	Inserts a link to the dashboard page.

Custom Decorator Templates

About Decorators

Confluence is built on top of the Open Source SiteMesh library, a web-page layout system that provides a consistent look and feel across a site. SiteMesh works through "decorators" that define a page's layout and structure, and into which the specific content of the page is placed. If you are interested, you can read more in the SiteMesh documentation.

What this means for Confluence is that you can customize 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 customization.

Editing Custom Decorators

To edit Confluence decorators you will need a good knowledge of HTML, and some understanding of the Velocit y 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.

Customizing 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 <baseU rl>/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:

- 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), its 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.

Customizing the Login Page

This page gets you started on customizing the Confluence login page, to add your own logo or custom text. This will not customize the login *process*, just what users sees when they log in.

Notes:

- Customizations 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
- Customizing Site and Space Layouts
- Changing the Look and Feel of Confluence
- Modify Confluence Interface Text

To change the login page:

- 1. Shut down your Confluence server.
- 2. In the Confluence installation directory, find the file confluence/login.vm.
- 3. Make a copy of this file as a backup.
- 4. Edit the file with a text editor to make the required changes. The content contains a mixture of HTML and Velocity. See Velocity Template Overview (in our developer documentation).
- 5. Start Confluence and test your changes.

The same process can be applied to modify most of the templates in the Confluence web application. Be careful to test your changes before applying them to a live site. The templates contain code that is vital for Confluence to function, and it is easy to accidentally make a change that prevents use of your site.

Modify Confluence Interface Text

All Confluence UI text is contained in a single Java properties file. This file can be modified to change the default text, and also to translate Confluence into languages other than English.

The UI text file is ConfluenceActionSupport.properties. From your Confluence install directory:

```
\confluence\WEB-INF\lib\confluence-x.x.x.jar

Replace "x.x.x" with your Confluence version, for example for 4.3.2, it will be named "confluence-4.3.2. jar".

Within this File, the relevant file to edit is :\com\atlassian\confluence\core\ConfluenceActionSupport. properties.
```

Refer to Editing jar files for reference.

The file contains parameters with name=value pairs, in the format:

```
parameter.name=Parameter value
```

Parameter names are any text before the '=' character and should never be modified. Any text after the '=' character is the parameter value, which can be modified freely and can also contain variables. An example involving variables is:

```
popular.labels=The three most popular labels are \{0\}, \{1\} and \{2\}.
```

For more information on replacing values, check out Translating ConfluenceActionSupport Content. Note that plugins store their text internally, so you must modify plugin text individually.

Steps For Modification

- 1. Stop Confluence
- 2. Under your install directory, open \confluence\WEB-INF\lib\confluence-x.x.x. jar\com\atlassian\confluence\core\ConfluenceActionSupport.properties
- 3. Search for the text you wish to modify, replace it and save the file in <Confluence-Install>\confluence\WEB-INF\classes\com\atlassian\confluence\core. Please create this folder structure, if it does not exist already.



If you re-bundle the JAR file, rather than re-deploy the class in the WEB-INF\classes directory, make sure to move the backup JAR file out of the /lib directory, or the backup may be deployed by mistake.

4. Restart Confluence

Modify Keyboard Shortcuts

Confluence provides a set of keyboard shortcuts. You could customize 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 Attachmentsone 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.

Customizing Email Templates



Customizing 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 customizations 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-notificationsplugin, which is a system app (plugin) that is installed automatically when you install Confluence.

Only administrators with access to the Confluence installation directory can modify the Confluence email templates.

Confluence uses Soy templates (also known as Closure templates) for email notifications. You can find out more in the Google Developer docs or see our developer tutorial which contains a short introduction to using Soy templates.

To change the email notification templates:

- 1. In the Confluence web application folder, find the file/confluence/WEB-INF/atlassian-bundledplugins/confluence-email-notifications-plugin-x.x.jar Note: This plugin is independently versioned, the version number will not necessarily match Confluence's version number.
- 2. Copy this file to a working location and extract the jar file. Find out more about how to edit files within .jar archives.
- 3. Within the jar file, templates are stored in the /templates / folder. Edit the Soy templates to make your changes.
- 4. Zip all the files and change the file extension to .jar (or refer to theguide on editing files within .jar archives for other methods).
- 5. Drop the new jar file into the /confluence / WEB-INF / atlassian bundled plugins folder (replacing the original file - you might want to make a copy of the original file for easy roll back) and then restart your instance.
- 6. Test your changes carefully before installing the updated plugin in production.

We strongly recommend you use a test instance for editing the templates contained within the plugin. If you are unable to enable the plugin, check the Confluence logs for information, it may be that there are problems with your edits to the Soy templates.

RELATED TOPICS

- Customizing Site and Space Layouts
- Changing the Look and Feel of Confluence
- Modify Confluence Interface Text

Changing the Default Behavior 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 customize the behavior and the default content of their Confluence site:

Related pages:

 Changing the Look and Feel of Confluence

- Administering Site Templates
- Changing the Site Title
- Choosing a Default Language
- Configuring the Administrator Contact Page
- Configuring the Site Home Page
- Customizing Default Space Content
- Editing the Site Welcome Message

Administering Site Templates

A template is a predefined page that can be used as a prototype when creating new pages. Templates can be created by users, or provided by a blueprints. See Page Templates and Blueprints.

Confluence also provides 'system templates' which contain default content for the site welcome message (see E diting the Site Welcome Message) and default space content (see Customizing 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:

- Go to General Configuration>Global Templates and Blueprints.
- Select**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.

Changing the Site Title

The site title appears in your browser's title bar. By default, it is set to 'Confluence'. The site title can't be empty.

To change the title of your Confluence site:

- 1. Go to General Configuration.
- 2. Choose **Edit**at the top of the **Site Configuration**section.
- 3. Enter a new title for your site then choose **Save**.

Related pages:

- Changing the Site Logo
- Editing the Site Welcome Message
- Customizing your Confluence Site

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.

Setting the default language

To change the default language for the Confluence site:

- 1. Choose the cog icon , then choose Genera I Configuration
- 2. Select 'Languages' in the 'Configuration' section of the left-hand panel.
- Choose Edit and select the language you want to use as the default language for your Confluence site.

Related pages:

- Edit Your User Settings
- Recognized System Properties
- Configuring Indexing Language
- Installing a Language Pack

Confluence comes with the following languages installed and ready to use:

- etina (esk republika | Czech Republic)
- Dansk (Danmark | Denmark)
- Deutsch (Deutschland | Germany)
- Eesti (Eesti | Estonia)
- English (UK)
- English (US)
- Espaol (Espaa | Spain)
- Franais (France)
- slenska (sland | Iceland)
- Italiano (Italia | Italy)
- Magyar (Magyarorszg | Hungary)
- Nederlands (Nederland | The Netherlands)
- Norsk (Norge | Norway)
- Polski (Polska | Poland)
- Portugus (Brasil | Brazil)
- Romn (Romnia | Romania)
- Slovenina (Slovensk republika | Slovak Republic)
- Suomi (Suomi | Finland)
- Svenska (Sverige | Sweden)
- (| Russia)
- (| China)
- (| Japan)
- (| Republic of Korea)

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 prioritized list of languages. Confluence will use the first supported language in that list. Confluence administrators can disable this option by setting the confluence.browser.language.enabled system property to false.

• 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 troubleshooting translations of the Confluence user interface. After opening the Confluence dashboard, you can add the following action to the end of your Confluence URL:

?il8ntranslate=on

For example http://myconfluencesite.com?il8ntranslate=on

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.

The key names are displayed with a 'lightning bolt' graphic. Here's an example from a space sidebar:



To turn off the translation view, add the following to the end of the Confluence URL:

?il8ntranslate=off

Configuring the Administrator Contact Page

The administrator contact page is a form that allows a user of Confluence to send a message to the administrators of their Confluence site. (In this context, administrators are the members of the default administrators group.)

See the explanation of Confluence Groups for Administrators.

The title of the administrator contact page is 'Contact Site Administrators'. Typically, Confluence users may get to this page by clicking a link on an error screen such as the '500 error' page.

On this page:

- Customizing the Administrator Contact Message
- Disabling the Administrator Contact Form
- Configuring Spam Prevention

Related pages:

Configuring Captcha for Spam Prevention

Customizing the Administrator Contact Message

You can customize 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
- 2. ChooseGeneral 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 Messagebox. You can enter any text or Confluence wiki markup.
- 5. ChooseSave.

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			
	on about your request for the site admini lude information on what you were doing	, ,	
То	Confluence Administrators		
From	user@email.com		
Subject*	Administrator Request		
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 O, then choose General Configuration
- 2. Choose **General Configuration** in the left-hand panel.
- 3. Choose **Edit** at the top of the 'Site Configuration' section.
- 4. Select on or off for the 'Contact Administrators Form'.
- 5. Choose Save.

Configuring Spam Prevention

You can configure Confluence to use Captcha to help prevent spam, including the spamming of Confluence administrators. The administrator contact form is covered by the site-wide Captcha settings as documented in Configuring Captcha for Spam Prevention.

Configuring the Site Home Page

The dashboard is the default home page for your site, but you can choose to use a space homepage as the landing page for your site.

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

Users can also choose to override the site homepage and use the dashboard or a different page as their landing page in theirpersonal settings.

Related pages:

- Editing the Site Welcome Message
- Changing the Site Title
- Changing the Site Logo

To use a page as your site home page:

- 1. Go to General Configuration>Further Configuration.
- 2. ChooseEdit.
- 3. Select a space from the **Site Homepage**dropdown menu. When users log in or click the site logo, Confluence will go to the home page of the space you choose here.
- 4. ChooseSave.

Note about permissions

Before changing the site homepage you should check that the default 'confluence-users' or 'users' groups have permissions to view the space the page was created in, and that the page itself is not restricted to particular people or groups.

If your site is public, you'll also need to make sure anonymous users have permissions to view the space, otherwise anonymous users will be directed to the dashboard instead.

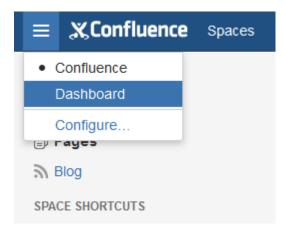
Accessing the dashboard with a site homepage set

If you choose to set a page as your site homepage but would like your users to still be able to access the Confluence dashboard, you can add a link to the Application Navigator.

To add the Confluence Dashboard to the Application Navigator:

- 1. Go to General Configuration>Application Navigator.
- 2. Enter the name for your link, for example, 'Dashboard'.
- Enter the URL for your site dashboard, for example, https://yoursite.com/wiki/dashboard.action.
- 4. ChooseAdd.

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



Customizing 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
- 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 customize 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 saw 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:

- Spaces
- Page Templates

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
- 2. Choose Global Templates and Blueprints in the left-hand panel.
- 3. Choose **Reset to default**next to the template you wish to reset.

From this point on, all new space home pages will be created with the original default content.

Editing the Site Welcome Message

Give your site's landing page some personality by editing the site welcome message.

The site welcome message appears on the right hand side of the dashboard and is perfect for adding announcements, useful links, or a fun photo from your last office party or team outing.

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

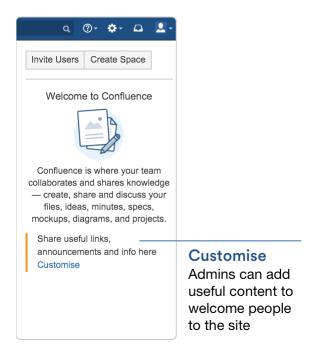
To edit the site welcome message:

Confluence administrators can either click the **Edit** link below the site welcome message on the dashboard, or:

- 1. Go to General Configuration>Global Templates and Blueprints.
- $2. \ \, {\sf Scroll \ down \ to \ the \ System \ templates \ and \ choose} \textbf{Edit} next \ to} \textbf{Default \ Welcome \ Message}.$
- 3. Add your content and chooseSave.

You can go back to the original welcome message at any time - choose **Reset to Default** next to the **Default** welcome message template.

Screenshot: Default site welcome message



Hints for using the template editor

The site welcome message is a template, not a page, so you'll be using the template editor to make your changes.

You can add text, links and macros, as you would in any confluence page, but the process for adding files, including images is a little different.

On this page:

- Hints for using the template editor
- Allowing other people to edit the site welcome message

Related pages:

- Configuring the Site Home Page
- Changing the Site Title
- Changing the Site Logo

You can't upload an image or other file into a template directly. First you'll need to upload the file to a page in your site, then in your template, choose**Insert** > **Files** > **Search on other pages**to embed the file or image.

You can't use template variables in the site welcome message.

Allowing other people to edit the site welcome message

You can allow people who are not Confluence administrators to edit the site welcome message by using the include Include Pagemacro to include content from elsewhere in your site, rather than adding content directly to the template.

To include content from a page in the site welcome message:

- 1. Create a new page in a space that is visible to all users. It's important that all users can see content in that space if a person does not have permissions to view the space where you've created the page, they won't be able to see the page content on the dashboard.
- 2. Add some text, images or macros, then save the page.
- 3. Restrictwho can edit the page (this is optional, but useful if you only want to allow some people to change the content).
- 4. Edit the site welcome message template (as described above) and use the Include pagemacro to include the contents of your newly created page.
- 5. Save the template.

People with permission to edit the page will now be able to make changes at any time, and their changes will be visible on the dashboard as soon as the page is saved.

Integrating Confluence with Other Applications

You can integrate Confluence with other applications using **Application Links**. The Application Links feature allows you to link Confluence to applications such as JIRA Software or JIRA Service Management.

Linking two applications allows you to share information and access one application's functions from within the other. For example, you can display a list of issues on a Confluence page using the Jira Issues Macro.

Related Topics

- Linking to Another Application
- Configuring Workbox Notifications
- Integrating Jira and Confluence
- Registering External Gadgets
- Configuring the Office Connector
- Managing Webhooks

Linking to Another Application

Application Links (sometimes called *app links*) is a bundled app that allows you to set up links, share information, and provide access to certain resources or functionality across multiple Atlassian products. We recommend using OAuth authentication for application links because of the greater security inherent with that protocol. Weno longer recommend the Trusted Applications and Basic authentication types.

Linking Confluence to other applications allows you to include information from those applications in pages or blogs that you create in Confluence. For example, you could link Confluence to Jira Software and display issues on a Confluence page using the Jira Issues Macro.

- 1. Go to General Configuration Application links.

 The Application Links configuration page appears and lists any links you already have set up.
- 2. Enterthe URL of the application you want to link to then select Create new link.
 - (i) Cloud to Cloud integration

When adding the Cloud URL, use the **https://**instance.atlassian.net. If the https:// is not used, the link will not be completed.

- If you check**The servers have the same set of users...**,then this link will be configured using OAuth (with impersonation) authentication.
- If you are *not*an admin on both servers you won't be able to set up a 2-way (reciprocal) application link. If you want to go ahead and create a 1-way link anyway, clear the lam an administrator on both instances checkbox.
- Use the wizard to finish configuring the link. If the application you are linking to does not have the Application Links plugin, you must supply additional information to set up a link with OAuth authentication.

When you complete the wizard, the Application Links plugin will create the link between your applications using the most secure authentication method that is supported between the two applications. See the Application Links User Guidefor more information.

The new link will appear on the Configure Application Links page, where you can:

- Edit the settings of the application link (for example, tochange the authentication typeof the link) using theEdit () icon.
- Specify the default instance if you have multiple links to the same type of application (for example, to
 multiple Jira servers) using the Make Primarylink. SeeMaking a primary link for links to the same
 application typefor more information.
- (i) Having trouble integrating your Atlassian products with application links?

We've developed aguide to troubleshooting application links, to help you out. Take a look at it if you need a hand getting around any errors or roadblocks with setting up application links.

Configuring Workbox Notifications

You can view and manage in-app notifications and tasks in yourConfluence workbox. In addition, you can receive notifications from Jira applications and other Confluence servers in your Confluence workbox. To make this possible, your Confluence server must be linked to the other server(s) via application links.

Possible configurations:

- Your Confluence server provides in-app notifications and displays them in its own workbox. There are two sub-configurations here:
 - This Confluence server is the only server involved.
 - Alternatively, this Confluence server displays its own in-app notifications, and also displays notifications from Jira and/or other Confluence servers.
- Your Confluence server does not provide or display in-app notifications.
- Your Confluence server sends in-app notifications to another Confluence server.

On this page:

- Which notifications are included?
- Configuring the polling intervals
- Including notifications from Jira
- Stopping Jira applications from sending notifications to Confluence
- Including notifications from another Confluence server
- Sending Confluence notifications to another Confluence server
- Disabling workbox and in-app notifications in Confluence

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.

Which notifications are included?

The workbox displays a notification when someone does one of the following in Confluence:

- Shares a page or blog post with you.
- Mentions you in a page, blog post, comment or task.
- Comments on a page or blog post that you are watching.
- Likes a page or blog post that you are watching.

The workbox does **not** show notifications triggered because you are watching a space. Only watches on pages and blog posts are relevant here.

The notification in your workbox appears as 'read' if you have already viewed the page or blog post.

If your Confluence site is linked to a Jira application, you will also see the following Jira notifications in your workbox:

- Comments on issues that you are watching.
- Mentions.
- Shares of issues, filters and searches.

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

If your Confluence site is connected to a Jira application, you can include notifications from your Jira application, for example Jira Software or Jira Service Management.

To include notifications from a Jira application:

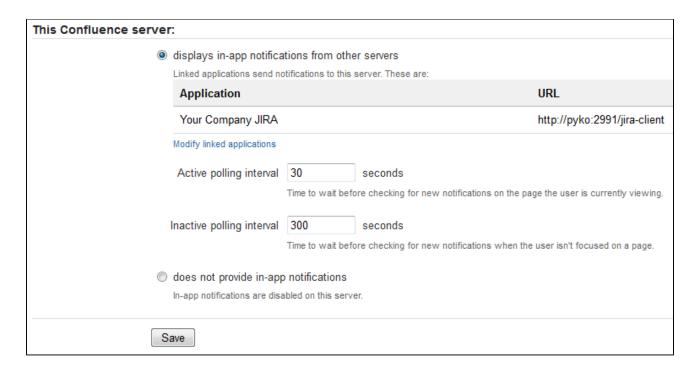
Your Jira application and Confluence must be connected via an application link. See Linking to Another Application.

- 1. Choose the cog icon , then choose General Configuration
- 2. Choose **In-app Notifications** in the left-hand panel of the Confluence administration console.
- 3. Choose displays in-app notifications from other servers.
 - Your Jira application will appear in the list of linked applications below this option.
 - People will see Jira notifications in their workbox, as described in Workbox Notifications.

Notes:

- Jira sends its notifications to the Confluence server that is configured as the primary application link.
- Your Jira server must be runningJira 5.2 or later.
- The following system apps must be present and enabled in Jira. The apps are shipped with Jira 5.2 and later:
 - 'Workbox Common Plugin'
 - 'Workbox Jira Provider Plugin'
- You do not need to configure Jira. The system apps are enabled by default in Jira, and Jira will automatically send notifications to Confluence.
- The application link must use OAuth authentication. If you don't see your Jira application listed, you will need to edit the application link (in both applications) to change the authentication type.
- Confluence can display notifications from more than one server.

Screenshot: This Confluence server displays in-app notifications from itself and from Jira



Stopping Jira applications from sending notifications to Confluence

You may wish to configure Confluence to display its own notifications in its workbox, but prevent notifications from Jira applications from appearing in the workbox, even when JIRA applications and Confluence are linked via application links.

The Jira administration interface does not offer a way of disabling notifications sent to Confluence.

To stop Jira applications from sending notifications to Confluence: Disable the following plugins in Jira. (See the Universal Plugin Manager guide to disabling plugins.)

- 'Workbox Common Plugin'
- 'Workbox Jira Provider Plugin'

Including notifications from another Confluence server

Confluence workbox can include notifications from another Confluence server.

Let's assume that you have two Confluence servers, ConfluenceChattyandConfluenceQuiet. Let's also assume that you wantConfluenceChattyto display a workbox, and to include notifications fromConfluenceQuiet.

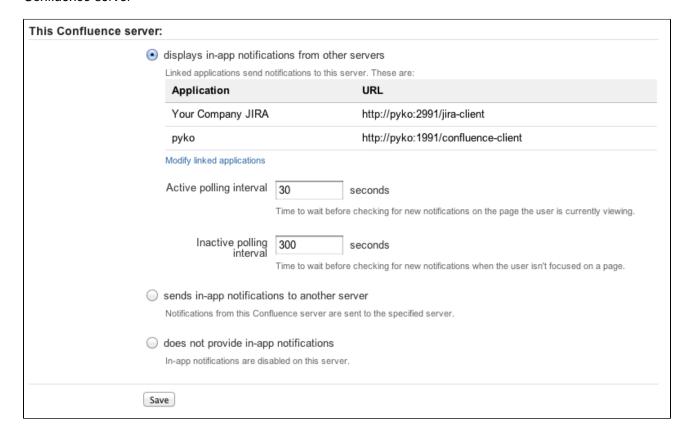
To include notifications from other Confluence servers:

- 1. Connect Confluence Chattyand Confluence Quietvia application links. In Confluence Chatty:
 - Choose thecog icon O, then chooseGeneral Configuration
 - ChooseApplication Linksin the left-hand panel.
 - Set up the link as described inLinking to Another Application.
- 2. Configure the notification settings in Confluence Chatty:
 - ChooseIn-app Notifications in the left-hand panel of the Confluence administration console.
 - Choosedisplays in-app notifications from other servers.
- 3. Configure the notification settings in Confluence Quiet.
 - ChooseIn-app Notifications in the left-hand panel of the Confluence administration console.
 - Choosesends in-app notifications to another server.
 - Select the Confluence server that will display the workbox in our example, this is ConfluenceCh atty. (The entry for ConfluenceChattywill appear here only if you have already configuredConfluenceChattyto display in-app notifications.)

Notes:

- Your Confluence servers must be runningConfluence 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

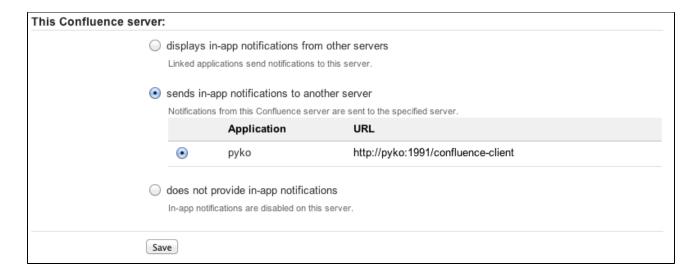


Sending Confluence notifications to another Confluence server

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 *Confluence Quiet* above.

Screenshot: This Confluence server sends its in-app notifications to another Confluence server



Disabling workbox and in-app notifications in Confluence

If you choose does not provide in-app notifications:

- The Confluence workbox icon will no longer be visible and people will be unable to access their workboxes on this server.
- This Confluence server will no longer send notifications to its workbox, and will not send notifications to any other Confluence server.

Integrating Jira and Confluence

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

Learn more about what you can do with Jira and Confluence

Here's some ways you can get Jira and Confluence working together.

Installing Jira and Confluence together

We recommend running Jira and Confluence in separate stand-alone instances behind an Apache Web Server. The following documentation will guide you through the installation processes:

- Installing Confluence
- Installing Jira applications
- Running Confluence behind Apache
- Integrating Jira with Apache

We don't support deploying Confluence and any other application (including Jira) in the same Tomcat container. SeeCan Multiple Atlassian Products Be Deployed in a Single Tomcat Container?for more information.

Use Jira and Confluence together

This is the fun stuff. Check outUse Jira applications and Confluence togetherto find out about all the integration points, great time saving features, and to check exactly which Jira application and version you'll need.

Delegate user management to Jira

If you already have a Jira application you can choose to delegate user management to Jira, and manage all your users in one place. You can control which Jira groups also have permissions to use Confluence. Your license tiers for each application do not need to be the same.

SeeConfiguring Jira Integration in the Setup Wizardto delegate user management to Jira when installing Confluence for the first time.

SeeConnecting to Crowd or Jira for User Managementto delegate user management to Jira for an existing Confluence site.

Connect Jira and Confluence with an application link

SeeLinking to Another Application to find out how to connect Confluence to your Jira application using an application link. This only needs to be done once.

If you delegated user management to Jira as part of Confluence's setup process, an application link to Jira will be all set up and ready to go.



Having trouble integrating your Atlassian products with application links?

We've developed aguide to troubleshooting application links, to help you out. Take a look at it if you need a hand getting around any errors or roadblocks with setting up application links.

On this page:

- Installing Jira and Confluence together
- Use Jira and Confluence together
- Delegate user management to Jira
- Connect Jira and Confluence with an application link

Registering External Gadgets

You can register gadgets from external sites (such as Jiraapplications), so the gadgets appear in the macro browser and people can add them to Confluence pages using the gadget macro.

There's two ways to register external gadgets:

- Subscribe to all of the external application's gadgets: You can add all the gadgets from your Jiraapplication, Bamboo, Fi shEye 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.

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 Allowlist
- The big list of Atlassian gadgets
- Linking to Another 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.

If you want to subscribe a third-party gadget, that doesn't require an application link, you will also need to add the gadget URL to the allowlist.

Subscribing to all of the application's gadgets

You can add all the gadgets from your Jira, Bamboo, FishEye or Crucible site or from another Confluence site to your Confluence gadget directory. People can then pick and choose the gadgets to add to their Confluence pages.

To subscribe to another site's gadgets:

- 1. Go to General Configuration > External Gadgets
- 2. Choose the **Gadget Feeds** tab.
- 3. Enter the base URL of the application you want to subscribe to, for example, http://example.com/jira or http://example.com/confluence.
- 4. Choose **Add**. Confluence will convert the URL to a gadget feed and place it in the list of 'Added Gadget Feeds'.

Screenshot: Subscribing to a gadget feed

External Gadgets						
Only add gadgets that you trust! Gadgets can allow unwanted or malicious code onto your web page.						
Gadget Specifications	Gadget Feeds	Gadget Whitelist				
Add a new Gadget F	eed					
Gadget F	eed URL					
	Add					
Added Gadget Feed	S					
Gadget Feed URL						Actions
https://pug.jira.com/rest/gad	dgets/1.0/g/feed					Delete
https://pug.jira.com/rest/gad	dgets/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 get the gadget URL and copy it to your clipboard.

Getting a gadget's URL from an Atlassian application

If your application is another Atlassian application:

A gadget's URL points to the gadget's XML specification file. In general, a gadget's URL looks something like this:

```
http://example.com/my-gadget-location/my-gadget.xml
```

If the gadget is supplied by a plugin, the URL will have this format:

For example:

 $\label{lem:mycompany.com/jira/rest/gadgets/1.0/g/com.atlassian.streams.streams-jira-plugin:activitystream-gadget/gadgets/activitystream-gadget.xml$

To find a gadget's URL in JIRA:

- Go to your dashboard by clicking the **Dashboards** link at the top left of the screen.
- Click Add Gadget to see the list of gadgets in the directory.
- Find the gadget you want, using one or more of the following tools:
 - Use the scroll bar on the right to move up and down the list of gadgets.
 - Select a category in the left-hand panel to display only gadgets in that category.
 - Start typing a key word for your gadget in the **Search** textbox. The list of gadgets will change as you type, showing only gadgets that match your search term.
- Right-click the Gadget URL link for that gadget and copy the gadget's URL into your clipboard.

To find a gadget's URL in Confluence:

- Choose Help> Confluence Gadgets to see the list of available Confluence gadgets.
- Find the gadget you want.
- Right-click the Gadget URL link for that gadget and copy the gadget's URL into your clipboard.

Getting a gadget's URL from another application

If the gadget comes from a non-Atlassian web application or web site, please consult the relevant documentation for that application to get the gadget URL.

Registering the gadget for use in Confluence

Now that you have the gadget's URL, you can register it in Confluence, so that people can add it to their pages. You need system administrator permissions to register a gadget.

To register the gadget in Confluence:

- 1. Go to General Configuration>External Gadgets
- 2. Paste your gadget's URL into the Gadget Specification URL field in the 'Add a new Gadget' section.
- 3. Choose **Add**. Your gadget will be shown in the list of registered gadgets below and it will also become available in the macro browser.

Screenshot: Registering external gadgets one by one

External Gadgets				
Only add gadgets that you trust! Gadgets can allow unwanted or malicious code onto your web page.				
Gadget Specifications	Gadget Feeds	Gadget Whitelist		
You can add gadgets from Atlassian applications such as Confluence, JIRA and others. You can also add gadgets from other websites such as iGoogle. Many public gadgets will work on a Confluence page. Some gadgets may rely on specific OpenSocial features that will not work properly in Confluence.				
using Trusted Applications (r	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 somet	hing like this: http:/	//example.com/my-g	adget-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				

Removing access to external gadgets

To remove a single gadget from Confluence, click the **Delete**button next to the gadget URL.

If you have subscribed to an application's gadgets, you will need to remove the entire subscription. You cannot unregister a single gadget. Click the **Delete**button next to the gadget feed URL.

The gadget(s) will no longer be available in the macro browser, and people will not be able to add them using the Gadget macro. Any pages that already use the gadget will show a broken gadget link.

Configuring the Office Connector

The Office Connector allows Confluence users to view and import content from Microsoft Office and Open Office files attached to a page.

The Office Connector system app is bundled with Confluence, but a System Administrator can enable or disable parts of the Office Connector and can configure options.

Enabling and disabling the Office Connector

If you want to limit access to all or part of the Office Connector you can disable the system app, or some modules in the app.

To enable or disable the Office Connector modules:

On this page:

- Enabling and disabling the Office Connector
- Configuring the Office Connector options

Related pages:

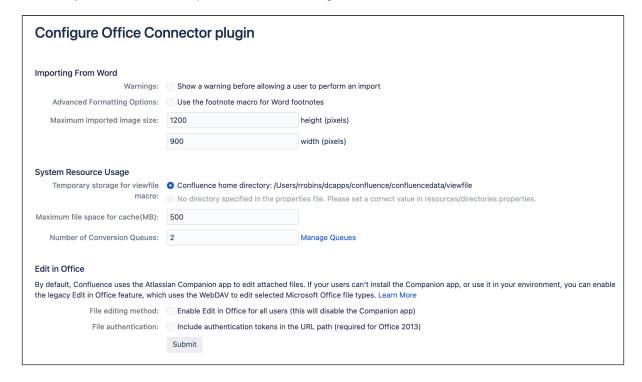
- Edit in Office using the Office Connector
- Office Connector Limitations and Known Issues

- 2. ChooseSystemfrom the filter drop down and then search forOffice Connector
- 3. Expand the Office Connector listing. From here you can:
 - Choose**Configure**to specify preferences for the Office Connector (this opens the configuration screen describedbelow)
 - Click Disable to disable all modules of the app
 - Expand the modules list to enable or disable selected Office Connector modules.

Note: only someOffice Connector modules can be disabled. Modules that are integral to the operation of the Office Connector cannot be disabled, and do not have an **Enable**or **Disable** button. Modules that can be disabled include the button and provide a brief description of the module.

Configuring the Office Connector options

Users with System Administrator permissions can change the behavior of the Office Connector.



To set the configuration options for the Office Connector:

- Go to General Configuration > Office Connector
 Set the configuration options as described in the table below.

Option	Default Value	Description		
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	Note: This feature requires a third party app. 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 Content Formatting for Confluence app from Adaptavist to get this macro.		
Maximum imported image size	1200x900	Users will be prevented from importing a Word document if it contains images that exceed the maximum imported image size. The size is based on the size the image is displayed in Word, not the resolution of the image. Defaults to 1200 wide x 900 high.		
Temporary storage for viewfile macro	The Confluen ce 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 directories.properties file You can specify a location by editing the Office Connector's directories.properties file: 1. Locate the OfficeConnector-x.xx.jar file (where x.xx is the version number) in your Confluence Home directory and copy it to a temporary location 2. Unzip the JAR file and find the resources/directories.properties file. 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.benryan.confluence. word.edit.cacheDir= 3. Edit the last line, adding the path to your required temporary location directly after the '=' character. For example: On Windows:com.benryan.confluence.word.edit. cacheDir=c: \my\path\ On Linux:com.benryan.confluence.word.edit. cacheDir=/home /myusername/my/path 4. Save the file, recreate the JAR and put it back in your Confluence Home directory, overwriting the original JAR.		
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	2	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.
File editing method:E nable Edit in Office for all users (this will disable the Companion app)	Disabled	This allows administrators to disable the Companion app method for editing files, and instead use the Office Connector to edit compatible files.
File authentication:In clude authentication tokens in the URL path (required for Office 2013)	Disabled	If this option is enabled, the Office Connector will use authentication tokens in the URL. This needs to be enabled to edit Office 2013 documents.

Managing Webhooks

Webhooks allow you to notify an application, or other external service, when certain events occur in Confluence. For example, you can configure webhooks to update an issue tracker, or trigger notifications in a chat tool.

On this page:

- Securing the webhook
- Create a new webhook
- Triggering webhooks
- Event payloads
- Circuit breaking

A webhook consists of:

- One or moreevents such as page creation, or space removed. Youcan select multiple eventsto trigger the webhook.
- AURL the endpointwhere you want Confluence to send the event payloads when a matching event happens.

Once created, Confluence will listen for these events, and send the event payload, in JSON format, to the URL you specified.

Securing the webhook

Confluence uses webhook secrets to authenticate the payload. Combined with HTTPS, it helps ensure the message transmitted is the one that Confluence intended to send, and that the contents were not tampered with.

When you define a secret for a webhook, each request is signed via a Hash-based Message Authentication Code (HMAC). The default for this algorithm is HMACSha256. The header X-Hub-Signature is defined and contains the HMAC.

To authenticate the validity of the message payload, the receiver can perform the HMAC algorithm on the received body with the secret as the key to the HMAC algorithm. If the results don't match, it may indicate there was a problem with transmission that has caused the message payload to change.

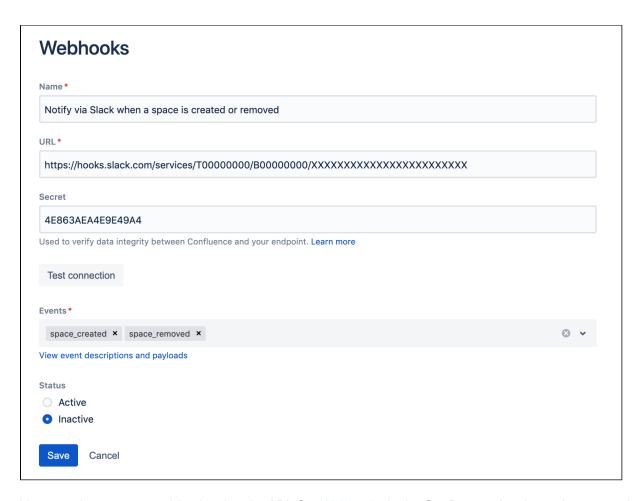
Create a new webhook

You need Confluence Administrator or System Administrator global permissions to create a webhook.

To create a new webhook:

- 1. Go to General Configuration > Webhooks.
- 2. Enter a title for your webhook.
- 3. Enter the **URL** of the application or server.
- 4. Enter a **secret**. This is a string of up to 255 characters that you define.
- 5. Select **Test connection** to check you can reach the application or server.
- 6. Choose the events that should trigger the webhook.
- 7. Select **Active** to make your webhook available immediately.
- 8. Select Create.

Screenshot: Creating a webhook to notify a chat application when a space is created or removed.



You can also create a webhook using the API. SeeWebhooks in the Confluence developer documentation.

Triggering webhooks

You can configure your webhook to be triggered by the following events.

Event	Triggered when		
attachment_created	a file is attached to a page or blog post		
attachment_remov	a file is deleted (sent to the trash) from the attachments page		
eu	(not triggered when a version is deleted from the file history)		
attachment_restored	a file is restored from the trash		
attachment_trashed	a file is purged from the trash		
attachment_updated	a new file version of is uploaded directly or by editing the file		
blog_created	a blog post is published		
blog_removed	a blog post is deleted (sent to the trash)		
blog_restored	a blog post is restored from the trash		
blog_trashed	a blog post is purged from the trash		
blog_updated	a blog post is edited		
blueprint_page_cre ated	a page is created from a blueprint (such as meeting notes, decision, or how-to)		

comment_created	a page comment, inline comment or file comment is made
comment_removed	a page comment, inline comment, or file comment is deleted
comment_updated	a page comment, inline comment, or file comment is edited
content_created	a page, blog post, attachment (file), comment (page, inline, or file), or other file (such as a space logo) is created or uploaded.
content_restored	a page, blog post, or attachment (file) is restored from the trash
content_trashed	a page, blog post, or attachment (file) is purged from the trash
content_updated	a page, blog post, attachment (file), or comment (page, inline, and file) is edited.
content_permission s_updated	a view or edit restriction is applied or removed from a page or blog post
group_created	a new group is created
group_removed	a group is deleted
label_added	an existing label is applied to a page, blog post, or space
label_created	a label is added for the first time (did not already exist)
label_deleted	a label is removed from the last page, blog post, or space, and so ceases to exist
label_removed	a label is removed from a page, blog post, or space
page_children_reor dered	the default ordering of pages is changed to alphabetical in the Space Tools > Reorder pages tab
	(is not triggered when you drag a page, or move a page, to change the page order)
page_created	a page is published for the first time, including pages created from a template or blueprint
page_moved	a page is moved to a different position in the page tree, to a different parent page, or to another space
	1
page_removed	a page is deleted (sent to the trash)
page_removed page_restored	
	a page is deleted (sent to the trash)
page_restored	a page is deleted (sent to the trash) a page is restored from the trash
page_restored page_trashed	a page is deleted (sent to the trash) a page is restored from the trash a page is purged from the trash
page_restored page_trashed page_updated	a page is deleted (sent to the trash) a page is restored from the trash a page is purged from the trash a page is edited (triggered at the point the unpublished changes are published)
page_restored page_trashed page_updated space_created space_logo_updat ed space_permissions	a page is deleted (sent to the trash) a page is restored from the trash a page is purged from the trash a page is edited (triggered at the point the unpublished changes are published) a new space is created
page_restored page_trashed page_updated space_created space_logo_updated	a page is deleted (sent to the trash) a page is restored from the trash a page is purged from the trash a page is edited (triggered at the point the unpublished changes are published) a new space is created a new file is uploaded for use as the logo of a space
page_restored page_trashed page_updated space_created space_logo_updat ed space_permissions	a page is deleted (sent to the trash) a page is restored from the trash a page is purged from the trash a page is edited (triggered at the point the unpublished changes are published) a new space is created a new file is uploaded for use as the logo of a space space permissions are changed in the Space Tools > Permissions tab
page_restored page_trashed page_updated space_created space_logo_updat ed space_permissions _updated	a page is deleted (sent to the trash) a page is restored from the trash a page is purged from the trash a page is edited (triggered at the point the unpublished changes are published) a new space is created a new file is uploaded for use as the logo of a space space permissions are changed in the Space Tools > Permissions tab (is not triggered when you edit space permissions using Inspect Permissions)
page_restored page_trashed page_updated space_created space_logo_updat ed space_permissions _updated space_removed	a page is deleted (sent to the trash) a page is restored from the trash a page is purged from the trash a page is edited (triggered at the point the unpublished changes are published) a new space is created a new file is uploaded for use as the logo of a space space permissions are changed in the Space Tools > Permissions tab (is not triggered when you edit space permissions using Inspect Permissions) a space is deleted the space details (title, description, status) is updated in the Space Tools >

user_created	a new user account is created
user_deactivated	a user account is disabled
user_followed	someone follows a user
user_reactivated	a disabled user account is enabled
user_removed	a user account is deleted

Event payloads

Here's an example of the event payload for the page_trashed event. This is the raw data that's sent, in JSON format, to your endpoint.

```
{
  "timestamp":1596182511300,
  "event":"page_trashed",
  "userKey":"ff80818154ec9913015501e194f601d8",
  "page":{
      "id":309264476
  }
}
```

You'll note that the content is comprised mostly of IDs. This is to ensure that identifiable information is not stored by third party services, or leaked to users who do not have permission to see it.

Once received, you can use the REST API to interpret these IDs. SeeConfluence Server Rest API.

Circuit breaking

To help protect your Confluence site, any webhooks that fail consistently, are skipped for a period of time. By default, if a webhook fails five times, it is considered unhealthy and is skipped, initially for 10 seconds. If it continues to fail, it will be gradually shipped for longer periods, up to 10 hours.

A webhook may also be skipped if there are too many webhooks in flight. If there are 500 webhooks being invoked, further requests will be skipped until the number in flight drops below 500.

Managing your Confluence License

Your license entitles you to run Confluence and be eligible for support and upgrades for a specified period. It also defines the number of users who are entitled to use Confluence.

To quickly check the status of your license you can go to General Configuration > Troubleshooting and support tools.

You'll need need Confluence Administrator or System Administrator permissions to view and edit your license.

Viewing your license details

To view your Confluence license:

- 1. Go to Configuration.
- 2. ChooseLicense Detailsin the left-hand panel.

The License Details page tells you:

On this page:

- Viewing your license details
- Updating your license
- Understanding the user count for your license
- Exceeding your licensed user count
- Reducing your user count
- Downgrading your license
- Finding your Support Entitlement Number (SEN)
- What happens when your maintenance or subscription expires

Related pages:

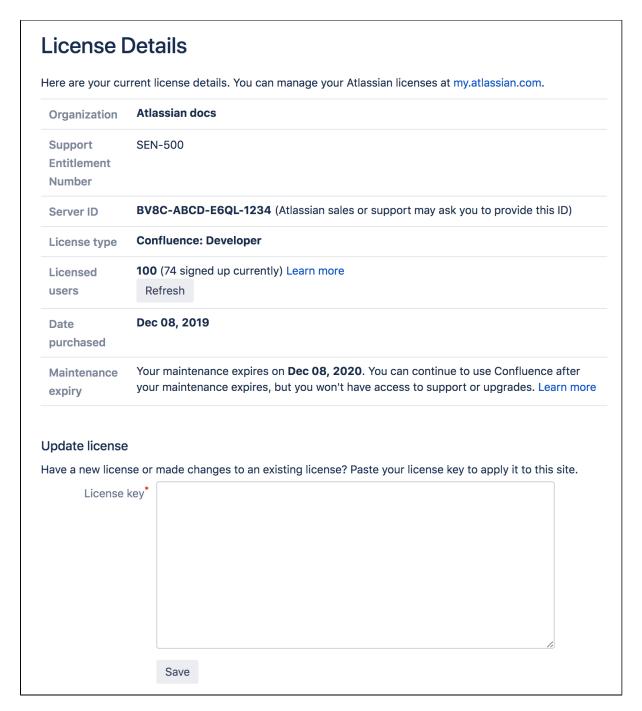
- Upgrading Beyond Current Licensed Period
- Confluence installation and upgrade guide
- Confluence administrator's guide
- The type of license (for example: Commercial, Academic, Community, or Evaluation).
- Number of users you are licensed for, and how many are currently in use.
- Your license expiry date, for support and upgrade eligibility.
- Your server ID which is generated when you install Confluence for the first time and remains the same for the life of the installation (including after upgrades or changes to your license).
- Your support entitlement number (SEN).

Updating your license

If you change your license (for example to a license with more users), or migrate from Confluence Cloudand you will need to update your license.

To update your Confluence license:

- 1. Go to General Configuration>License Details
- 2. Enter your new license in the License field.
- 3. ChooseSave.



⚠If you run Data Center in a cluster, you may need to apply the license to each node individually, if it does not automatically propagate to all nodes.

Understanding the user count for your license

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

The License Details page will indicate the number of users currently signed up (your registered user count). It:

- includes only users who have the 'can use'global permissionsfor the Confluence site.
- does not include anonymous users, who may access your Confluence site if you haveallowed anonymous access.
- does not includedeactivated users.

Exceeding your licensed user count

If you exceed the number of users included in your license, your Confluence instance will become read-only, that means no users will be able to create or edit content until you reduce the number of users.

Reducing your user count

You can reduce your user count by removing or deactivating users who do not require access to Confluence. SeeDelete or Disable Users.

If you haveconnected Confluence to an LDAP directory, you may want configure Confluence to only synchronize a subset of users from LDAP rather than all users. SeeHow to change the number of users synchronized from LDAP to Confluencein the Knowledge Base. This can be a complicated process and we recommend that you only use this method if necessary.

Downgrading your license

If you decide to downgrade your Confluence license to pay for fewer users you need to ensure that the number of users currently signed up (as shown on the License Details page) is lower that the number allowed by your new license before your apply the new license.

If you have more users than your new license allows you will need to reduce your user count before applying the new license.

Finding your Support Entitlement Number (SEN)

You can find your Support Entitlement Number (SEN) in three places:

- In Confluence go to General Configuration > License Details)
- At my.atlassian.com
- On your Atlassian invoice.

See How to find your Support Entitlement Number (SEN) for more general information about how Atlassian Support uses this number.

What happens when your maintenance or subscription expires



(i) Important changes to our server and Data Center products

Weve ended sales for new server licenses, and will end support for server on February 2, 2024. Were continuing our investment in Data Center with several key improvements. Learn what this means for you

Server licenses

If you have a Confluence Server license, your maintenance entitles you to access Atlassian support, and upgrade Confluence.

When your maintenance expires, you can still use Confluence, but you'll no longer be able to contact Atlassian support, or upgrade to a version of Confluence released after your maintenance expiry date.

Data Center licenses

Confluence Data Center is offered as a subscription (also known as a fixed term license), which includes access to support and version upgrades.

If your subscription expires, Confluence will become read-only, which means you'll be able to view pages, but not create or edit them.

Our licensing policy can change from time to time, so it's best to check our Purchasing and Licensing FAQ for the latest information.

Need more information about your Server or Data Center license? Get in touch

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 or Pages
- Data Import and Export
- Import a Text File
- Auditing in Confluence
- Set retention rules to delete unwanted data
- Data pipeline

Check the amount of data in your site

To check the amount of data in your site:

- 1. Go to General Configuration > System information.
- 2. Scroll down to the Confluence usage section.

This will tell you the number of:

- Total Spaces total number of spaces, including site spaces and personal spaces
- Site Spaces
- Personal Spaces
- Content (All Versions) total number of page and blog post versions (includes the latest version and all historical versions)
- Content (Current Versions) total number of pages and blog posts (only counts the latest version)
- Local Users
- · Local Groups.

Clean up unwanted data

As your team grows, so does the data being created and stored in Confluence. Find outthe ways you can declutter and reduce the overall size of your Confluence site in our Cleanup guide.

Related pages:

- Managing System and Marketplace Apps
- Integrating Confluence with Other Applications
- Getting Started as Confluence Administrator
- Confluence administrator's guide

Database Configuration

This document provides information on connecting Confluence to an external database.

Choosing an external database

Note: Take time to choose your database wisely. The XML backup built into Confluence is not suited for migration or backup of large data sets. If you need to migrate later, you may need to use a third party database migration tool.

Below is more information on selecting and migrating to an external database:

- Migrating to a Different Database
- Supported Databases
- Database Troubleshooting

On this page:

- Choosing an external database
- About the embedded H2 database
- Database setup
- Database drivers
- Database connection methods
- Database troubleshooting

Related pages:

- Database JDBC Drivers
- Supported Platforms
- Embedded H2 Database
- Managing Confluence Data

About the embedded H2 database

The embedded H2 database is**only**supported for testing and app development purposes onnon-clustered (single node)Confluence Data Center installations.

To find out if you are still using the embedded database, go to General Configuration > Troubleshooting and support tools.

Database setup

To find out how to set up your database, see:

- Database Setup for Oracle
- Database Setup For MySQL
- Database Setup for PostgreSQL
- Database Setup for SQL Server
- Configuring Confluence Data Center to work with Amazon Aurora

Database drivers

You must use a supported database driver. SeeDatabase JDBC Drivers for the drivers we support.

If you attempt to use an unsupported or custom JDBC driver (or adriverClassNamefrom an unsupported or custom driver in your JINDI datasource connection) collaborative editing will fail.

Database connection methods

You can connect Confluence to your database using a JDBC URL or a JNDI datasource.

By default the setup wizard only provides the option to use a JDBC connection, as this is the recommended connection method.

If you want to use a JNDI datasource,seeConfiguring a datasource connectionfor the steps you'll need to take before you set up Confluence, asthe setup wizard will only provide the option to use a datasource if it detects a datasource in your Tomcat configuration.

Database troubleshooting

For database-related problems seeDatabase Troubleshooting.

If you need more help, check outTroubleshooting Problems and Requesting Technical Support.

Database JDBC Drivers

This page provides the download links for the JDBC drivers for all supported databases.

Due to licensing constraints, we are not able to bundle MySQL or Oracle database drivers with Confluence, so you will need to manually download and install the driver listed below before you can set up Confluence.

If you use PostgreSQL or Microsoft SQL Server, the drivers are bundled with Confluence, so you're ready to.

Adding your database driver (MySQL and Oracle)

The Confluence setup wizard will stop you at the Database configuration step if it can't find an appropriate driver for the database you select.

To make your database driver available to Confluence:

- 1. Stop Confluence.
- 2. Download and extract the appropriate driver from the list below.
- 3. Drop the .jar file in your<installation-directory> /confluence/WEB-INF/lib directory.
- 4. Restart Confluence then go to http://localhost:<port>in your browser to continue the setup process.

The setup wizard will return to the database configuration step, and you're back on your way.

Supported drivers

Database	Driver bundled?	JDBC drivers	Notes	More information
PostgreSQL	•	Postgres JDBC driver download (latest)	We recommend that you use the bundled JDBC 4 driver. If you want to use a later driver, you can download it from the PostgreSQL website.	Database Setup for PostgreSQL
Microsoft SQL Server	•	Microsoft JDBC Driver for SQL Server download	We recommend that you use the bundled Type 4 JDBC driver. If you decide to use a later version, we may not be able to provide support for any problems you encounter.	Database setup for Microsoft SQL Server
	•	jTDS 1.3.1 driver download	This driver is deprecated. New Confluence installations use the Microsoft JDBC Driver for SQL Server (above). If you're upgrading an existing Confluence site to Confluence 6.4 you should continue to use the bundled jTDS driver. We'll help you migrate to the Microsoft driver in a later release.	

On this page:

- Adding your database driver (MySQL and Oracle)
- Supported drivers

Related pages:

- Database Configuration
- Supported Platforms

MySQL 5.7	⊗	Connector\J 5.1 driver download	Due to licensing constraints, MySQL drivers are not bundled with Confluence. Confluence is currently tested with the 5.1.48 driver. You can't use the latest driver (8.x) with Confluence and MySQL 5.7.	Database setup for MySQL
MySQL 8.0	8	Connector\J 8.0 driver download	Due to licensing constraints, MySQL drivers are not bundled with Confluence. Confluence is currently tested with the 8.0.22 driver.	Database setup for MySQL
Oracle	⊗	JDBC driver downloads	Due to licensing constraints, Oracle drivers are not bundled with Confluence. For Oracle 12c R2 use the 12.2.0.x driver (ojdbc8.jar) For Oracle 19c you can use either ojdbc8.jar or ojdbc10.jar. We recommend using the thin drivers only. See the Oracle JDBC driver FAQ.	Database setup for Oracle

If you attempt to use an unsupported or custom JDBC driver (or a driverClassNamefrom an unsupported or custom driver in your JINDI datasource connection) collaborative editing will fail. You must use a supported driver.

Database Setup for Oracle

This page provides instructions for configuring Confluence to use an Oracle database.

Before you start

- SeeSupported Platformstocheck your version of Oracle is supported. You may need to upgrade your database before installing Confluence.
- If you're switching from another database, including the embedded evaluation database, readMigrating to Another Databasebefore you begin.
- You'll need an experienced Oracle database administrator (DBA) to set up and maintain your database.

Our support team can assist with Confluence problems, but are unable to help you administer your Oracle database.

If you don't have access to an experienced Oracle DBA, consider using a different supported database.

On this page:

- Before you start
- 1. Install Oracle
- 2. Create database user
- 3. Install Confluence
- 4. Download and install the Oracle thin driver
- 5. Enter your database details
- 6. Test your database connection
- Troubleshooting

Related pages:

- Database Configuration
- Known Issues for Oracle
- Confluence installation and upgrade guide

1. Install Oracle

If you don't already have an operational Oracle server, download and install it now. See the Oracle documentation for instructions.

When setting up your Oracle server:

- Character encoding must be set to AL32UTF8 (this the Oracle equivalent of Unicode UTF-8).
- Collation should be set to BINARY.

2. Create database user

To create the user and assign its privileges:

1. Use the sqlplus command to access Oracle via the command line

```
sqlplus user/password <as sysdba|as sysoper>
```

If you're logging in with the user 'sys' you'll need to include the "as sysdba" or "as sysoper" to determine which sys role you want to use.

2. Create a Confluence user (for example confluenceuser), and grant the following only to that user:

```
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>;
```

- It isvery important that the user is granted the exact privileges indicated above. Confluence requires only these privileges soyou should grant specific privileges to the user create table, creat e sequence, and create triggers houldn't be assigned as part of a role.
- Do not grant the user theselect any tablepermission. That permission can cause problems with other schemas.
- When you create a user, specify thetablespacefor the table objects as shown above.

3. Install Confluence

Check out the Confluence Installation Guidefor step-by-step instructions on how to install Confluence on your operating system.

4. Download and install the Oracle thin driver

Due to licensing restrictions, we're not able to bundle an Oracle driver with Confluence. To make your database driver available to Confluence:

- 1. Stop Confluence.
- 2. Head to Database JDBC Drivers and download the appropriate driver. The driver file will be called something likeojdbc8.jar
- 3. Drop the .jar file in your < installation-directory > /confluence / WEB-INF / libdirectory.
- 4. Restart Confluence then go tohttp://localhost:<port>in your browser to continue the setup process.

5. Enter your database details

The Confluence setup wizard will guide you through the process of connecting Confluence to your database.

Use a JDBC connection (default)

JDBC is the recommended method for connecting to your database.

The Confluence setup wizard will provide you with two setup options:

- Simple- this is the most straightforward way to connect to your database.
- By connection string- use this option if you want to specify additional parameters and are comfortable constructing a database URL.

Depending on the setup type, you'll be prompted for the following information.

Setup type	Field	Description
Simple	Hostna me	This is the hostname or IP address of your database server.
Simple	Port	This is the Oracle port. If you didn't change the port when you installed Oracle, it will default to 1521.
Simple	Service name	This is the service name (of your confluence database.
By connection string	Databas e URL	The database URL is entered in this format: jdbc:oracle:thin:@// <host>:<port>/<service></service></port></host>
		<pre><service> can be either the SID or Service Name. For example: jdbc: oracle:thin:@//localhost:1521/confluence</service></pre>
		By default, we use the new style URL provided by the thin driver. You can also use the tnsnames style.

Both	Userna me	This is the username of your dedicated database user. In the example above, this is confluenceuser.
Both	Passwo rd	This is the password for your dedicated database user.

To determine the host, port, service name, and/or SID, execute the following command as the user running Oracle (usually 'Oracle'):

```
lsnrctl status
```

Here's an example of the 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
                       ON: Local OS Authentication
Security
SNMP
                        OFF
                       XE
Default Service
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
```

- The host and port are determined by the line containingPROTOCOL=tcp(the line withoutPresent ation=HTTP).
- UnderServices Summary, each service which has an instance with READY status is a
 connectable service. The name followingService a service name for connecting to the
 database name followingInstanceon the next line.
- The SID is the name of the database instance, as defined by the \$ORACLE_SIDvariable when you
 have sourced the Oracle environment to your shell.

For example, if you are running Confluence on the same server as the Oracle database, with the abovels nrctlstatus 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 a datasource.

Seethe Oracle JDBC FAQ for more information on Oracle JDBC URLs.

Use a JNDI datasource

If you want to use a JNDI datasource,seeConfiguring a datasource connectionfor the steps you'll need to take before you set up Confluence, asthe setup wizard will only provide the option to use a datasource if it detects a datasource in your Tomcat configuration.

6. Test your database connection

In the database setup screen, hit the **Test connection** button to check:

- that Confluence can connect to your database server
- that the database character encoding is correct
- that your database user has appropriate permissions for the database
- that your database user has NOT been granted the SELECT ANY TABLE privilege

Once the test is successful, hit Next to continue with the Confluence setup process.

Troubleshooting

- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- The following page contains common issues encountered when setting up your Oracle database to work with Confluence:Known Issues for Oracle.
- There's a known issue when username or schema names contain dots. See
 CONFSERVER-60274 GATHERING IMPACT for more information.

Database Setup for PostgreSQL

This page provides instructions for configuring Confluence to use a PostgreSQL database.

Before you start

- SeeSupported Platformstocheck your version of PostgreSQL is supported. You may need to upgrade your database before installing Confluence.
- If you're switching from another database, including the embedded evaluation database, readMigrating to Another Databasebefore you begin.

On this page:

- Before you start
- 1. Install PostgreSQL
- 2. Create a database user and database
- 3. Install Confluence
- 4. Enter your database details
- 5. Test your database connection
- Troubleshooting

Related pages:

- Database Configuration
- Known issues for PostgreSQL

1. Install PostgreSQL

If you don't already have PostgreSQL installed, download and install it now.

A few tips when installing PostgreSQL:

- The **password** you provide during the installation process is for the 'postgres' account, which is the database root-level account (the super user). Remember this username and password asyou'll 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, make sure it does not conflict with any other services running on that port.
- Choose the **locale** that best matches your geographic location.
- Don't launch **Stack Builder** at the completion of the installer.

Create a database user and database

Once you've installed PostgreSQL:

- 1. Create a database user, for example confluence user.
 - Your new user must be able tocreate database objects and must have can login permission.
- 2. Next, create a database (for example confluence):
 - Owneris your new database user (for exampleconfluenceuser)
 - Character encodingmust be set toutf8encoding.
 - Collation must also be set to utf8. Other collations, such as "C", are known to cause issues with Confluence.
 - If you are running PostgreSQL on Windows use the equivalent character type and collation for your locale, for exampleEnglish_United States.1252
 - In Linux systems, if the locale is not utf8, include LC_CTYPE as utf8 during database creation.

You can usepgAdminas an alternative to the command line to complete this step.

3. Install Confluence

Check out the Confluence Installation Guidefor step-by-step instructions on how to install Confluence on your operating system.

4. Enter your database details

The Confluence setup wizard will guide you through the process of connecting Confluence to your database. Be sure to select "**My own database**".

Use a JDBC connection (default)

JDBC is the recommended method for connecting to your database.

The Confluence setup wizard will provide you with two setup options:

- **Simple** this is the most straightforward way to connect to your database.
- By connection string- use this option if you want to specify additional parameters and are comfortable constructing a database URL.

Depending on the setup type, you'll be prompted for the following information.

Setup type	Field	Description
Simple	Hostname	This is the hostname or IP address of your database server.
Simple	Port	This is the PostgreSQL port. If you didn't change the port when you installed Postgres, it will default to 5432.
Simple	Database name	This is the name of your confluence database. In the example above, this is confluence
By connection string	Database URL	The database URL is entered in this format: jdbc:postgresql:// <server>:<port>/<database> For example: jdbc:postgresql://localhost:5432/confluence If you need to connect to an SSL database, add thesslmode=requireparam eter in the database URL. For example: jdbc:postgresql://localhost:5432/confluence? sslmode=require</database></port></server>
Both	Username	This is the username of your dedicated database user. In the example above, this is confluenceuser.
Both	Password	This is the password for your dedicated database user.

Use a JNDI datasource

If you want to use a JNDI datasource,seeConfiguring a datasource connectionfor the steps you'll need to take before you set up Confluence, asthe setup wizard will only provide the option to use a datasource if it detects a datasource in your Tomcat configuration.

5. Test your database connection

In the database setup screen, hit the **Test connection** button to check:

- that Confluence can connect to your database server
- that the database character encoding is correct
- that your database user has appropriate permissions for the database

Once the test is successful, hit **Next** to continue with the Confluence setup process.

If Confluence and PostgreSQL are hosted on different servers, seethePostgreSQL documentation on how to set up pg_hba.confto make sure Confluence and PostgreSQL can communicate remotely.

Troubleshooting

- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If you're 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.

Database Setup for SQL Server

This page provides instructions for configuring Confluence to use a Microsoft SQL Server database.

Before you start

Check the following before you start:

- SeeSupported Platformstocheck your version of SQL Server is supported. You may need to upgrade your database before installing Confluence.
- If you're switching from another database, including the embedded evaluation database, readMigrating to Another Databasebefore you begin.

On this page:

- Before you start
- 1. Install SQL Server
- 2. Create a database and database user
- 3. Install Confluence
- 4. Enter your database details
- 5. Test your database connection
- Database driver changes
- Troubleshooting

Related pages:

- Database Configuration
- Known issues for SQL Server
- Confluence installation and upgrade guide

1. Install SQL Server

If you don't already have Microsoft SQL Server installed, download and install it now. SeeInstallation for SQL Server on MSDN for step-by-step instructions.

SQL Server allows two types of authentication: SQL Server Authentication and Windows Authentication. To make sure Confluence will be able to connect to your database you'll need to set your SQL server to allow Mixed Authentication (both SQL Server and Windows modes). This setup is generally found under Properties > Security > Server Authentication.

Create a database and database user

Once you've installed SQL Server, create a database user and database for Confluence as follows:

- 1. Using your SQL administrator permissions, create a new database (for exampleconfluence)
- 2. Set the default collation for the database to SQL_Latin1_General_CP1_CS_AS(case sensitive).

```
ALTER DATABASE <database-name> COLLATE SQL_Latin1_General_CP1_CS_AS
```

If you see 'database could not be exclusively locked to perform the operation' error, 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;
```

3. Check the database isolation level of READ_COMMITTED_SNAPSHOT is ON.

```
SELECT is_read_committed_snapshot_on FROM sys.databases WHERE name= 'database-name'
```

If this query returns 1, then READ COMMITTED SNAPSHOT is ON, and you're good to go.

If this query returns **0**,READ_COMMITTED_SNAPSHOT option is OFF and you will need to turn it on as follows:

```
ALTER DATABASE <database-name>
SET READ_COMMITTED_SNAPSHOT ON
WITH ROLLBACK IMMEDIATE;
```

- 4. Using your SQL administrator permissions, create a new SQL user account for Confluence (for example, confluenceuser).
- 5. Give this user the default schema as follows:

```
ALTER USER <confluenceuser> WITH DEFAULT_SCHEMA = dbo
```

6. Make sure this user has full create, read and write permissions for the database tables. Confluence must be able to create its own schema, and have the ability to create/drop triggers and functions. Refer to the SQL Server documentation for more information.

3. Install Confluence

Check out the Confluence Installation Guidefor step-by-step instructions on how to install Confluence on your operating system.

4. Enter your database details

The Confluence setup wizard will guide you through the process of connecting Confluence to your database.

Use a JDBC connection (default)

JDBC is the recommended method for connecting to your database.

The Confluence setup wizard will provide you with two setup options:

- **Simple** this is the most straightforward way to connect to your database.
- By connection string- use this option if you want to specify additional parameters and are comfortable constructing a database URL.

Depending on the setup type, you'll be prompted for the following information.

Setup type	Field	Description
Simple	Hostname	This is the hostname or IP address of your database server.
Simple	Port	This is the SQL Server port. If you didn't change the port when you installed SQL Server, it will default to 1433.
Simple	Database name	This is the name of your confluence database. In the example above, this is confluence
Simple	Instance name	To find out your instance name, connect to your database and run one of the following:
		select @@SERVICENAME;
		SELECT SERVERPROPERTY('InstanceName');
		If you have a default named instance setup in SQL Server, you won't need to specify this parameter.

By connection string	Database URL	The database URL is entered in this format: jdbc:sqlserver:// <hostname>:<port>; databaseName=<database> For example: jdbc:sqlserver://yourserver:1433;databaseName=confluence</database></port></hostname>
Both	Username	This is the username of your dedicated database user. In the example above, this is confluenceuser.
Both	Password	This is the password for your dedicated database user.

Use a JNDI datasource

If you want to use a JNDI datasource,seeConfiguring a datasource connection for the steps you'll need to take before you set up Confluence, as the setup wizard will only provide the option to use a datasource if it detects a datasource in your Tomcat configuration.

5. Test your database connection

In the database setup screen, hit the **Test connection** button to check:

- Confluence can connect to your database server
- the database collation and isolation level is correct
- your database user has appropriate permissions for the database

Once the test is successful, hit Next to continue with the Confluence setup process.

Database driver changes

In Confluence 6.6 we replaced the open source jTDS driver for Microsoft SQL Server with the official Microsoft JDBC Driver for SQL Server. You will be automatically migrated to the new driver when you upgrade to 6.6 or later.

If for some reason the automatic migration fails (for example, if you're using a datasource connection), you'll need to make this change manually. See Migrate from the jTDS driver to the supported Microsoft SQL Server driver in Confluence 6.4 or later.

Troubleshooting

• If you get the following error message, check you've 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)'
```

- You may need to open additional ports. See this Microsoft KBabout the ports required for SQL Server.
- The following page contains common issues encountered when setting up your SQL Server database to work with Confluence:Known Issues for SQL Server.

Database Setup For MySQL

This page provides instructions for configuring Confluence to use a MySQL database.

Before you start

- SeeSupported Platformstocheck your version of MySQL is supported. You may need to upgrade your database before installing Confluence.
- If you're switching from another database, including the embedded evaluation database, readMigrating to Another Databasebefore you begin.
- Confluence will not work on MySQL variants such as MariaDB (CONFSERVER-29060) and Percona Server (CONFSERVE R-36471)

On this page:

- Before you start
- 1. Install MySQL Server
- 2. Configure MySQL Server
- 3. Create database and database user
- 4. Install Confluence
- 5. Download and install the MySQL driver
- 6. Enter your database details
 - Use a JDBC connection (default)
 - Use a JNDI datasource
- 7. Test your database connection
- Upgrade your database and driver
- Troubleshooting

Related pages:

- Database Configuration
- Database Troubleshooting for MySQL
- Confluence installation and upgrade guide

1. Install MySQL Server

If you don't already have MySQL installed, downloadand install it now. See the MySQL documentation for step-by-step instructions.

2. 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.

1 These instructions apply to Confluence 7.3 and later. Using an earlier version? See Database Setup For MySQL in Confluence 7.2 and earlier.

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 thetype 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 useInnoDB as its default storage engine.
 - You must use the InnoDB storage engine with Confluence. Using the MyISAM storage enginec an lead to data corruption in Confluence.
 - e. Set the InnoDB Tablespace settings to your requirements. (The default settings are acceptable.)

- f. Set the approximatenumber 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 internationalization. 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 themy.cnffile (my.inion Windows operating systems) in your MySQL server. Locate the [mysqld] section in the file, and add or modify the following parameters:

(Refer toMySQL Option Filesfor detailed instructions on editingmy.cnfandmy.ini.) Locate the [mysqld] section in the file, and add or modify the following parameters:

Specify the default character set to be utf8mb4:

```
[mysqld]
...
character-set-server=utf8mb4
collation-server=utf8mb4_bin
...
```

Set the default storage engine to InnoDB:

```
[mysqld]
...
default-storage-engine=INNODB
...
```

Specify the value of max_allowed_packet to be at least 256M:

```
[mysqld]
...
max_allowed_packet=256M
...
```

Specify the value ofinnodb_log_file_sizeto be at least 2GB:

```
[mysqld]
...
innodb_log_file_size=2GB
...
```

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
```

 Ensure that the global transaction isolation level of your Database had been set to READ-COMMITTED.

```
[mysqld]
...
transaction-isolation=READ-COMMITTED
...
```

• Check that the binary logging format is configured to use 'row-based' binary logging, and that your database user can create and alter stored functions.

```
[mysqld]
...
binlog_format=row
log-bin-trust-function-creators = 1
...
```

• If you're using MySQL 5.7, turn off the 'derived merge' optimizer switch, as this can cause the dashboard to load slowly.

```
optimizer_switch = derived_merge=off
```

- 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'.

Create database and database user

Once you've installed and configured MySQL, create a database user and database for Confluence as follows:

- 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 (for example confluence):

```
CREATE DATABASE <database-name> CHARACTER SET utf8mb4 COLLATE utf8mb4_bin;
```

3. Create a Confluence database user (for example confluenceuser):

```
GRANT ALL PRIVILEGES ON <database-name>.* TO '<confluenceuser>'@'localhost' IDENTIFIED BY '<password>';
```

If Confluence is not running on the same server, replace localhost with the hostname or IP address of the Confluence server.

4. Install Confluence

Check out the Confluence Installation Guidefor step-by-step instructions on how to install Confluence on your operating system.

5. Download and install the MySQL driver

Due to licensing restrictions, we're not able to bundle the MySQL driver with Confluence. To make your database driver available to Confluence follow the steps below for your MySQL version.

MySQL 5.7

- 1. Stop Confluence.
- 2. Head to Database JDBC Drivers and download the appropriate driver. The driver file will be called something likemysql-connector-java-5.1.xx-bin.jar
- 3. Drop the .jar file in your<installation-directory>/confluence/WEB-INF/libdirectory.
- 4. Restart Confluence then go tohttp://localhost:<port>in your browser to continue the setup process.

MySQL 8.0

You can't use MySQL 8.0 with Confluence 7.1 or earlier.

- 1. Stop Confluence.
- 2. Head to Database JDBC Drivers and download the appropriate driver for MySQL 8. The driver file will be called something likemysql-connector-java-8.0.xx-bin.jar
- 3. Drop the .jar file in your < installation-directory > /confluence / WEB-INF / libdirectory.
- 4. Restart Confluence then go tohttp://localhost:<port>in your browser to continue the setup process.

6. Enter your database details

The Confluence setup wizard will guide you through the process of connecting Confluence to your database.

Use a JDBC connection (default)

JDBC is the recommended method for connecting to your database.

The Confluence setup wizard will provide you with two setup options:

- **Simple-** this is the most straightforward way to connect to your database.
- By connection string- use this option if you want to specify additional parameters and are comfortable constructing a database URL.

Depending on the setup type, you'll be prompted for the following information.

Setup type	Field	Description
Simple	Hostname	This is the hostname or IP address of your database server.
Simple	Port	This is the MySQL port. If you didn't change the port when you installed MySQL, it will default to 3306.
Simple	Database name	This is the name of your confluence database. In the example above, this is confluence
By connection string	Database URL	The database URL is entered in this format: jdbc:mysql:// <hostname>:<port>/<database> For example: jdbc:mysql://localhost:3306/confluence</database></port></hostname>
Both	Username	This is the username of your dedicated database user. In the example above, this is confluenceuser.
Both	Password	This is the password for your dedicated database user.

Use a JNDI datasource

If you want to use a JNDI datasource,seeConfiguring a datasource connectionfor the steps you'll need to take before you set up Confluence, asthe setup wizard will only provide the option to use a datasource if it detects a datasource in your Tomcat configuration.

7. Test your database connection

In the database setup screen, hit the **Test connection** button to check:

- Confluence can connect to your database server
- the database character encoding, collation, isolation level and storage engine are correct
- your database user has appropriate permissions for the database.

Once the test is successful, hit Next to continue with the Confluence setup process.

Upgrade your database and driver

If you upgrade MySQL you may also need to upgrade the database driver Confluence uses to connect to your database. Always use the driver recommended on the Database JDBC Drivers page.

Before you begin, back up your database, Confluence installation directory and Confluence home directory. We strongly recommend you test your changes in astaging environment first.

To upgrade your database driver:

- 1. Stop Confluence.
- 2. Go to<installation-directory>/confluence/WEB-INF/lib/ and delete your existing driver. It will be called something likemysql-connector-java-x.x.xx-bin.jar
- 3. Drop the new driver .jar file in your<installation-directory>/confluence/WEB-INF/libdire ctory.
- 4. Upgrade your MySQL server.
- 5. Restart Confluence.

If you're using a datasource connection, you may need to also update the driver classname in the datasource.

Troubleshooting

- There is a known issue with MySQL 5.7 where parts of the dashboard can take a very long time to load. You may need to turn the "derived merge" optimizer switch off in your database configuration.
 See CONFSERVER-54984 SHORT TERM BACKLOG for details of the workaround.
- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- 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:Database Troubleshooting for MySQL

Embedded H2 Database

From Feb 2, 2021 (server end of sale date) you will only be able to generate Data Center evaluation licenses. This means the Confluence Setup Wizard wont include an option to use an embedded H2 database.

On this page:

- Connect to the embedded H2 database using DB Visualizer
- Connectto the embedded H2 databaseusing the H2 console
- Remote connections
- Migrate to a supported external database

Related pages:

- Confluence Home and other important directories
- Database Configuration

The embedded H2 database is**only**supported fortesting and app development purposesonnon-clustered (single node)Confluence Data Center installations.

To find out if you are still using the embedded database, go to General Configuration > Troubleshooting and support tools.

The embedded database files are stored in your Confluence home directory < confluence - home > /database.

Connect to the embedded H2 database using DB Visualizer

If you need to make changes directly in the database, and you're using the H2 database, here's how you can connect to it using DBVisualizer.

DBVisualizer is just one database administration tool. You can use any administration tool that supports embedded H2 databases. The steps will be similar.

- 1. Shut down Confluence.
- 2. Back up your < confluence home > / database directory.
- 3. LaunchDBVisualizer.
- 4. Choose**Create new database connection**and follow the prompts to set up the connection. The information you'll need is:
 - Database driver:H2 embedded
 - Database Userid:sa
 - Database password:leave this field blank
 - Database filename:<confluence-home>/database/h2db
 - Aleave off the . h2 . dbfile extension.
- 5. Connect to the database.

Refer to the DBV isualizer documentation for help using DBV isualizer.

Connectto the embedded H2 databaseusing the H2 console

Alternatively you can connect using the browser based H2 console. The easiest way to access the console is to double click the H2 database jar file at<installation-directory>\confluence\WEB-INF\lib\h2-x.x.x.jar.

Remote connections

Remote connections to the embedded H2 database are not permitted. You can only connect to H2 from the server on which Confluence is installed.

Plugin vendors can connect remotely when Confluence is running in dev mode, but admins should not use this as a workaround, and instead should migrate to a supported external database.

Note:The H2 database doesnt work on amulti-node Confluence cluster. A shared database is required for a multi-node cluster.

Migrate to a supported external database

If you're using the H2 database, but running Confluence as a production system, you should start planning to migrate to a supported database as soon as possible.

To migrate to a supported external database:

- 1. CheckSupported Platforms to find out which databases and versions are supported.
- 2. Head to Migrating to Another Database for a step-by-step 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:

- Moving from the embedded, trial database to a supported external database.
- · Moving from one external database to another, for example from Oracle to PostgreSQL (provided your dataset is not large)
- Upgrading to a new version of the same external database. Note: you don't need to migrate your data if you're upgrading the database in place.



If you are moving your database from one server to another you can change the JDBC URL in <confluence-home> /confluence.cfg.xml (if you are using a direct JDBC connection) or in the definition of your datasource (if you are

connecting via a datasource).

On this page:

- Limitations of database migration
- Database migration
- Migrating to an Amazon Aurora database
- Method one standard procedure
 - Step 1: Take note of your Marketplace apps
 - 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 Marketplace
 - 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 Marketplace apps
 - 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 wizardand copy your data to your new database
 - Step 7: Copy your attachments
 - Step 8. Re-install your Marketplace
 - Step 9. Check settings for new machine
- A note about case sensitivity in your database
 - Setting up a new Confluence instance
 - Migrating an existing Confluence instance to a different database
- Troubleshooting

Related pages:

- Database Configuration
- Confluence Home and other important directories

Note: The XML export built into Confluence is not suited for the backup or migration of large data sets. There are a number of third party tools that may be able to assist you with the data migration. If you would like help in selecting the right tool, or help with the migration itself, we can put you in touch with one of the Atl assian 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.

Migrating to an Amazon Aurora database

If you plan to migrate to an Amazon Aurora database, seeConfiguring Confluence Data Center to work with Amazon Aurora. This guide explains how to migrate to an Amazon Aurora cluster and connect it to Confluence Data Center.

Method one standard procedure

Step 1: Take note of your Marketplace apps

Take note of the apps (also knowns as plugins or add-ons) currently installed and enabled in Confluence, so that you can reinstate them later. Make a note of the following for each app:

- App name and vendor
- 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

- Create an XML backup of your existing data. SeeManually 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. Stop Confluence.
- 3. Make a copy of the Confluence Homedirectory. 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.
- ChooseMy own databasethen select your particular database from theDatabase typedropdown
 menu.
- When prompted to chooseMy own database, then select your newDatabase type.
- Enter your database details. Usetest connection to check your database is set up correctly.
- 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. This is the recommended method for
 large XML files.

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 Marketplace apps

Re-install any apps (also known as plugins or add-ons) that are not bundled with Confluence.

- Use the same version of the app as on your old Confluence site.
- The data created by the app will already exist in your new Confluence site, because it is included in the XML backup.

Step 8. Check settings for new machine

If you are moving Confluence to a different machine, you need to check the following settings:

- Configure your new base URL. See Configuring the Server Base URL.
- Check your application links. See Linking to Another Application.
- Update any gadget subscriptions from external sites pointing to this Confluence site. For example, if your Jira site subscribes to Confluence gadgets, you will need to update your Jira site.
- Review any other resources that other systems are consuming from Confluence.

Method two for installations with a large volume of attachments

Before you start

These instructions only apply to attachments stored in the file system. If you store attachments in the database seeAttachment Storage Configuration to find out how to migrate between different attachment storage methods.

Step 1: Take note of your Marketplace apps

Take note of the apps (also knowns as plugins or add-ons) currently installed and enabled in Confluence, so that you can reinstate them later. Make a note of the following for each app:

- App name and vendor
- 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

- Create an XML backup of your existing data. SeeManually 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. Stop 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 wizardand 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.
- ChooseMy own databasethen select your particular database from theDatabase typedropdown menu.
- When prompted to chooseMy own database, then select your newDatabase type.
- Enter your database details. Usetest connection to check your database is set up correctly.
- 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. This is the recommended method for
 large XML files.

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 Marketplace apps

Re-install any apps (also known as plugins or add-ons) that are not bundled with Confluence.

- Use the same version of the app as on your old Confluence site.
- The data created by the app will already exist in your new Confluence site, because it is included in the XML backup.

Step 9. Check settings for new machine

If you are moving Confluence to a different machine, you need to check the following settings:

- Configure your new base URL. See Configuring the Server Base URL.
- Check your application links. See Linking to Another Application.
- Update any gadget subscriptions from external sites pointing to this Confluence site. For example, if your Jira site subscribes to Confluence gadgets, you will need to update your Jira site.
- Review any other resources that other systems are consuming from Confluence.

A note about case sensitivity in your database

'Collation' refers to a set of rules that determine how data is sorted and compared. Case sensitivity is one aspect of collation. Other aspects include sensitivity to kana (Japanese script) and to width (single versus double byte characters).

Setting up a new Confluence instance

For new Confluence instances, we recommend using **case sensitive** collation for your Confluence database. This is the default collation type used by many database systems.

Note: Even if the database is configured for case sensitive collation, Confluence reduces all usernames to lower case characters before storing them in the database. For example, this means that 'joebloggs', 'joeBloggs' and 'JoeBloggs' will be treated as the same username.

Migrating an existing Confluence instance to a different database

The default Confluence configuration uses case sensitive database collation. This is typical of databases created under default conditions. If you are migrating from this type of configuration to a new database, we recommend that the new database uses case sensitive collation. If you use case insensitive collation, you may encounter data integrity problems after migration (for example, via an XML import) if data stored within your original Confluence site required case sensitive distinctions.

Troubleshooting

See our troubleshooting guideif you're unable to restore your XML backup.

Configuring Database Character Encoding

Confluence and your database must be configured to use the same character encoding.

Confluence uses UTF-8 character encoding, so your database will also need to be configured use UTF-8 (or the equivalent for your database, for example, AL32UTF8 for Oracle databases, or utf8mb4 for MySQL).

On this page:

- New installations
- Existing installations

Related pages:

- Troubleshooting Character Encodings
- Configuring Character Encoding
- Database Troubleshooting for MySQL

New installations

When installing Confluence for the first time you will need to consider character encoding:

- when creating your database
- when connecting to the database via a JDBC connection string or datasource (if you use the simple setup method in the Confluence setup wizard, we'll take care of this for you).

The Confluence setup wizard will alert you if there is a problem with your character encoding, this will make sure you don't experience problems down the track. It is much easier to solve problems now, than later when you have Confluence data in your database.

The setup guide for each of our supported databases outlines how to configure character encoding correctly when creating your database:

- Database Setup for PostgreSQL
- Database Setup For MySQL
- Database Setup for SQL Server
- Database Setup for Oracle

Existing installations

For existing Confluence sites, where the first version of Confluence installed was 6.4 or earlier, we many not have checked the collation or character encoding of your database during the initial setup.

If your database is not correctly configured to use UTF-8 character encoding (or the equivalent for your database, for exampleAL32UTF8 for Oracle databases):

- you may see a health check warning while using Confluence
- you may not be able to start Confluence after an upgrade.

If this happens, you'll need to change the character encoding for your existing database. The way you do this will depend on your database.

Also seeTroubleshooting Character Encodings for help diagnosing character encoding problems.

MySQL

SeeHow to Fix the Collation and Character Set of a MySQL Database manuallyfor details of what you'll need to do to fix the character encoding in your database. You should also make sure the collation is correct.

Microsoft SQL Server

SeeHow to fix the collation of a Microsoft SQL Server Confluence databasefor details of what you'll need to do to fix the character encoding in your database.

PostgreSQL

If you use PostgreSQL, the best option is to recreate your database.

SeeDatabase Setup for PostgreSQL for how to create your database using the correct character encoding, then follow the steps inMigrating to Another Database.

Oracle

If you use Oracle, the best option is to recreate your database.

SeeDatabase Setup for Oraclefor how to create your database using the correct character encoding, then follow the steps inMigrating to Another Database.

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:

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.

Surviving Database Connection Closures

When a database server reboots or a network failure has occurred, all connections in the database connection pool are broken. To overcome this issue, Confluence would normally need to be restarted. To avoid this situation Confluence uses a validation query to check a database connection is alive before attempting to use it. If a broken connection is detected in the pool, a new one is created to replace it.

This validation query is**enabled by default on new installations**(from Confluence 6.5 and later), but if you've upgraded from an older Confluence version you can enable this manually by following the steps below.

While there are several different ways to perform this validation query, we recommend letting the database driver choose how to validate if a connection is still alive, rather than overriding the driver configuration with a specific validation query.

Enable validation query with a direct JDBC connection

To ensure Confluence validates database connections in the database connection pool:

- 1. Stop Confluence.
- 2. Edit the <home-directory>confluence.cfg.xmlfile.
- 3. Insert the following propertyin the cproperties block.

- 4. Saveconfluence.cfg.xml
- 5. Restart Confluence.

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.

Enable validation query with a datasource connection

To ensure Confluence validates database connections in the database connection pool:

- 1. Stop Confluence.
- 2. Edit the<installation-directory>/conf/server.xmlfile (or wherever you have configured your datasource).
- 3. Find the Resource element for your data source, and add the "testOnBorrow" parameter as in the example for PostgreSQL below. Remember to give it the appropriate value for your database type.

- 4. Saveconf/server.xml
- 5. Restart Confluence.

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.

Configuring a datasource connection

This guide covers how to configure a JNDI datasource connection to your database. With this type of connection, Confluence asks the application server (Tomcat) for your database connection information.

If you'd prefer to use a JDBC connection see the guide for your database:

- Database Setup for PostgreSQL
- Database Setup for SQL Server
- Database Setup For MySQL
- Database Setup for Oracle

Direct JDBC is the most common way to connect Confluence to your database and is the easiest method when it comes time to upgrade Confluence.

On this page:

- New Confluence installation
- Existing Confluence installation
- Upgrading Confluence with a datasource
- Known issues

Related pages:

 Database JDBC Drivers

New Confluence installation

The Confluence setup wizard will only provide an option to use a datasource if it detects one in your Tomcat configuration. If you want to use a datasource, follow the steps below.

1. Stop Confluence

In the Confluence setup wizard, you'll be prompted to choose your database. At this point, you should:

- 1. Stop Confluence.
- 2. Back up the following files, in case you need to revert your changes:
 - <installation-directory>/conf/server.xml
 - <installation-directory>/confluence/WEB-INF/web.xml
 - <home-directory>/confluence.cfg.xml

2. Add your database driver

Copy your database driver into the <installation-directory>/libdirectory.

Here's where to find the driver for your database:

- **PostgreSQL**: bundled with Confluence at<installation-directory>/confluence/WEB-INF/lib/postgresql-x.x.x.jar
- **Microsoft SQL Server**: bundled with Confluence at<installation-directory>/confluence /WEB-INF/lib/mssql-jdbc-x.x.x.x.jar
- MySQL: head to Database JDBC Drivers todownload the driver
- Oracle:head toDatabase JDBC Drivers todownload the driver

3. Configure the datasource in Tomcat

Next, add the datasource configuration to Tomcat.

- Edit <installation-directory>/conf/server.xml
- 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 following DataSourceResourceelement for your specific database directly after the lines above (inside theContextelement, directly after the opening<Context.../>line, beforeManager).

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
    driverClassName="org.postgresql.Driver"
    url="jdbc:postgresql://<host>:5432/<database-name>"
    maxTotal="60"
    maxIdle="20"
    testOnBorrow="true"/>
```

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
    driverClassName="com.microsoft.sqlserver.jdbc.SQLServerDriver"
    url="jdbc:sqlserver://<host>:1433;database=<database-name>"
    maxTotal="60"
    maxIdle="20"
    testOnBorrow="true"/>
```

If you're using Confluence 6.3 or earlier, you'll need to specify the jTDS driver for SQL Server. See Configuring a SQL Server Datasource in Apache Tomcat in our 6.3 documentation for a sample configuration.

If you are using the 5.1.x driver (for MySQL 5.7):

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
    driverClassName="com.mysql.jdbc.Driver"
    url="jdbc:mysql://<host>:3306/<database-name>?useUnicode=true&amp;characterEncoding=utf8"
    maxTotal="60"
    maxIdle="20"
    defaultTransactionIsolation="READ_COMMITTED"
    testOnBorrow="true"/>
```

If you're using the 8.0.x driver (for MySQL 8):

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
    driverClassName="com.mysql.cj.jdbc.Driver"
    url="jdbc:mysql://<host>:3306/<database-name>?useUnicode=true&amp;characterEncoding=utf8"
    maxTotal="60"
    maxIdle="20"
    defaultTransactionIsolation="READ_COMMITTED"
    testOnBorrow="true"/>
```

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    driverClassName="oracle.jdbc.OracleDriver"
    url="jdbc:oracle:thin:@<host>:1521:<SID>"
    username="<database-user>"
    password="<password>"
    connectionProperties="SetBigStringTryClob=true"
        accessToUnderlyingConnectionAllowed="true"
    maxTotal="60"
    maxIdle="20"
    maxWaitMillis="10000"
    testOnBorrow="true"/>
```

Seehow to find your Oracle URL.

Replace <database-user>, <password>, <host>and <database-name>(or <SID>for Oracle) with details of your own database. You may also need to change the port, if your database server is not running on the default port.

 Configure the connection pool and other properties. See the Apache Tomcat 9 Datasource documentation for more information.

Here are the configuration properties for Tomcat's standard data source resource factory (org. apache.tomcat.dbcp.dbcp.BasicDataSourceFactory):

- driverClassName Fully qualified Java class name of the JDBC driver to be used.
- maxTotal The maximum number of active instances that can be allocated from this pool at the same time.
- maxIdle The maximum number of connections that can sit idle in this pool at the same time.
- maxWaitMillis The maximum number of milliseconds that the pool will wait (when there are no available connections) for a connection to be returned before throwing an exception.
- password Database password to be passed to the JDBC driver.
- url Connection URL to be passed to the JDBC driver. (For backwards compatibility, the property driverName is also recognized.)
- user Database username to be passed to the JDBC driver.
- validationQuery We don't recommend you set a validation query explicitly. Instead, we recommend you set testOnBorrow, which will use the validation query defined by your database driver. See Surviving Database Connection Closures for more information.
- 5. If you plan to use collaborative editing, you'll need to make sure:
 - You're using a supported database driver. Collaborative editing will fail if you're using an
 unsupported or custom JDBC driver ordriverClassNamein your datasource. SeeDatabase
 JDBC Driversfor the list of drivers we support.
 - Your database connection pool allows enough connections to support both Confluence and Synchrony (which defaults to a maximum pool size of 15)
 - You're using simple username and password authentication for your database.

4. Configure the Confluence web application

Configure Confluence to use this datasource:

- 1. Edit<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>
```

5. Restart Confluence and continue setup process

Now that your datasource is configured, you can continue with the setup wizard.

- 1. Start Confluence.
- 2. Go tohttp://localhost:8090 to return to the setup wizard.
- 3. When prompted choose My own database(datasource).
- 4. Enterthe JNDI name of your datasource, for example, java:comp/env/jdbc/confluence
- *5.* Follow the prompts to finish setting up Confluence.

Existing Confluence installation

If you want to switch from using a direct JDBC connection to a datasource:

- Stop Confluence.
- Back up the following files, in case you need to revert your changes:
 - <installation-directory>/conf/server.xml
 - $^{\bullet} \ \mbox{<installation-directory>/confluence/WEB-INF/web.xml}$
 - ullet <home-directory>/confluence.cfg.xml
- Follow the steps above for a new installation and copy your driver and add the datasource to the appropriate files. You can find the details of your current database connection in <home-directory> /confluence.cfg.xml.
- Edit the < home-directory > / confluence.cfg.xml file and remove any line that contains a property that begins withhibernate.
- Insert the following at the start of the cpropertiessection.

Start Confluence.

Upgrading Confluence with a datasource

If you're upgrading Confluence (manually or using the installer) you will need to:

- Stop Confluence (if you have attempted to start it).
- Copy your database driver into the <installation-directory>/libdirectory.
- Edit <installation-directory>/conf/server.xmland add your datasource resource.
- Edit <installation-directory>/confluence/WEB-INF/web.xmlto configure Confluence to use this datasource.

If you forget to do these steps, Confluence will not start up after upgrade and you'll see the following error:

```
HTTP Status 500 - Confluence is vacant, a call to tenanted [public abstract org.hibernate.Session org. hibernate.SessionFactory.getCurrentSession() throws org.hibernate.HibernateException] is not allowed.
```

Known issues

• If you experience a lot ofContentUpdatedEvent errors in the logs, you may need to add adddefau ltAutoCommit="false" to the datasource in the server.xml file. See

- CONFSERVER-59524 SHORT TERM BACKLOG for more information and full details of the workaround.
- There's a known issue whereSynchrony does not start if Confluence connects to the database using a datasource. See
 - CONFSERVER-60120 Synchrony not starting with datasource after upgrade to Confluence 7.5.2, 7.6.0, 7.6.1 & 7.6.2 CLOSED

for more information and a workaround.

Configuring Confluence Data Center to work with Amazon Aurora

On this page:

- Deploying Confluence Data Center with Amazon Aurora
- Connecting an existing Quick Start deployment to Amazon Aurora
- Manually setting up an Amazon Aurora Database
 - AWS documentation
- Connecting Confluence Data Center to a new Amazon Aurora database
 - Step 1: Shut down Confluence Data Center
 - Step 2: Update the database URL each Confluence node uses
 - Step 3: Configure collaborative editing
 - Step 4: Restart Confluence

Confluence Data Center supports the use of a single-writer, PostgreSQL-compatible Amazon Aurora clustered database. A typical production-grade cluster includes one or more readers in a different availability zone. If the writer fails, Amazon Aurora will automatically promote one of the readers to take its place. For more information, see Amazon Aurora Features: PostgreSQL-Compatible Edition.

Deploying Confluence Data Center with Amazon Aurora

To create a new Confluence Data Center deployment with Amazon Aurora, we recommend that you use the AW S Quick Start for Confluence . This Quick Start lets you configure a PostgreSQL-compatible Amazon Aurora cluster with one writer and two readers in separate availability zones. See Running Confluence Data Center in AWS for more information.

Connecting an existing Quick Start deployment to Amazon Aurora

If you deployed Confluence Data Center using the Quick Start before 11 June 2019, you wont be able to connect it to a new Amazon Aurora cluster. The Quick Start version prior to that date applied some settings that are incompatible with Aurora.

Instead, youll have to migrate your existing data to a new Confluence Data Center deployment:

- 1. Use the latest AWS Quick Start for Confluence to create a new Confluence Data Center deployment.
- 2. Shut down Confluence on the application nodes of both old and new deployments. If you use a standalone Synchrony cluster, shut down all the nodes in that cluster too.
- 3. Migrate your data from the old deployment to the new one:
 - EFS: EFS-to-EFS Backup explains how you can use an easy-to-deploy backup solution to perform a backup of your old EFS and restore it in the new deployment.
 - Database: Migrating Data to Amazon Aurora PostgreSQL contains instructions for migrating from Amazon RDS to a PostgreSQL-compatible Amazon Aurora cluster.

Once you finish the migration, re-start Confluence on all application nodes in the new deployment. If you use a standalone Synchrony cluster, re-start all its nodes.

We strongly recommend your content index after performing a migration, to ensure Confluence search works as expected.

Manually setting up an Amazon Aurora Database

Confluence Data Center specifically supports the use of an Amazon Aurora cluster with the following configuration:

- It must have only one writer, replicating to one or more readers.
- Your PostgreSQL engine must be version 9.6 or higher.

CheckSupported Platformsfor more details.

AWS documentation

AWS has some helpful guides for setting up an Aurora database and migrating to it:

- Modular Architecture for Amazon Aurora PostgreSQL: a Quick Start that guides you through the deployment of a PostgreSQL-compatible Aurora Database cluster. This cluster is similar to the one set up by the AWS Quick Start for Jira (one writer and two readers, preferably in different availability zones).
- Upgrading the PostgreSQL DB Engine for Amazon RDS: shows you how upgrade your database engine to a supported version before migrating it to Amazon Aurora.
- Migrating Data to Amazon Aurora PostgreSQL: contains instructions for migrating from Amazon RDS to a PostgreSQL-compatibleAmazon Aurora cluster.
- Best Practices with Amazon Aurora PostgreSQL: contains additional information about best practices and options for migrating data to a PostgreSQL-compatible Amazon Aurora cluster.

Amazon also offers anAWS Database Migration Serviceto facilitate a managed migration. This service offers minimal downtime, and supports migrations to Aurora from a wide variety of source databases.

If you deployed Confluence Data Center through our AWS Quick Start before 11 June 2019, you cant connect it to a new Amazon Aurora cluster. Rather, youll need to re-deploy Confluence Data Center using our updated Quick Start and migrate your data across. See Connecting an existing Quick Start deployment to Amazon Aurora for more information.

Connecting Confluence Data Center to a new Amazon Aurora database

After deploying an Aurora cluster and migrating your database to it, youll need to properly connect it to Confluence. This will involve updating the database URL used by Confluence Data Center.

Confluence Data Center should point to the the Aurora cluster writer endpoint URL, and include thetargetServ erTypeparameter. This parameter allows Confluence to target the writer database instance, which ensures the application can reconnect to it after a failover.

Your database URL will look something like this:

jdbc:postgresql://<CLUSTER_WRITER_ENDPOINT>:<CLUSTER_WRITER_PORT>/<DATABASE_NAME>?targetServerType=master



If you deployed your Aurora cluster through the Modular Architecture for Amazon Aurora PostgreSQL Quick Start, you can find then the cluster writer details from the Outputs tab in AWS. The RDSEndpoin tAddress and RDSEndpointAddressPort values will be your CLUSTER_WRITER_ENDPOINT and CL USTER_WRITER_PORT, respectively.

The following steps will walk you through the process of connecting Confluence and Aurora.

Step 1: Shut down Confluence Data Center

To safely reconfigure Confluence Data Centers database connection, we recommend a full outage. To do this, stop Confluence on all application nodes.

If you have a standalone Synchrony cluster, stop Synchrony on each node there.

Step 2: Update the database URL each Confluence node uses

How you perform this step depends on how Confluence currently connects to your database.

If you use a direct JDBC connection

- 1. On the first node, edit the <local home > / confluence cfg.xmlfile.
- 2. Update thehibernate.connection.urlproperty with your cluster writer endpoint URL as follows:

```
<CLUSTER WRITER PORT>/<DATABASE NAME>?targetServerType=master
```

- 3. Repeat this change on all other nodes.
- 4. Start Confluence, one node at a time.



This change must be made in the local homedirectory on each node, not in the copy of the confluenc e-cfg.xml that can be found in the shared home.

If you use a datasource connection

- 1. Stop Confluence on all nodes.
- 2. On the first node, edit the <install-directory>/conf/server.xml file.
- 3. Update the url parameter in the datasource Resource element with your cluster writer endpoint URL as follows:

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"</pre>
username="<database-user>"
password="<password>"
driverClassName="org.postgresgl.Driver"
url="jdbc:postgresql://<CLUSTER_WRITER_ENDPOINT>:<CLUSTER_WRITER_PORT>/<DATABASE_NAME>?
targetServerType=master"
maxTotal="60"
maxIdle="20"
validationQuery="select 1"/>
```

- 4. Repeat this change on all other nodes.
- 5. Start Confluence, one node at a time.

Step 3: Configure collaborative editing

Synchrony, the engine that powers collaborative editing, can be deployed in two different ways, which affects how you pass the database URL:

- 1. Managed by Confluence Confluence will automatically launch a Synchrony process on the same node, and manage it for you.
- 2. Standalone Synchrony cluster You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need. This is the default method when you deploy Confluence in AWS using our Quick Start.

If Synchrony is managed by Confluence, you don't need to do anything. Confluence will pass the URL to Synchrony for you.

If you run a Standalone Synchrony cluster, you will need to provide the cluster writer endpoint URL in your startup script. This script will be either < synchrony - home > / start - synchrony . shorstart - synchrony . bat, depending on your operating system. Edit your script as follows:

start-synchrony.sh (Linux)

DATABASE_URL="jdbc:postgresq1://<CLUSTER_WRITER_ENDPOINT>:<CLUSTER_WRITER_PORT>/<DATABASE_NAME>? targetServerType=master"

start-synchrony.bat (Windows)

set DATABASE_URL=jdbc:postgresq1://<CLUSTER_WRITER_ENDPOINT>:<CLUSTER_WRITER_PORT>/<DATABASE_NAME>? targetServerType=master

SeeSet up a Synchrony cluster for Confluence Data Centerfor more information about setting up Synchrony standalone cluster.



If you run Synchrony as a Linux service, you'll need to reinstall the service.

Step 4: Restart Confluence

After making the necessary database URL updates, you can now restart Confluence on each application node, one node at a time.

If you have a standalone Synchrony cluster, restart it on each of the clusters nodes.

Site Backup and Restore

When setting up your Confluence site, it's important to consider how you will back up your data, and restore it, if things go wrong.

Recommended backup strategy

Having a robust backup strategy for your Confluence site is essential. You should back up your database, install directory, and home directories (including attachments) on a regular basis, using the database administration and/or backup tool of your choice.

SeeProduction Backup Strategy for more information.

Scheduled XML backup

Confluence provides a scheduled XML backup option, which backs up your site by performing a full site XML export each day. This method can be useful for small sites, test sites, or in addition to your database and directory backups. We don't recommend you rely solely on this backup method for your production site.

There are a number of reasons why XML site backups are unsuitable for large Confluence sites:

- As the number of pages in your site increases, the XML backup takes progressively longer to complete, and in extreme cases the process of generating the export can cause an outage.
- XML backups can consume a lot of disk space rapidly. For example a 1GB Confluence site will create 30GB worth of backups in a month, if unattended.
- If the XML export file is very large, restoring your site can take a long time, or may time out.
- Marketplace and other user-installed apps are not included in the XML backup. After importing your backup into a new Confluence site, you will need to re-install all user installed apps.

SeeConfiguring Backupsto find out more about the scheduled XML backup, including how to disable this backup, or change how often this job runs.

If you have a Confluence Data Center license, the scheduled XML backup is disabled by default.

Manual XML backup

You can manually create an XML site backup at any time.

Restoring your site from a backup

In the event do you need to restore your site from a backup, the way you do this depends on your backup method.

- SeeRestoring Data from other Backupsfor tips on how to restore Confluence from a database backup.
- SeeRestoring a Site to find out how to import data from an XML site export file into an existing Confluence site.

Migrate to Confluence Cloud

If you're migrating from Confluence Server to Confluence Cloud, you can use the Confluence Cloud Migration Assistantto migrate your content and spaces.

Production Backup Strategy

Although Confluence provides a scheduled XML backup, this backup method is only suitablefor small sites, test sites, or in addition to database and directory backups.

Establishing a production system backup solution

We recommend establishing a robust database backup strategy:

On this page:

- Establishing a production system backup solution
- Which files need to be backed up?
- How do I back up?
- How do I restore?
- Other processes

Related pages:

- Site Backup and Restore
- Create a backup or dump of your database using tools provided by your database.
 If your database doesn't support online backups, you will need to stop Confluence while you do this.
- Create a file system backup of your home directory (both local home and shared home for Data Center)

Once this is in place, you can disable the daily backupscheduled job.

Having a backup of your database and home directories is more reliable, and easier to restore than a large XML backup.

Which files need to be backed up?

Backing up the whole home directory is the safest option, however most files and directories are populated on startup and can be ignored. At minimum, these files/directories *must* be backed up:

- <conf-home>/confluence.cfg.xml
- <conf-home>/attachments(you can exclude extracted text files if space is an issue)

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

- <conf-home>/configif you have modified your ehcache.xml file.
- <conf-home>/index if your site is large or reindexing takes a long time this will avoid the need for a
 full reindex when restoring.

The location of the home directory is configured on installation and is specified in the confluence.init. properties file. For installation created with the automatic installer the default locations are:

- Windows C:\Program Files\Atlassian\Application Data\Confluence
- Linux /var/atlassian/application-data/confluence

For Clustered instances only:Backing up the whole shared home directory is the safest option, however some files and directories are populated at runtime and can be ignored:

- <conf-home>/thumbnails
- <conf-home>/viewfile.

How do I back up?

The commands to back up your database will vary depending on your database vendor, for example the command for PostgreSQL is pg_dump_dbname > outfile.

You should refer to the documentation for your particular database to find out more.

How do I restore?

Our guide on Migrating Confluence Between Servershas instructions on restoring a backup using this technique.

Other processes

XML site backups can be used for other processes in Confluence, for example moving servers or switching to a different database. Using the backup strategy described above will work for those processes too.

- Ourmigrate server procedure used to set up a test server can use a SQL dump as well.
- Thedatabase migrationprocedure 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 Atl assian Experts.

Configuring Backups

Confluence can automatically back up your data by performing a full site export at a scheduled time each day.

- If you have a Confluence Server license, the scheduled XML backup happens every day at 2:00am by default.
- If you have a Confluence Data Center license, the scheduled backup is disabled, as it is not suitable for large sites.

The zipped XML backup file will be named'backup-yyyy_MM_dd', and stored in thebackupsdirectory of yourConfluence Home directory.

On this page:

- Configuring automated backups
- Performing manual backups
- Scheduled XML backups in Confluence Data Center

You'll need System Administrator global permissions to make changes to the scheduled XML backups. You can choose to:

- disable the scheduled backup altogether
- change the naming convention
- include or exclude attachments
- schedule the backup at a different time
- store the backup files in a different location (disabled by default, find out how to enable it below).

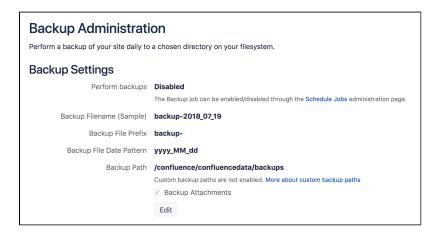


We don't recommend relying on these automatic backups in production sites. You should instead back up your database, installation directory and home directory manually. SeeProduction Backup Strategyfor more information.

Configuring automated backups

To configure scheduled XML backups:

- 1. Go to Seneral Configuration Backup administration.
- 2. Choose Edit to:
 - Change the backup file name prefix.
 - Use a different date format (uses the syntax described insimple date format).
 - Choose whether to include or exclude attachments from backups (attachments are included by default).
 - Choose to store backup files in a custom location (this is disabled by default seeEnabling backup path configuration below).
- 3. Saveyour changes.



Enabling Backup Path Configuration

For security reasons, the ability to change the backup file location **Backup administration** screen is disabled by default.

To enable custom backup paths:

- 1. Stop Confluence.
- 2. Edit the <confluence-home>/confluence.cfg.xmlfile.
- 3. Change the value of the following property to true:

- 4. Restart Confluence to pick up the change.
- 5. Go to General Configuration>Backup administration to enter the new path.

The directory must be on either a local drive or amounted network drive. Make sure the mounted drive is on a physical server and not a Virtual Machine image.

If you migrate Confluence to a new server or change your architecture, you will need to update this path. Changing your home directory location will not automatically update your backup file path if you've enabled a custom path.

Disable automatic backups

If you have an appropriate Production Backup Strategy, you may want to disable automatic backups to save on disk space.

To disable automatic backups entirely:

- 1. Go to General Configuration>Scheduled jobs.
- 2. Choose Disable next to the Back up Confluence job.

Change the backup schedule

To change the frequency of backups, or to change the time the backup runs each day:

- 1. Go to General Configuration>Scheduled jobs.
- 2. ChooseEditnext to theBack up Confluencejob
- 3. Enter the new schedule using acron expression.

The time zone used for the scheduled job is taken from the server on which Confluence is running. Go to General Configuration System Information to look up the System Time.

Performing manual backups

If you need a one-off XML backup, you can manually perform a site export. SeeManually Backing Up the Site for more information.

These files are not saved to the same location as the automated backups, they are saved in the temp directory. You can change where the zipped XML files are saved by changing the location of your <Conflue nce-home>/tempdirectory. SeeConfluence Home and other important directories for more information on how to do this.

Scheduled XML backups in Confluence Data Center

The scheduled XML backup is **disabled by default** in Confluence Data Center, as the job is known to cause outages in sites with a lot of data.

We recommend that your backup strategyincludes full backups of your database, local home and shared home directories on a regular basis.

If you do choose to enable the scheduled XML backup (for example on a staging or test site), the default backup path is<shared-home>/backups. You can find the location of your shared home in theconfluence.cluster.homeproperty.

User Submitted Backup & Restore Scripts

These scripts are user-submitted and should be used with caution as they are not covered by Atlassian technical support. If you have questions on how to use or modify these scripts, please post them to Atlassian Answers.

Delete Old Backups - Wscript Script On Windows

This script examines backup filename and deletes them if necessary, it may need to be edited.

```
'If you want 3 day old files to be deleted then insert 3 next to Date - "your number here"
'This script will search out and delete files with this string in them ".2005-12-04-" This of course
depends on the number you enter.
'You can always do a wscript.echo strYesterday or strFileName to see what the script thinks you are
searching for.
dtmYesterday = Date - 3
strYear = Year(dtmYesterday)
strMonth = Month(dtmYesterday)
If Len(strMonth) = 1 Then
    strMonth = "0" & strMonth
End If
strDav = Dav(dtmYesterdav)
If Len(strDay) = 1 Then
   strDay = "0" & strDay
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
```

Manually Backing Up the Site

Confluence can be configured automatically back up your data by performing a full site export at a scheduled time each day.

You can also manually back up Confluence at any time, by performing a full site export.

You'll needSystem Administrator permissions to do this.

Related pages:

- Restoring a Site
- Configuring Backups
- Production Backup Strategy



Good to know:

- We don't recommend you rely on XML site exports as your main backup method. Instead, you should regularly back up your database, install and home directories. SeeProduction backup strategy for more information.
- Marketplace and user-installed apps are not included in the XML export. After importing your site export file into a new Confluence site, you'll need to re-install all apps that are not bundled with Confluence as theplugindatatable is not backed up in a manual backup.
- You can't import a site export XML file into an earlier version of Confluence.

Create the site export file

To create an XML export of your site:

- 1. Go to Configuration Backup & Restore.
- 2. ChooseAlso save a copy to the backups directory to store a copy of the backup in the same folder as Confluence'sbackups.
 - 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.
- 3. Choose**Include attachments**to include attachments in your backup.
- 4. Choose Export.

The process can take some time.

If you have a large site, you should temporarily disable the Clean Temporary Directory scheduled job, as this job deletes the contents of the /temp directory at 4am each day, and will delete your export files. You should re-enable this job once the export is complete.

If you repeatedly experience timeout errors, try creating the export directly from the server. This will speed up the process and prevent timeouts.

For example, use this URL: http://localhost:8090/confluence/admin/backup.action . directly from your server.

What's included in the export?

The site export includes spaces (including pages, blogs, comments, attachments, and unpublished changes), users and groups. Essentially everything in your site except add-ons.

Retrieving the site export file

Confluence will create the backup as a zipped XML file in your<home-directory>/backups>directory. You'll need access to the Confluence server itself in order to retrieve this file.

Allow export files to be downloaded from within Confluence

By default, you can't retrieve the backup file from within Confluence. This feature is disabled for security reasons, but you can choose to enable it. Once enabled, Confluence will prompt you to download the backup file when the backup process finished. We recommend that you keep this feature offin production environments.

To enable download of the backup file from within Confluence:

- 1. Stop Confluence.
- 2. Edit the < confluence home > \confluence.cfg.xmlfile.
- 3. changeadmin.ui.allow.manual.backup.downloadto true.
- 4. 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'.

Restoring the site export file

There are some restrictions on which Confluence versions you will be able to import this file into. The most important is that you **can't import into an earlier version of Confluence**. See Restoring a Site for more information and troubleshooting tips.

Restoring a Site

This page describes how to restore data from an XML site export file into an existing Confluence site.

If you want to import data into a new site, seerestorin g from backup during setup.

You needSystem Administratorpermissions in order to perform this function.

On this page:

- Before you start
- Troubleshooting
- Note about using site exports as backups

Related pages:

- Production Backup Strategy
- Exporting a site
- Importing a Space



Importing a site export file will:

- Overwrite all existing Confluence content in your database. Back up your database before
- Log you out of Confluence. Make sure you know the login details contained in the file you're about to import.

Before you start

- All content replaced. Importing a site will replace all your content and users. Back up your database before you start.
- Selective space restoration not possible. You can't select a single space to restore from the entire site backup.
- Version compatibility. Confluence accepts site backups from many previous Confluence versions. You can check which versions are accepted in the Backup and Restore screen. You can only import into a later version of Confluence, not an earlier one.
- For best results, export from and import into the same Confluence version.
- XML export files should not be used to upgrade Confluence. Upgrade Confluence by following Up grading Confluence

Check your export is compatible

To check that your site export can be successfully restored:

- 1. Start up the Confluence site you'll be importing into.
- 2. Go to General Configuration>Backup and Restore.
- 3. Check the accepted Confluence version it's listed underImport Confluence data.



You can't import into an earlier version of Confluence.

For example, if your site export was generated from Confluence 6.12, you can't import it into Confluence 6.6.

If your export is from Confluence Cloud you can only import it into Confluence 6.0 or later.

Import a Confluence site

There are two ways to import a site - by uploading a file, or from a directory on your Confluence server. Uploading a file is only suitable for small sites. For best results, we recommend importing from the restore directory.

To upload and import a small site:

- 1. Go to General Configuration>Backup and Restore.
- 2. Under Upload a site or space export file, click Choose File and browse for your space export file.
- 3. UncheckBuild Indexif you want to create the index at a later stage.
- 4. Choose Upload and import.

To import a site from the home directory:

- 1. Copy your export file to < confluence home > / restore. (If you're not sure where this directory is located, the path is listed in the **Backup and Restore**screen)
- 2. Go to So Seneral Configuration>Backup and Restore.
- 3. Select your site export file under Import from home directory
- 4. UncheckBuild Indexif you want to create the index at a later stage.
- 5. ChooseImport.

Building the index is optional during the import process. The content of your site won't be searchable until the index is created, but if you have a very large site, you may choose to rebuild the index manually after the import is complete.



(i) Using Confluence Data Center?

If you're using Confluence Data Center, and you run a Synchrony standalone cluster there are a few extra steps. You need to stop Synchrony completley, and we also recommend performing the import with just one Confluence node running, and directing traffic away from that node.

Once the import is complete, you can restart your Synchrony cluster, and then restart your remaining nodes (one at a time).

This is not required if you allow Confluence to manage Synchrony for you.

Troubleshooting

If you have problems importing a site, check out these hints.

Is your file too large to upload?

This is a very common problem. It happens when the file can't be uploaded to the server in time. To avoid this problem, drop your export file into the <confluence-home > / restoredirectory and import

Are you trying to import into an earlier version of Confluence?

This is not possible. You can only import a site into the same version or a later compatible version.

Is the import timing-out or causing out of memory errors? If the site to be imported is large, you may need to temporarily increase the memory available to Confluence. SeeHow to fix out of memory errors by increasing available memory.

Is your username or password not recognized?

All user data was overwritten during the import process. You need to log in with a system administrator account from the site that was exported. If you don't know the password, you'll need to reset it from the database. SeeRestore Passwords To Recover Admin User Rights.

Is your site export from Confluence Cloud?

You can only import into Confluence 6.0 or later. The Cloud export does not include a system administrator account, so you will need to start Confluence in recovery mode, create a new system administrator account, and make it a member of the confluence-administrators group. See Restore Passwords To Recover Admin User Rights for more.

Did you download the export file on a Mac?

If you get an error saying that Confluence can't find the export Descriptor. properties file, chances areOS X has unzipped the backup for you and sent the original zipped file to the trash. You need to retrieve the original zip file from the trash and then try the import again.

• Importing into a site with a Synchrony standalone cluster?
You must stop your Synchrony cluster before commencing the site import.

Note about using site exports as backups

- Production backup strategy preferred. We recommend that you follow the Production Backup Strategy(which involves backing up your database and home directory) for your production Confluence site, because Confluence XML exports are not recommended as the sole backup mechanism
- Restoring from other backups. If your daily backup zip files can't be restored for some reason, but
 you have backups of both your database and your Confluence home directory, you'll be able torestore
 from these backups.

Restoring a Space

You can export a spaceincluding pages, comments and attachments to a zip that contains an XML file and, optionally, all the attachments in the space. To import the space to another Confluence site, restore the zip as described below.

You need System Administrator permissions in order to restore a space from an XML zip file.

Export and import compatibility

To find out which versions your current Confluence version can accept space exports from, go to $^{\circ} >$ General Configuration > Backup and Restore.

If you need to import a space from Confluence 5.3 or earlier, you'll need to f ollow a workaround.

To find out what is included in an XML export, see Export Content to Word, PDF, HTML and XML.



You can't import into an earlier version of Confluence.

For example, if you export a space from Confluence 5.9, you can't import it into Confluence 5.5.

If your export is from Confluence Cloud, you can only import it into Confluence 6.0 or later.

On this page:

- Export and import compatibility
- Importing a space from Confluence Cloud
- Importing a space from Confluence Server or Data Center
 - Upload a site or space export file
 - Import from the home directory
- Groups and permissions
- Check the size of a space before importing
- Troubleshooting
- Workaround for restoring spaces from Confluence 5.3 and earlier

Related pages:

Restoring a Site

Importing a space from Confluence Cloud

As the way users are managed is different in Confluence Cloud there are a few more considerations when importing a space from Confluence Cloud into Confluence Server or Data Center.

SeeImport a space from Confluence Cloud for a step-by-step guide.

Importing a space from Confluence Server or Data Center

AWe recommend performing a full backup of your database before importing a space. Occasionally the space import process may fail, and a backup will make it easier for you to roll back.

There are two ways to import a space by uploading a file, or from a directory on your Confluence server. Uploading a file is only suitable for small spaces. For best results, we recommend importing from the restore directory.

Upload a site or space export file

To upload and import a small space:

- 1. Go to General Configuration Backup and Restore.
- 2. Under Upload a site or space export file, clickChoose File and browse for your space export file.
- 3. Uncheck**Build Index**if you want to create the index at a later stage.
- 4. Choose Upload and import.

Once the import is complete, you can either navigate directly to the space, or head to the **Space Directory**.

Import from the home directory

Importing from the home directory is a great alternative for large spaces, as you don't need to upload the file via your browser.

To import a space from the home directory:

- Copy your space export file to < confluence home > / restore.
 (If you're not sure where this directory is located, the path is listed in the Backup and Restore screen)
- 2. Go to O Seneral Configuration Backup and Restore.
- 3. Select your space export file under **Import from theHome Directory**.
- 4. Uncheck **Build Index** if you want to create the index at a later stage.
- 5. Choose **Import**.

Building the index is optional during the import process. The content of your imported space won't be searchable until the index is created, but, if you have a very large site, rebuilding the index can take a long time and impact your site'sperformance. Alternatively, you can rebuild the index manually at a low peak time.

Groups and permissions

Importing a space will not import any users or groups that may have been granted specific space permissions in your source Confluence site. This means that if any pages are restricted to these groups, you may not be able to see them until you recreate these groups in your destination site.

Check the size of a space before importing

It's important to check the size of your exported space before you attempt it import it as it may mean you need to use a particular method, or complete the import at a quieter time.

The total file size of the zip will determine whether you can upload the file through the browser, or if you need to use the home directory method.

The size of the entities.xml file provides an indication of the amount of entities such as pages, versions, and comments. A very large time may take quite some time to import.

To check the size of the entities.xml in an exported space:

- 1. Unzip the space export file.
- 2. Navigate to the entities.xml file and check the size of that file. The way you do this will depend on your operating system.

Troubleshooting

If you have problems importing a space, check out these hints.

Is your file too large to upload?

This is a very common problem. It happens when the file can't be uploaded to the server in time. To avoid this problem, drop your export file into the <confluence-home>/restore directory and import it from there.

- Are you trying to import into an earlier version of Confluence?
 - This is not possible. You can only import a space into the same version or a later compatible version.
- Is your space export file from Confluence Cloud?
 - You can only import this file into Confluence 6.0 or later. Trying to import into earlier versions can cause major problems. Seelmport a space from Confluence Cloudfor other considerations.
- Does a space with the same space key already exist?

 Space keys are unique, soif you already have a space with
 - Space keys are unique, soif you already have a space with the same key, you'll need to delete the existing space before importing the new one.
- Is the import timing-out or causing out of memory errors?

 If the space to be imported is very large, you may need to temporarily increase the memory available to Confluence. SeeHow to fix out of memory errors by increasing available memory.

- Did you download the export file on a Mac?
 - If you get an error saying that Confluence can't find the exportDescriptor.properties file, chances are OS X has unzipped the backup for you and sent the original zipped file to the trash. You need to retrieve the original zip file from the trash and then try the import again.
- Did your import fail? Sometimes importing a space may fail because the process times out or runs
 out of memory. This can lead to data being left behind in your database. SeeAfter a failed space
 import, it's not possible to re-import because of leftover space datafor more information.
- Do you have adequate disk space?
 Confluence will need to make copies of the file being imported at various points during the import process. Make sure you have enough disk space to temporarily accommodate multiple copies of the site export file.

Workaround for restoring spaces from Confluence 5.3 and earlier

If you need to import a space from a version prior to Confluence 5.3 into a later version of Confluence, 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 yourtemporary site to the same version as the site where you want to import the space (see Upgrading Confluencefor instructions).
- 5. Export the space from your temporary Confluence site (it'll now have the right version number).
- 6. Import the space into your production Confluence site.

Restoring a Test Instance from Production



See Migrating Confluence Between Servers for a more comprehensive explanation.

Many Confluence administrators will have a production instance running the "live" version of Confluence, as well as a test instance for testing upgrades and so on. In this situation, it's quite common that the two instances are running different versions of Confluence. This document describes how to copy the data from a production instance to a test instance, where the production version may be different to the test version.

Before proceeding with this guide, ensure you have read and understood the normal procedure for upgrading Confluence.

Updating 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';
```

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.

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 an zipped XML backup file, it is still possible to restore provided you have:

- 1. A backup of your home directory.
- 2. A backup of your database (if you're using an external database).

Instructions for this method of restoring differ depending on whether you are using the embedded database or an external database (like Oracle, MS SQL Server, MySQL or Postgres).

Embedded Database

If you are running against the embedded database, the database is located inside the database folder of your Confluence Home Directory. Hence, all you need to do is:

- 1. Retrieve the most recent backup of your home directory.
- 2. Unpack the Confluence distribution and point the confluence-init.properties file to this directory.

External Database

If you're using an external database, you need to do the following.

- 1. Prepare backups of your home directory and database (preferably backups that are dated the same). That is, make sure the home directory is accessible on the filesystem and the database available to be connected to.
- 2. If this database happens to have a different name, or is on a different server, you need to modify the jdbc url in the confluence.cfg.xml file inside the Confluence Home Directory. The value of this property is specified as hibernate.connection.url.
- 3. Unpack the Confluence distribution and point the confluence-init.properties file to the home directory.

Retrieving File Attachments from a Backup

File attachments on pages can be retrieved from a backup without needing to import the backup into Confluence. This is useful for recovering attachments that have been deleted by users.

Both automated and manual backups allow this, as long as the 'Include attachments' property was set. If you want to restore pages, spaces or sites, see the Confluence Administrator's Guide instead.

Before following the instructions for recovering attachments below, we will review how backups store file and page information.

```
⚠ The information on this page does not apply to Confluence
Cloud.
```

How Backups Store File and Page Information

The backup zip file contains entities.xml, an XML file containing the Confluence content, and a directory for storing attachments.

Backup Zip File Structure

Page attachments are stored under the attachments directory by page and attachment id. Here is an example listing:

```
Listing for test-2006033012_00_00.zip \attachments\98\10001 \attachments\98\10002 \attachments\99\10001 entities.xml
```

Inside the attachment directory, each numbered directory inside is one page, and the numbered file inside is one attachment. The directory number is the page id, and the file number is the attachment id. For example, the file \attachments\98\10001 is an attachment with page id 98 and attachment id 10001. You can read entities.xml to link those numbers to the original filename. Entities.xml also links each page id to the page title.

Entities.xml Attachment Object

Inside the entities.xml is an Attachment object written in XML. In this example, the page id is 98, the attachment id is 10001 and the filename is myimportantfile.doc. The rest of the XML can be ignored:

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>
cproperty name="title"><![CDATA[Editing Your Files]]>
...
</object>
```

Instructions for Recovering Attachments

Each file must be individually renamed and re-uploaded back into Confluence by following the instructions below. Choose one of the three methods:

Choice A - Recover Attachments By Filename

Best if you know each filename you need to restore, especially if you want just a few files:

- 1. Unzip the backup directory and open entities.xml.
- 2. Search entities.xml for the filename and find the attachment object with that filename. Locate its page and attachment id.
- 3. Using the page and attachment id from entities.xml, go to the attachments directory and open that directory with that page id. Locate the file with the attachment id.
- 4. Rename the file to the original filename and test it.
- 5. Repeat for each file.
- 6. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

Choice B - Restore Files By Page

Best if you only want to restore attachments for certain pages:

- 1. Unzip the backup directory and open entities.xml.
- 2. Search entities.xml for the page title and find the page object with that title. Locate its page id.
- 3. Go to the attachments directory and open that directory with that page id. Each of the files in the directory is an attachment that must be renamed.
- 4. Search entities.xml for attachment objects with that page id. Every attachment object for the page will have an attachment id and filename.
- 5. Rename the file with that attachment id to the original filename and test it.
- 6. Repeat for each page.
- 7. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

Choice C - Restore All Files

Best if you have a small backup but want to restore many or all the attachments inside:

- Following process is applicable to space export only. Site xml backups do not require page id to be updated manually due to the nature of persistent page_id's.
 - 1. Unzip the backup directory and open entities.xml.
 - 2. Go to the attachments directory and open any directory. The directory name is a page id. Each of the files in the directory is an attachment that must be renamed.
 - 3. Search entities.xml for attachment objects with that page id. When one is found, locate the attachment id and filename.
 - 4. Rename the file with that attachment id to the original filename and test it.
 - 5. Find the next attachment id and rename it. Repeat for each file in the directory.
 - 6. Once all files in the current directory are renamed to their original filenames, search entities.xml for the page id, eg directory name. Find the page object with that page id and locate its page title.
 - 7. Rename the directory to the page title and move on to the next directory. Repeat for each un-renamed directory in the attachments directory.
 - 8. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

Troubleshooting failed XML site backups



XML site backups are only necessary for migrating to a new database. Setting up a test server or Establishing a reliable backup strategy is better done with an SQL dump.

Related pages:

Enabling detailed SQL logging

Seeing an error when creating or importing a backup?

Problem	Solution
Exception while creating backup	Follow instructions below
Exception while importing backup	Follow Troubleshooting XML backups that fail on restore instead

Common problems

• Is the export timing out or causing out of memory errors? If your site is large, you may need to temporarily increase the memory available to Confluence. SeeHo w to fix out of memory errors by increasing available memory.

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

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.propertiesand add this to the bottom and save:

```
log4j.logger.com.atlassian.hibernate.extras.XMLDatabinder=DEBUG, confluencelog
log4j.additivity.com.atlassian.hibernate.extras.XMLDatabinder=false
```

- 5. Find your application logs. Move or delete all existing Confluence application logs to make it easier to find the relevant logging output. You could also choose to mark the application logs after restarting Confluence, to indicate when you started the export.
- 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
         at net.sf.hibernate.ObjectNotFoundException.throwIfNull(ObjectNotFoundException.java:24)
10
         at net.sf.hibernate.impl.SessionImpl.immediateLoad(SessionImpl.java:1946)
11
         at net.sf.hibernate.proxy.LazyInitializer.initialize(LazyInitializer.java:53)
12
         at net.sf.hibernate.proxy.LazyInitializer.initializeWrapExceptions(LazyInitializer.java:
60)
13
         at net.sf.hibernate.proxy.LazyInitializer.getImplementation(LazyInitializer.java:164)
14
          at net.sf.hibernate.proxy.CGLIBLazyInitializer.intercept(CGLIBLazyInitializer.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.toGenericXML(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 atlassian-confluence.log, 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 2 535 (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.
BucketPropertySetItem@a70067d3]; SQL []; Violation of PRIMARY KEY constraint
'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate key in object 'OS_PROPERTYENTRY'.; nested
exception is java.sql.SQLException: Violation of PRIMARY KEY constraint 'PKOS_PROPERTYENTRY_314D4EA8'.
Cannot insert duplicate key in object 'OS_PROPERTYENTRY'.
```

this indicates that the Primary Key constraint 'PK_OS_PROPERTYENTRY_314D4EA8' has duplicate entries in table 'OS_PROPERTYENTRY'.

You can locate the constraint key referring to 'PK_OS_PROPERTYENTRY_314D4EA8' in your table 'OS_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS_PROPERTYENTRY' table is:

SELECT ENTITY_NAME, ENTITY_ID, ENTITY_KEY, COUNT(*) FROM OS_PROPERTYENTRY GROUP BY ENTITY_NAME, ENTITY_ID, ENTITY_KEY HAVING COUNT(*)>1

To Help Prevent This Issue From Reoccurring

- 1. If you are using the embedded database, be aware that it is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss, which is why an external database is recommended for production use. You should migrate to an external database.
- 2. If you are using an older version of Confluence than the latest, you should consider upgrading at this point.

Troubleshooting XML backups that fail on restore

XML site backups are only necessary for migrating to a new database. Upgrading Confluence, Setting up a test server or Prod uction Backup Strategy is better done with an SQL dump.

Seeing an error when creating or importing a site or space backup?

Problem	Solution
Exception while creating backup	See Troubleshooting failed XML site backups
Exception while importing backup	See instructions below

On this page:

- Common problems
- 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

Common problems

To upload and import a small site:

- 1. Go to General Configuration>Backup and Restore.
- 2. Under Upload a site or space export file, click Choose File and browse for your space export file.
- 3. UncheckBuild Indexif you want to create the index at a later stage.
- 4. Choose Upload and import.

To import a site from the home directory:

- 1. Copy your export file to < confluence home > / restore. (If you're not sure where this directory is located, the path is listed in the Backup and Restorescreen)
- 2. Go to General Configuration>Backup and Restore.
- 3. Select your site export file under Import from home directory
- 4. UncheckBuild Indexif you want to create the index at a later stage.
- 5. Choose**Import**.

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 application log files and thecatalina. <datestamp>.log (in your installation directory) to find out what object could not be converted into XML format.
- 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: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 the log filescheck 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 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 cpl.ID, cpl.CP_TYPE, cpl.USERNAME, cpl.GROUPNAME, cpl.CPS_ID,
cpl.CREATOR, cpl.CREATIONDATE, cpl.LASTMODIFIER, cpl.LASTMODDATE
FROM CONTENT_PERM cpl, CONTENT_PERM_SET cpsl, 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 cpl.CPS_ID, cpl.CP_TYPE, cpl.USERNAME, cpl.CREATIONDATE;
SELECT DISTINCT cpl.ID, cpl.CP_TYPE, cpl.USERNAME, cpl.GROUPNAME, cpl.CPS_ID,
cpl.CREATOR, cpl.CREATIONDATE, cpl.LASTMODIFIER, cpl.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.
BucketPropertySetItem@a70067d3]; SQL []; Violation of PRIMARY KEY constraint
'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate key in object 'OS_PROPERTYENTRY'.; nested
exception is java.sql.SQLException: Violation of PRIMARY KEY constraint 'PKOS_PROPERTYENTRY_314D4EA8'.
Cannot insert duplicate key in object 'OS_PROPERTYENTRY'.
```

This indicates that the Primary Key constraint 'PK_OS_PROPERTYENTRY_314D4EA8' has duplicate entries in table 'OS_PROPERTYENTRY'.

You can locate the constraint key referring to 'PK_OS_PROPERTYENTRY_314D4EA8' in your table 'OS_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS_PROPERTYENTRY' table is:

```
SELECT ENTITY_NAME, ENTITY_ID, ENTITY_KEY, COUNT(*) FROM OS_PROPERTYENTRY GROUP BY ENTITY_NAME, ENTITY_ID, ENTITY_KEY HAVING COUNT(*)>1
```

Troubleshooting "net.sf.hibernate.PropertyValueException: not-null" related problems

If you're receiving a message like:

ERROR [Importing data task] [confluence.importexport.impl.ReverseDatabinder] endElement net.sf.hibernate. PropertyValueException: not-null property references a null or transient value: com.atlassian.user.impl. hibernate.DefaultHibernateUser.name

This means there's an unexpected null value in a table. In the above example, the error is in the name column in the USERS table. We've also seen them in the ATTACHMENTS table.

Remove the row with the null value, redo the xml export, and reimport.

To help prevent this issue from recurring

- 1. If you are using the embedded database, be aware that it is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss, which is why an external database is recommended for production use. You should migrate to an external database.
- 2. If you are using an older version of Confluence than the latest, you should consider upgrading at this



⚠ 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

Import a space from Confluence Cloud



As of 16 July 2019, usernames are no longer included in space exports from Confluence Cloud.

Email addresses will be included, regardless of profile visibility settings, if the person performing the export is a Site Admin.

Email matching is available in selected Confluence versions. See below for more information.

The user base of your Confluence Cloud and Confluence Server or Data Center sites are seperate. Although the same people may have accounts in both sites, the way information is stored about them is different. For example, in Confluence Cloud usernames have been replaced by email addresses, and they have an additional ID (a random string of characters) that acts as a unique identifier.

When you import a space into Confluence we attempt to attribute content based on username. If the two usernames match, we will attribute content to the user.

In spaces exported from Cloud, where there is no username, we willattempt to match users by their email address. To reduce the risk of making restricted pages visible to the wrong person, imported content will be attributed to 'unknown user' if:

- the email address is used by multiple user accounts (with different usernames), or
- the user account doesn't have an email address (for example if it is marked private, and the space was exported by someone who is not a Site Admin, the email address would not be included in the export).



Email matching is only available in the following Confluence Server and Data Center versions:

- 6.6.14 and any later 6.6 Long Term Support release version
- 6.13.5 and any later 6.13 Long Term Support release version
- 6.15.4 and later

In all other versions content will be attributed to 'unknown user' if we're unable to match by username.

Import a space from Confluence Cloud

To import a small space from Confluence Cloud:

- 1. In Confluence Cloud, export the space to XML.
- 2. In Confluence Server, go to Seneral Configuration Backup & Restore.
- 3. Scroll down to Restore Confluence Data and select your export file.
- 4. ChooseImport.

To import a large space, the steps are the same, however we recommend dropping the export file into your home directory, rather than uploading it via your browser. SeeRestoring a Spacefor more details.

You may need to restore some permissions to the space, if any users or groups aren't present if your destination site.

About unknown users

Any Cloud user accounts found in the space export, that are not reconciled with an existing Server / Data Center user, will appear in the Unsynced from directorylist. They may be listed by email address, or by ID (depending on whether the Cloud user has chosen to keep their email address private).

Permissions and restrictions are respected, so if a space or page is restricted to just one of these users, it will not be visible to other people. An administrator will need to restore permissions after the import is complete.

Here's an example of how unknown users may appear in a page.

Restore permissions and restrictions

If the content you import is not attributed to existing users, there will be some work to do to restore the correct permissions to the right people. People may not be able to see the space until you do this.

Restore space admin permissions

First step is to make sure the space has at least one space administrator. To do this:

- 1. Go to General Configuration>Space Permissions.
- 2. ChooseRecover Permissions beside the newly imported space.
- 3. Choose Manage Permissions. This will take you to the Permissions page in that space.
- 4. Grant a user or group Space Admin permission for the space and save your changes.

If you're a member of the confluence-administrators super group, you can skip steps 2 and 3, and navigate directly to the space.

Restore space permissions

Now that the space has at least one space admin, they can restore any other permissions.

- 1. Go to the space and choose **Space tools>Permissions** from the bottom of the sidebar
- 2. Grant each user the desired permissions. It can be useful to have the space permissions screen in Confluence Cloud open while you do this.

As long as any groups are named the same in Cloud and Server / Data Center you shouldn't need to make any changes to groups. If your groups aren't named the same, you can add any relevant groups at this point.

Restore page restrictions

Pages with view restrictions applied in Confluence Cloud may be associated with unknown users in Server / Data Center. This means the pages won't be visible.

Space admins can remove individual page restrictions

- 1. Go to the space and choose Space tools > Permissions from the bottom of the sidebar
- 2. Go to the Restricted Pages tab. Any pages with view or edit restrictions will be listed.
- 3. Click the padlock icon beside the page to remove one of the View restrictions.
- 4. If the page is still restricted, use your browser back button and click the padlock beside another View restriction. Repeat this process until enough restrictions have been removed that you can see the page (you'll land on Page Information).
- 5. Choose ••• > Restrictions.
- 6. You can now reinstate the view and edit restrictions. It can be useful to have the Cloud page open to refer to.

Removing restrictions so that you can see the page may mean that the page becomes temporarily visible to others. If this is a concern you can either apply a temporary view restriction to a parent page, or perhaps remove space permissions until you've finished restoring the right view restrictions.

Understanding the risks

When importing a space from Confluence Cloud, there's a small risk that content is attributed to the wrong user, which would make any restricted pages visible to the wrong person. This is because the only information we can use to match the user is their email address, which can be changed by the user themselves or by an administrator.

It's essential that email addresses are associated with the correct user accounts. Content may be attributed to the wrong user account if the email address has been changed maliciously, or accidentally, for example if a username and email combination has been reused, so that a former and current employee share the same username and email address.

We mitigate this risk by only associating content to user accounts that have a unique email address. We don't match accounts with no email address, or where the same email address has been used for multiple user accounts with different usernames, even if they exist in different user directories.

However, if the space you are importing is sensitive, you may want to manually check whether there have been any changes to email addresses recently, before importing a space from Confluence Cloud.

Attachment Storage Configuration

By default Confluence stores attachments in the home directory (e.g. in a file system).

If you have upgraded from an earlier Confluence version you may still be storing attachments in your database or WebDAV, and should consider migrating to a supported storage method.

On this page:

- Attachment storage methods
 - Local file system
 - Database (deprecated)
 - WebDav (deprecated)
- Migrating to a supported attachment storage option

Related pages:

Working with Confluence Logs

Attachment storage methods

Earlier Confluence versions allowed attachments to be stored in WebDav or the Confluence database. This is not an option for new installations.

Local file system

By default, Confluence stores attachments in the attachments directory within the configured Confluence home folder.

Database (deprecated)

Confluence 5.4 and earlier gives administrators the option to store attachments in the database that Confluence is configured to use.

We recommend migrating to storing attachments in the file system.



Mhile storing attachments in the database can offer some advantages (such as ease of backup, and avoiding issues with some characters in attachment filenames), please be aware that the amount of space used by the database will increase because of greater storage requirements.

WebDav (deprecated)

WebDav is no longer available as an attachment storage option.

This has no impact on your ability toconfiguring a WebDAV clientto access spaces, pages or attachments in your Confluence site.

Migrating to a supported attachment storage option

If you are storing attachments in WebDay or your database, you can migrate to storing attachments in the file system. When migrating attachments from your database to a filesystem, the attachments are removed from the database after migration.



(1) 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.

To improve logging during the migration, add the package com.atlassian.confluence.pages. persistence.daowith levelDEBUG. See Configuring Loggingfor more information.

To migrate, follow the steps below:

- 1. Go to General Configuration>Attachment storage.
- 2. Click**Edit**to modify the configuration.
- 3. SelectLocally in Confluence home directory.
- 4. ClickSaveto save the changes.
- 5. 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?



If you're already storing attachments in a file system, the **Attachment Storage** option won't appear in the admin console - this is because you're already using the only supported storage method, and don't need to migrate.

Configuring Attachment Size

You can limit the size of files that can be uploaded and attached in Confluence.

To configure the maximum file size that can be uploaded:

- 1. Go to General Configuration.
- 2. Choose Edit.
- 3. Enter the maximum size next to **Attachment Maximum Size**.
- The default is 100 MB.
- 4. Choose Save.

Related pages:

- Recognized System Properties
- Files

How attachments are indexed

When a file is uploaded, Confluence will attempt to extract and index its text. This allows people to search for the content of a file, not just the filename. This process is quite memory intensive and can cause out of memory errors when very large files are uploaded. Confluence has a number of safeguards to prevent this happening:

- If the uploaded file is larger than 100 MB, Confluence will not attempt to extract text or index the file contents. Only the filename will be searchable.
- If the uploaded file is one of the following types, Confluence will only extract up to:
 - 1 MB of text from Excel (.xlsx) and PowerPoint (.pptx)
 - 8 MB of text from PDF (.pdf)
 - 10 MB of text from other text files (including .txt, .xml, .html, .rtf etc)
 - 16 MB of text from Word (.docx)

Note that this is based on the size of the file when it's uncompressed. As .xlsx and .docx files are compressed, text extraction may fail even though the size of the file appears to be under the limit.

• If the text extracted from the file was greater than 1 MB, it will be searchable, but Confluence will not show this text as an excerpt with the search result.

If Confluence stops extracting text, only a portion of the file's content will be searchable.

Confluence will only attempt to extract and index the file once. If it fails, it will not try again.

Some of the values above are configurable via system properties. If you experience out of memory errors when people upload large files, you may want to reduce these limits further, using the following properties:

- atlassian.indexing.attachment.maxsize
- officeconnector.excel.extractor.maxlength
- officeconnector.textextract.word.docxmaxsize
- atlassian.indexing.contentbody.maxsize
- officeconnector.powerpoint.extractor.maxlength

Hierarchical File System Attachment Storage



The way attachments are stored changed significantly in Confluence 3.0. If you are upgrading from Confluence 2.10 or earlier see Upgrading Confluence for recommended upgrade paths, and read the version of the Hierarchical File System Attachment Storage page in our Confluence 3.0 documentation which provides more detail about migrating to the new file system structure.

Confluence stores attachments, such as files and images, in a file system. Confluence's attachment storage layout is designed to:

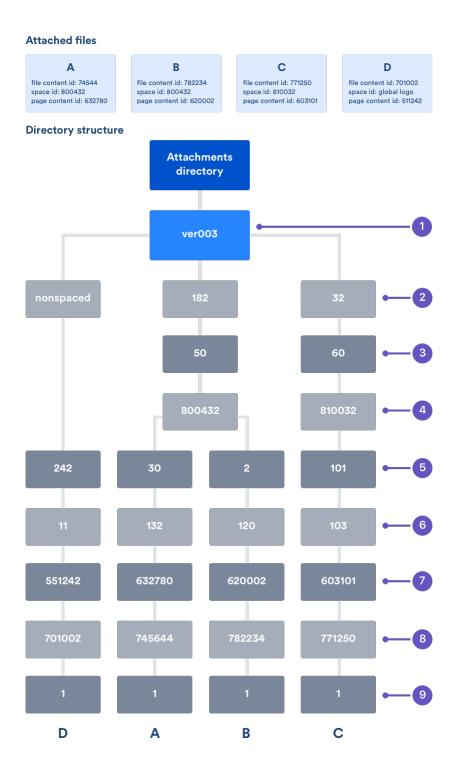
- 1. Limit the number of entries at any single level in a directory structure (as some file systems have a limit on the number of files that can be stored in a directory).
- 2. Partition attachments per space making it possible for a system admin to selectively back up attachments from particular spaces.

Attachments in Confluence have a number of identifying attributes: contentid of the file itself, the space idendicant ent id of the page the file is attached to. This means the file logically belongs to a piece of content which logically belongs in a space (not all content belongs to a space). For files within a space in Confluence, the directory structure is typically 8 levels, with the name of each directory level based on the following algorithm:

level	Derived From
1 (top)	Always 'ver003' indicating the Confluence version 3 storage format
2	The least significant 3 digits of the space id, modulo 250
3	The next 3 least significant digits of the space id, modulo 250
4	The full space id
5	The least significant 3 digits of the content id of the page the file is attached to, modulo 250
6	The next 3 least significant digits of the <i>content id</i> of the page the file is attached to, modulo 250
7	The full content id of the page the file is attached to
8	The full content id of the attached file
9	These are the files, named with the version number of the file, e.g. 1, 2, 6.

Themodulocalculation is used to find the remainder after division, for example 800 modulo 250 = 50.

An example:



To find the directory where attachments for a particular space are stored, go to <confluence url>/admin /findspaceattachments.jspand enter a space key. It will return the directory on the file system where attachments for that space are stored.

File D in the above diagram is stored in a slightly different structure. Files that are not conceptually within a space replace the level 2 - 4 directories with a single directory called 'nonspaced'. Examples of such files are the global site logo and attachments on unsaved content.

Extracted text files

When a text based file is uploaded in Confluence (for example Word, PowerPoint, etc), its text is extracted and indexed so that people can search for the content of a file, not just the filename. We store the extracted text so that when that file needs to be reindexed, we don't need to re-extract the content of the file.

The extracted text file will be named with the version number, for example 2.extracted_text, and stored alongside the file versions themselves (within level 8 in the explanation above). We only keep the extracted text for the latest version, not earlier versions of a file.

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-x.x.x.jar).
- The tables, columns and other attributes are likely to change with each major release of Confluence. To find the exact DDL of your Confluence site, please run a query after installation.

On this page:

- Database diagrams
- Database tables and references
- Authentication
- Content
- Clustering
- System information
- Spaces
- Appearance
- Miscellaneous

Database diagrams

We find that creating your own visualization of the Confluence database can be useful if you want to focus on particular tables or relationships. There are a number of tools you can use to create a visualization. Your own database tool may have options to do this.

View our visualization (excludes some tables, including ActiveObjects tables)

We used DbVisualizer. SeeViewing Table Relationships in the DbVis documentation to find out how it's done.

Database tables and references

Expand the link below to see a table of the primary and foreign keys for each table.

Note that Marketplace apps can also add tables to your database.

Primary Key Table Name	Primary Key Column Name	Foreign Key Table Name	Foreign Key Column Name	Foreign Key Name	Primary Key Name
AUDITREC ORD	AUDITRECO RDID	AUDIT_AFFECTE D_OBJECT	AUDITRECO RDID	FK_AFFECTED_OBJ ECT_RECORD	PRIMARY _KEY_D
AUDITREC ORD	AUDITRECO RDID	AUDIT_CHANGE D_VALUE	AUDITRECO RDID	FK_CHANGED_VAL UE_RECORD	PRIMARY _KEY_D
CONTENT	CONTENTID	ATTACHMENTDA TA	ATTACHMEN TID	FKJNH4YVWEN017 6QSVH4RPSRY2J	PRIMARY _KEY_6
CONTENT	CONTENTID	BODYCONTENT	CONTENTID	FKMBYIAYESFP1EI Q6GMOL3MK3YL	PRIMARY _KEY_6
CONTENT	CONTENTID	CONFANCESTO RS	DESCENDE NTID	FKLMHSIPSWOL8IM EQSG906IH62X	PRIMARY _KEY_6
CONTENT	CONTENTID	CONFANCESTO RS	ANCESTORID	FKSQB1AF9H7FVQ TGY73O8JDCUUE	PRIMARY _KEY_6
CONTENT	CONTENTID	CONTENT	PARENTCO MMENTID	FKAL6O8XWYPD4M DGID9B9NW1Q51	PRIMARY _KEY_6

CONTENT	CONTENTID	CONTENT	PARENTCCID	FKFIYHKA48C7E776 QJ90KLBPM9Q	PRIMARY _KEY_6
CONTENT	CONTENTID	CONTENT	PREVVER	FKK6KBB7SUQELO J82NX7XDCD803	PRIMARY _KEY_6
CONTENT	CONTENTID	CONTENT	PARENTID	FKOXTT893WEUJK YH0IICOXSM37V	PRIMARY _KEY_6
CONTENT	CONTENTID	CONTENT	PAGEID	FKWJYN6091Q3L1G L7BH143MA2A	PRIMARY _KEY_6
CONTENT	CONTENTID	CONTENT_LABEL	CONTENTID	FKI8CVAHSU6D2Y2 85VTRP4NHC3W	PRIMARY _KEY_6
CONTENT	CONTENTID	CONTENT_PERM _SET	CONTENT_ID	FK2BUUNK1HOR0I3 K0M3NT03HW1W	PRIMARY _KEY_6
CONTENT	CONTENTID	CONTENT_RELA TION	SOURCECO NTENTID	FKE2A00URQYXMY AJ3JOP48UB8QD	PRIMARY _KEY_6
CONTENT	CONTENTID	CONTENT_RELA TION	TARGETCO NTENTID	FKIPR00838MKLN69 9CIMD7RG17X	PRIMARY _KEY_6
CONTENT	CONTENTID	CONTENTPROP ERTIES	CONTENTID	FK3FLY21XFK13RQ H63TXW2T6K2V	PRIMARY _KEY_6
CONTENT	CONTENTID	EXTRNLNKS	CONTENTID	FK5V5LW9X88VM27 RVUBSC130NJY	PRIMARY _KEY_6
CONTENT	CONTENTID	IMAGEDETAILS	ATTACHMEN TID	FK2301QICIUQ6SC3 2JAJ8TYSG3S	PRIMARY _KEY_6
CONTENT	CONTENTID	LIKES	CONTENTID	FKBDOIWI70I7O3TC 7HPBU4VNLMY	PRIMARY _KEY_6
CONTENT	CONTENTID	LINKS	CONTENTID	FKN8MYCKO8FRER NE7BRH5NR1CSR	PRIMARY _KEY_6
CONTENT	CONTENTID	NOTIFICATIONS	CONTENTID	FK_NOTIFICATIONS _CONTENT	PRIMARY _KEY_6
CONTENT	CONTENTID	SPACES	SPACEDESC ID	FK7NDEWMRL3HQ CPWC8EYDN9MV8J	PRIMARY _KEY_6
CONTENT	CONTENTID	SPACES	HOMEPAGE	FKJ4CU5838AQCB W57WY7CKT0T7O	PRIMARY _KEY_6
CONTENT	CONTENTID	TRACKBACKLIN KS	CONTENTID	FK1TO6OMJL8NHE VCJBVPT3ED7NT	PRIMARY _KEY_6
CONTENT	CONTENTID	USERCONTENT_ RELATION	TARGETCO NTENTID	FKPWGL85A266IIE5 I0ADU8BDBCV	PRIMARY _KEY_6
CONTENT_ PERM_SET	ID	CONTENT_PERM	CPS_ID	FKDE5WL1CUR1SE 9281GC0DSAWTB	PRIMARY _KEY_BF
CWD_APP_ DIR_MAPPI NG	ID	CWD_APP_DIR_ GROUP_MAPPING	APP_DIR_M APPING_ID	FK_APP_DIR_GRO UP_MAPPING	PRIMARY _KEY_2A
CWD_APP_ DIR_MAPPI NG	ID	CWD_APP_DIR_ OPERATION	APP_DIR_M APPING_ID	FK_APP_DIR_MAPP ING	PRIMARY _KEY_2A

CWD_APPL ICATION	ID	CWD_APP_DIR_ GROUP_MAPPING	APPLICATIO N_ID	FK_APP_DIR_GRO UP_APP	PRIMARY _KEY_5
CWD_APPL ICATION	ID	CWD_APP_DIR_ MAPPING	APPLICATIO N_ID	FKSTEKJ41875RGS W8OTFFRAYHPL	PRIMARY _KEY_5
CWD_APPL ICATION	ID	CWD_APPLICATI ON_ADDRESS	APPLICATIO N_ID	FK_APPLICATION_ ADDRESS	PRIMARY _KEY_5
CWD_APPL ICATION	ID	CWD_APPLICATI ON_ATTRIBUTE	APPLICATIO N_ID	FK_APPLICATION_ ATTRIBUTE	PRIMARY _KEY_5
CWD_DIRE CTORY	ID	CWD_APP_DIR_ GROUP_MAPPING	DIRECTORY _ID	FK_APP_DIR_GRO UP_DIR	PRIMARY _KEY_AA
CWD_DIRE CTORY	ID	CWD_APP_DIR_ MAPPING	DIRECTORY _ID	FK_APP_DIR_DIR	PRIMARY _KEY_AA
CWD_DIRE CTORY	ID	CWD_DIRECTOR Y_ATTRIBUTE	DIRECTORY _ID	FK_DIRECTORY_AT TRIBUTE	PRIMARY _KEY_AA
CWD_DIRE CTORY	ID	CWD_DIRECTOR Y_OPERATION	DIRECTORY _ID	FK_DIRECTORY_O PERATION	PRIMARY _KEY_AA
CWD_DIRE CTORY	ID	CWD_GROUP	DIRECTORY _ID	FK_DIRECTORY_ID	PRIMARY _KEY_AA
CWD_DIRE CTORY	ID	CWD_GROUP_A TTRIBUTE	DIRECTORY _ID	FK_GROUP_ATTR_ DIR_ID	PRIMARY _KEY_AA
CWD_DIRE CTORY	ID	CWD_USER	DIRECTORY _ID	FK_USER_DIR_ID	PRIMARY _KEY_AA
CWD_DIRE CTORY	ID	CWD_USER_ATT RIBUTE	DIRECTORY _ID	FK_USER_ATTR_DI R_ID	PRIMARY _KEY_AA
CWD_GRO UP	ID	CWD_GROUP_A TTRIBUTE	GROUP_ID	FK_GROUP_ATTR_I D_GROUP_ID	PRIMARY _KEY_C
CWD_GRO UP	ID	CWD_MEMBERS HIP	CHILD_GRO UP_ID	FK_CHILD_GRP	PRIMARY _KEY_C
CWD_GRO UP	ID	CWD_MEMBERS HIP	PARENT_ID	FK_PARENT_GRP	PRIMARY _KEY_C
CWD_USER	ID	CWD_MEMBERS HIP	CHILD_USE R_ID	FK_CHILD_USER	PRIMARY _KEY_A3
CWD_USER	ID	CWD_USER_ATT RIBUTE	USER_ID	FK_USER_ATTRIBU TE_ID_USER_ID	PRIMARY _KEY_A3
CWD_USER	ID	CWD_USER_CR EDENTIAL_REC ORD	USER_ID	FK2RFDH2AP00B8 MHOLDSY1B785B	PRIMARY _KEY_A3
EXTERNAL _ENTITIES	ID	EXTERNAL_MEM BERS	EXTENTITYID	FKADLKFU6A03U8F 8BS82LM4QLG1	PRIMARY _KEY_6D
GROUPS	ID	EXTERNAL_MEM BERS	GROUPID	FK47K0FUDQNBNS BW0YW44UCSU2R	PRIMARY _KEY_7D
GROUPS	ID	LOCAL_MEMBERS	GROUPID	FKI71UOMCF4F9SE SIBDHSMFDBGH	PRIMARY _KEY_7D
KEYSTORE	KEYID	TRUSTEDAPP	PUBLIC_KEY _ID	FKM7N581Y7GROA 49TYGAPKMNFIV	PRIMARY _KEY_4D

LABEL	LABELID	CONTENT_LABEL	LABELID	FK91V3V5NEMR532 QQ4GLA9SJ9TF	PRIMARY _KEY_44
LABEL	LABELID	NOTIFICATIONS	LABELID	FK4TCCRJAMRJVM D2AOGL3HKLPFJ	PRIMARY _KEY_44
OS_GROUP	ID	OS_USER_GRO UP	GROUP_ID	FKM2O7638OJNKI0 5I3U0N5OEPOP	PRIMARY _KEY_DB
OS_USER	ID	OS_USER_GRO UP	USER_ID	FK6W5BWO7289K9 47EE5FWEC30JV	PRIMARY _KEY_E6
PAGETEMP LATES	TEMPLATEID	CONTENT_LABEL	PAGETEMPL ATEID	FK28KIFOKT21QD9 GES0Q0WV0FB9	PRIMARY _KEY_BC
PAGETEMP LATES	TEMPLATEID	PAGETEMPLATES	PREVVER	FK4WGWY1DQCI8R CWAD4TNQBGLT8	PRIMARY _KEY_BC
SPACES	SPACEID	CONTENT	SPACEID	FKLMWEU06NFT59 G7MW1I1MYORYS	PRIMARY _KEY_92
SPACES	SPACEID	NOTIFICATIONS	SPACEID	FKMQE1PHE52XW QC4HK4IB8P9EH6	PRIMARY _KEY_92
SPACES	SPACEID	PAGETEMPLATES	SPACEID	FK18A1D37PVQ2O9 HU5X3TPS97MX	PRIMARY _KEY_92
SPACES	SPACEID	SPACEPERMISSI ONS	SPACEID	FKBI3X723M8FBGO KO3S84F9ODDL	PRIMARY _KEY_92
TRUSTEDA PP	TRUSTEDAP PID	TRUSTEDAPPRE STRICTION	TRUSTEDAP PID	FKJOFK5643721EFT OW3NJWR73AA	PRIMARY _KEY_DDB
USER_MAP PING	USER_KEY	CONTENT	CREATOR	FK_CONTENT_CRE ATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	CONTENT	LASTMODIFI ER	FK_CONTENT_LAS TMODIFIER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	CONTENT	USERNAME	FK_CONTENT_USE RNAME	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	CONTENT_LABEL	OWNER	FK_CONTENT_LAB EL_OWNER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	CONTENT_PERM	CREATOR	FK_CONTENT_PER M_CREATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	CONTENT_PERM	LASTMODIFI ER	FK_CONTENT_PER M_LASTMODIFIER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	CONTENT_PERM	USERNAME	FK_CONTENT_PER M_USERNAME	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	CONTENT_RELA TION	CREATOR	FK_C2CRELATION_ CREATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	CONTENT_RELA TION	LASTMODIFI ER	FK_C2CRELATION_ LASTMODIFIER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	EXTRNLNKS	CREATOR	FK_EXTRNLNKS_C REATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	EXTRNLNKS	LASTMODIFI ER	FK_EXTRNLNKS_LA STMODIFIER	PRIMARY _KEY_13

USER_MAP	USER_KEY	FOLLOW_CONN	FOLLOWEE	FK FOLLOW CONN	PRIMARY
PING	JOEN_NET	ECTIONS	· OLLOWLL	ECTIONS_FOLLOW EE	_KEY_13
USER_MAP PING	USER_KEY	FOLLOW_CONN ECTIONS	FOLLOWER	FK_FOLLOW_CONN ECTIONS_FOLLOW ER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	LABEL	OWNER	FK_LABEL_OWNER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	LIKES	USERNAME	FK_LIKES_USERNA ME	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	LINKS	CREATOR	FK_LINKS_CREATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	LINKS	LASTMODIFI ER	FK_LINKS_LASTMO DIFIER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	LOGININFO	USERNAME	FK_LOGININFO_US ERNAME	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	NOTIFICATIONS	CREATOR	FK_NOTIFICATIONS _CREATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	NOTIFICATIONS	LASTMODIFI ER	FK_NOTIFICATIONS _LASTMODIFIER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	NOTIFICATIONS	USERNAME	FK_NOTIFICATIONS _USERNAME	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	PAGETEMPLATES	CREATOR	FK_PAGETEMPLAT ES_CREATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	PAGETEMPLATES	LASTMODIFI ER	FK_PAGETEMPLAT ES_LASTMODIFIER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	SPACEPERMISSI ONS	CREATOR	FK_SPACEPERMIS SIONS_CREATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	SPACEPERMISSI ONS	LASTMODIFI ER	FK_SPACEPERMIS SIONS_LASTMODIFI	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	SPACEPERMISSI ONS	PERMUSER NAME	FK_SPACEPERMIS SIONS_PERMUSER NA	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	SPACES	CREATOR	FK_SPACES_CREA TOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	SPACES	LASTMODIFI ER	FK_SPACES_LAST MODIFIER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	TRACKBACKLIN KS	CREATOR	FK_TRACKBACKLIN KS_CREATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	TRACKBACKLIN KS	LASTMODIFI ER	FK_TRACKBACKLIN KS_LASTMODIFIER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	USER_RELATION	SOURCEUS ER	FK_RELATION_U2U SUSER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	USER_RELATION	TARGETUSER	FK_RELATION_U2U TUSER	PRIMARY _KEY_13

USER_MAP PING	USER_KEY	USER_RELATION	CREATOR	FK_U2URELATION_ CREATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	USER_RELATION	LASTMODIFI ER	FK_U2URELATION_ LASTMODIFIER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	USERCONTENT_ RELATION	SOURCEUS ER	FK_RELATION_U2C USER	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	USERCONTENT_ RELATION	CREATOR	FK_U2CRELATION_ CREATOR	PRIMARY _KEY_13
USER_MAP PING	USER_KEY	USERCONTENT_ RELATION	LASTMODIFI ER	FK_U2CRELATION_ LASTMODIFIER	PRIMARY _KEY_13
USERS	ID	LOCAL_MEMBERS	USERID	FKRCUYOPTNAD1P OS41GP1B1F3PI	PRIMARY _KEY_4D4

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_member	Mapping the membership of users to groups.
cwd_direct	The user directories in your Confluence site. Examples of directories are the Confluence internal directory, or an LDAP directory.
cwd_applic ation	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
attach mentda ta	The binary data for attached files. This table is only used in older Confluence versions where Confluence was configured to store attachments in the database. Otherwise, attachments are stored in the local file system.
attach ments	Metadata for the files attached to Confluence pages.
bodyco ntent	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.

conten t_label	Arbitrary text labels for content.
label	The other half of the content_label system.
conten t_perm	Content-level permissions objects.
conten t_perm _set	A one-to-many mapping for content items and their permissions, with added metadata.
pagete mplates	The back end of the templates feature.
likes	The pages and other content liked by a particular user.
follow _conne ctions	A mapping of users who are following other users.

Clustering

The following table contains information about clustered Confluence sites.

Table	Description
clust ersaf ety	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.
journ alent ry	The journal service keeps the search index in synch across each Confluence node.

System information

These tables store data related to the status and configuration of the Confluence site.

Table	Description
confver sion	Used by the upgrade system to determine what to expect from the database, so as to negotiate upgrades.
plugind ata	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.
diagnos ticaler ts	The diagnostics tool continuously checks for symptoms or behaviours that we know may contribute to problems in your site. Events are stored for a limited amount of time in this table.

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 customize Velocity layouts.

Miscellaneous

This section includes other tables worth commenting on.

Table	Description
os_pro pertye ntry	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.
extrnl nks	Referral links.
hibern ate_un ique_k ey	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.
indexq ueueen tries	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.
keysto re	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).
notifi cations	Stores page- and space-level watches.
trackb acklin ks	Trackback links.
confan cestors	Used to speed up permissions checks, by allowing quick lookup of all a page's ancestors.

Finding Unused Spaces or Pages

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?



If you have Confluence Data Center, you can view detailed analytics for spaces and pages. See Analytics for more information.

If you have Confluence Server, you'll need to query the database directly to find out this information.

The following query identifies the last date on which content was modified in each space within a single Confluence instance:

```
SELECT spaces.spacename, MAX(content.lastmoddate)
FROM content, spaces
WHERE content.spaceid = spaces.spaceid
GROUP BY spaces.spacename;
```

It returns a list of space names, and the last date and time at which any content was added or changed.

Alternatively, this query identifies spaces where the content hasn't changed since a specified date:

```
SELECT spaces.spacename, spaces.spacekey
FROM content, spaces
WHERE content.spaceid = spaces.spaceid
GROUP BY spaces.spacename, spaces.spacekey
HAVING MAX(content.lastmoddate) < '2006-10-10';</pre>
```

The result is a simple list of space names and corresponding space keys.

Similarly, to identify pages where content hasn't changed since a specified date:

```
SELECT c.title,
c.creationdate,
c.lastmoddate
FROM content c
WHERE c.prevver IS NULL
  AND c.contenttype = 'PAGE'
 AND c.content_status = 'current'
GROUP BY c.title, c.creationdate, c.lastmoddate
HAVING MAX(c.lastmoddate) < '2006-10-10';
```

The result will provide a list of page titles, created date, and last modified date. Change thelastmoddatevalue as per your needs.

Data Import and Export

Confluence administrators and users can import data into Confluence from a number of sources. The permissions required differ, depending on the scope of the import. See Import Content Into Confluence.

You can also export Confluence content to various formats. See Export Content to Word, PDF, HTML and XML.

Related pages:

- Managing Confluence Data
- Confluence Administrator's Guide

Import a Text File

Confluence allows you to import text files from a directory on the Confluence server, and convert them into Confluence pages. Each file is imported as a separate Confluence page with the same name as the file.



- The text file may contain plain text, HTML orConfluence storage format
- You need to be part of theconfluence-administratorsg roup or a System Administrator to import text files
- You can import pages from disk into site spaces, but not into personal spaces
 - Please seeSpacesfor information about differences between site spaces and personal spaces.

Related pages:

- Import Content Into Confluence
- Site Backup and Restore

To make sure Confluence maintains the formatting of the text document, add to the beginning and to the end. This will let Confluence know that it should treat the text as pre-formatted.

If you're working in a Unix-like environment, you can add the opening and closing tags to all files in a particular directory by following these steps:

- 1. Go to the directory containing the files
- 2. Run the following command in the terminal:

```
for i in $(ls); do echo "" >> m$i; cat $i >> m$i; echo "" >> m$i; mv m$i $i; done
```

To import text files:

- 1. Go to the space and choose Space tools > Content Tools from the bottom of the sidebar
- 2. Choose Import.
- 3. Type the directory path into the **Import directory** box.
- 4. Select **Trim file extensions** to remove file extensions from the page titles when converting the files to Confluence pages
 - The Confluence pages will take their titles from the files' names (including their extensions). To avoid having page titles with a suffix like '.txt' check this box.
- 5. Select **Overwrite existing pages** if you want to replace existing Confluence pages with the same title with the one you're importing.
- 6. ChooseImport.

Screenshot:	Importing	toyt files
Screenshor	111100111110	iexi illes

You can use this action to import text files from a directory on the Confluence server.		
These text files become	ne pages in Confluence, with the following features:	
 The page title is taken from the filename The content is the entire page body 		
Import directory		
Trim file extensions		
Overwrite existing pages		
	Import Cancel	

Auditing in Confluence

The audit log allows administrators to look back at changes that have been made in your site. This is useful when you need to troubleshoot a problem or if you need to keep a record of important events, such as changes to global permissions.

The feature is available in both Server and Data Center, however its scope is more extensive in Data Center.

With a Data Center license, Space admins can also view the audit log for their specific space.

On this page:

- Main differences between auditing in Server and Data Center
- View the audit log
- View the space audit log (Data Center)
- Search and filter the audit log
- Edit log settings
- Access the audit log file (Data Center)
- Integrate with external software (Data Center)
- Audit log and migration
- Auditing and the REST API

Main differences between auditing in Server and Data Center

The auditing feature works differently in Confluence Server and Confluence Data Center.

Functionality	Available in Server	Available in Data Center
Coverage areas	✓ Yes (fewer coverage areas than Data Center)	⊘ Yes
Selecting coverage areas	⊗ No (only Base or Off)	
Setting database log retention	✓ Yes	
Storing audit logs in two locations	⊗ No	
Integrating with 3rd party monitoring tools	⊗ No	⊘ Yes
Exporting latest 100,000 results	✓ Yes	
Filter by category and summary	⊗ No	
Exporting filtered results	✓ Yes	
Space level audit log	⊗ No	⊘ Yes

View the audit log

To view the global audit log in Confluence:

- 1. Go to Configuration
- 2. SelectAudit log
- 3. Select an event to expand it and see details.

Different details will be shown depending on the event itself. These can include:

- IP address IP address of the user who performed the action. This is not recorded for systemgenerated events.
- Load balancer/proxy IP address IP address of the load balancer or proxy server that forwarded the request.
- Node ID unique ID of the cluster node where the action was performed.
- Method depending on how the action was performed, this will be either Browser (end user) or System (system process).

View the space audit log (Data Center)

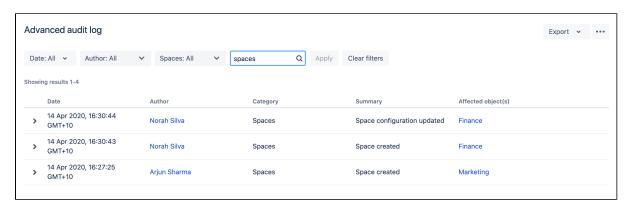
System admins, Confluence admins and space admins can also access audit logs for a specific space, if they have permission to administer that space.

The space audit log records events related to space permissions and configuration, user actions within the space, and some events related to space security (for example, events related to accessing and granting permissions to restricted pages with a particular space).

To view the audit log for a specific space, go to **Space tools** > **Audit log**.

Search and filter the audit log

You can search the log by keyword, and narrow your results by date, author, and space. If you have a Data Center license you can also filter by category and summary.





Your query can be up to 100 characters long. To speed up the search, we only search the most recent 1 million events. After this search is performed, you can choose to run a full database search. If you have a large or busy Confluence site, running a full search can take a while.

Can't find a specific event?

Changing coverage level changes the individual events that are logged. If you can't find a specific event, it might be because coverage level was changed, and these events were not logged for a period of time. Check the audit log configuration events to determine if this might be the case.

Edit log settings

In the audit log settings you can decide how long you want to retain the logged events in the database, and the areas from which you want to collect the logs.

Update database retention

The database retention is limited by the retention period, with a maximum of 10 million records.

To update the database retention period:

- 1. Select more options *** > Settings.
- 2. Enter the period of time. This can be in days, months or years.
- 3. SelectSave.

If you choose a long retention period, it can affect the size and performance of your database. Learn more aboutsetting an optimal retention periodfor your Confluence instance.

If you decide to lower the retention period, all the events that exceed the newly set period will be deleted, and disappear from the page. It's a good idea to create a backup before you lower the retention period.

If you migrated from a previous Confluence version, your default retention period is 20 years. If you have a new Confluence installation, its 3 years.

Select events to log

The events that are logged are organized in categories that belong to specific coverage areas.

For example, import and export-related events are logged in the Import/Export category, that belongs to the **L ocal configuration and administration** coverage area. For all coverage areas and events logged in each area, see Audit log events in Confluence.

To adjust the coverage:

- Go to more options *** > Settings.
- 2. In the **Coverage level** drop-down, select the level to log the events you need, or **Off** to stop collecting events from a particular area.

Coverage level definitions

Coverage levels reflect the number and frequency of events that are logged. Some coverage levels are only available with a Data Center license.

Coverage level	Definition
Off	Turns off logging for this coverage area.
Base	The lowest level of coverage. Logs only the core events. Base coverage provides a minimum level of insight into your sites activity. If you have a Confluence Server license, this is the only coverage level available.
Advanced (Data Center only)	Logs all the events covered in Base, plus additional events. Advanced coverage provides a more detailed record of your sites activity.
Full (Data Center only)	The highest level of coverage available. Logs all events in Base and Advanced. Depending on your site's activity, setting your coverage level to Full can generate a large volume of events, which can impact your database and disk space.

Export the audit log

You can export up to 100,000 latest or filtered events as a CSV file. If you have more than 100,000 events, only the 100,000 newest events are included in the export.

To export the audit log:

- 1. Go to Audit log, then choose Export.
- 2. Select to export the latest 100,000 or filtered results.
- 3. Confirm by clicking **Export**.

Space admins can also export from the space level audit log.

Access the audit log file (Data Center)

For Confluence Data Center, each node has its own log, which can be found in the local-home>/log
/auditdirectory. The log is stored as a JSON file.

Confluence creates a new log file every 24 hours, or once the current one reaches 100 MB, whichever occurs first. For more details on log rotation, see Audit Log Integrations in Confluence.

Change the audit log file retention

You can choose how many audit log files to store in the local home directory on each node. By default we store 100 files. Make sure you've provisioned enough disk space for these files, especially if you have set the logging level to Advanced or Full.

To change the file retention setting:

- 1. Go to General Configuration > Audit log.
- 2. SelectSettings.
- 3. Enter the maximum number of files to be stored and selectSave.

Once a node reaches the log file retention limit, the oldest one is deleted. If you need to keep these logs, for example for compliance purposes, you may want to manually back up the files in this directory on a regular basis, or send them to a third party logging platform. See Audit Log Integrations in Confluence.

Integrate with external software (Data Center)

You can use the log file to integrate with third-party tools such as ELK, Splunk, Sumologic, and Amazon CloudWatch. For more information on integrations, see Audit Log Integrations in Confluence.

Audit log and migration

Migrate database

If you have more that 10 million events stored in your database, and you move to a new database, only the latest 10 million will be migrated, and the remaining data will be removed.

To have access to your older events, you can create a backup before you migrate and access the data in the backup.

Migrate from a previous Confluence version

Migrating audit log records can take a while, depending on the size of the audit log and your database.

Auditing and the REST API

The audit log can also be accessed via the REST API.

Audit Log Events in Confluence

This page outlines the auditing events available in Confluence Server and Data Center, and which events fall into each coverage level.

For more information about how auditing works, see Auditing in Confluence.

On this page

- Definitions
- Global configuration and administration
- User management
- Permissions
- Local configuration and administration
- Security
- End user activity

Definitions

Coverage area

A coverage area is a grouping of events related to a similar theme.

Coverage area	Definition
Global configuration and administration	Logs instance or system admin activity around instance administration or configuration such as platform changes or upgrades to global settings.
User management	Logs activity around users, groups, memberships, and roles, such as adding and removing users and groups.
Permissions	Log activity around local and global permissions and configurations such as changing to anonymous access or updating group permissions.
Local configuration and administration	Logs admin activity around spaces, such as creating or deleting a space.
Security	Logs user actions related to security such as authentication, or granting access to a restricted page.
End user activity (Da ta Center only)	Logs end user activity on your site, such as user actions on a page (creating, editing, commenting), searching, or viewing pages.

Category

A category is a grouping of related events. Categories can belong to multiple coverage areas.

Category names change over time. You may find your audit log contains some categories not described on this page. These are usually associated with events logged prior to Confluence 7.5.

Coverage level

Coverage levels allow you to control which events are logged. Some levels are only available with a Data Center license.

Coverage level	Definition
Off	Turns off logging for this coverage area.

Base	The lowest level of coverage. Logs only the core events. Base coverage provides a minimum level of insight into your sites activity. If you have a Confluence Server license, this is the only coverage level available.
Advanced (Data Center only)	Logs all the events covered in Base, plus additional events. Advanced coverage provides a more detailed record of your sites activity.
Full (Data Center only)	The highest level of coverage available. Logs all events in Base and Advanced. Depending on your site's activity, setting your coverage level to Full can generate a large volume of events, which can impact your database and disk space.

Global configuration and administration

Category: Global administration

Coverage level	Events logged
Base	Mail server created Mail server deleted Mail server edited Max cache size changed Licence updated Global settings changed Site import Site export Site logo changed Favicon changed Favicon reset to default Custom stylesheet added Custom stylesheet removed Theme changed Custom decorator modified Color scheme modified Color scheme type changed Security configuration updated Global retention rule changed Retention rule exemption added Retention rule exemption removed

Banner configuration changed Collaborative editing mode changed Synchrony restarted Mail queue flushed Mail queue: error queue re-sent Mail queue: error queue deleted Allowlist turned on Allowlist turned off Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job run manually Application link created Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed		
Banner configuration changed Collaborative editing mode changed Synchrony restarted Mail queue flushed Mail queue: error queue re-sent Mail queue: error queue deleted Allowlist turned on Allowlist turned off Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job run manually Application link created Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed	Advanced	
Collaborative editing mode changed Synchrony restarted Mail queue flushed Mail queue: error queue re-sent Mail queue: error queue deleted Allowlist turned on Allowlist turned off Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job edited Scheduled job run manually Application link created Application link created Application link removed Rate limiting exemption added Rate limiting exemption edited		Read-only mode configuration changed
Synchrony restarted Mail queue flushed Mail queue: error queue re-sent Mail queue: error queue deleted Allowlist turned on Allowlist turned off Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption edited		Banner configuration changed
Mail queue: error queue re-sent Mail queue: error queue deleted Allowlist turned on Allowlist turned off Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job disabled Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption edited		Collaborative editing mode changed
Mail queue: error queue re-sent Mail queue: error queue deleted Allowlist turned on Allowlist turned off Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job disabled Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption edited		Synchrony restarted
Mail queue: error queue deleted Allowlist turned on Allowlist turned off Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job disabled Scheduled job run manually Application link created Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption edited		Mail queue flushed
Allowlist turned on Allowlist turned off Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job disabled Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption removed Rate limiting exemption edited		Mail queue: error queue re-sent
Allowlist turned off Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job disabled Scheduled job run manually Application link created Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption edited		Mail queue: error queue deleted
Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption edited		Allowlist turned on
Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed		Allowlist turned off
Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		Allowlist URL added
Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption edited		Allowlist URL removed
Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		·
Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption edited		
Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption edited		
Scheduled job enabled Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		
Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		, , ,
Scheduled job edited Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		•
Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		
Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		
Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		
Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		''
Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		
Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited		
Rate limiting exemption removed Rate limiting exemption edited		
Rate limiting exemption edited		
CDN configuration		
obit comigaration		CDN configuration
Full No events	Full	No events

Category: Apps

Coverage level	Events logged
Base	App installed App uninstalled App enabled App disabled App module enabled App module disabled
Advanced	No events
Full	No events

Category: Page templates

Coverage level	Events logged
Base	Page template updated Page template created Page template deleted
Advanced	No events
Full	No events

User management

Category: Users and groups

Coverage level	Events logged
Base	User created User deleted User renamed User details updated User requested password reset Group created Group deleted User added to group User removed from group User directory created User directory updated
Advanced	User was invited to join site
Full	No events

Permissions

Category: Permissions

Coverage level	Events logged
Base	Space permission removed Space permission added Global permission removed Global permission added
Advanced	No events
Full	No events

Local configuration and administration

Category: Pages and blogs

Coverage level	Events logged
Base	Page hierarchy copy started Page hierarchy delete started
Advanced	No events
Full	No events

Category: Import / export

Coverage level	Events logged
Base	Space import Space export Space exported to PDF
Advanced	No events

Full No events	Full
----------------	------

Category: Spaces

Coverage level	Events logged
Base	Space created Space deleted Space archived Space unarchived Space trash emptied Space logo uploaded Space logo enabled Space logo disabled Space logo deleted Unknown update to space logo Space configuration updated
Advanced	No events
Full	No events

Security

Category: Auditing

Coverage level	Events logged
Base	Audit log search performed Audit log exported Audit log configuration updated
Advanced	No events
Full	No events

Category: Authentication

Coverage level	Events logged
Base	No events
Advanced	Secure admin access granted Secure admin access request failed Secure admin access dropped User authorized external application access via OAuth tokens User de-authorized external application access via OAuth token User login failed*
Full	User login successful User logout

Category: Users and groups

Coverage level	Events logged	
Base	No events	
Advanced	Forgot password feature triggered Forgot password feature triggered for unknown user	

Full No events	
----------------	--

Category: Security

Coverage level	Events logged
Base	No events
Advanced	User tried to access restricted page User requested access to restricted page User requested access to restricted blog Owner of the page authorized access Owner of the blog authorized access
Full	No events



(i) *We only track **User login failed** events if the authentication does not involve a redirect to an external identity provider. If a user tries to log in using SSO and fails, this event will not be logged. Most identity providers track these events in their own audit logs.

End user activity

DATA CENTER ONLY

Category: Pages and blogs

Coverage level	Events logged	
Base	No events	
Advanced	Blog post edit restriction added Blog postedit restriction removed Blog postview restriction added Blog postview restriction removed Blog post deleted Blog post purged from trash Blog post version deleted Page view restriction added Page view restriction removed Page edit restriction added Page edit restriction removed Page restored Page deleted Page version deleted Page purged from trash Attachment deleted Attachment version deleted Comment deleted Page moved Blog post moved Page shared Blog post shared	

Full	Page created Page edited Blog post created Blog post edited Comment created Comment edited Inline comment created Inline comment edited Attachment downloaded Attachment uploaded Search performed
------	--

Note: The **Search performed** event records the search terms entered in search, advanced search, and in search macros (such as Livesearch and Page Tree Search). If you don't want to collect this data you can disable this event using theaudit.log.search.disabledsystem property.

Category: Import / export

Coverage level	Events logged
Base	No events
Advanced	Page exported to PDF Blog post exported to PDF Page exported to Word Blog post exported to Word
Full	No events

Category: Data pipeline

Coverage level	Events logged
Base	Full data export cancelled Full data export triggered Unauthorized full data export triggered Full data export failed
Advanced	No events
Full	No events

Audit Log Integrations in Confluence

Confluence Data Center writes audit logs to the database and log file. By itself, the log file saves you the effort of periodically exporting your audit logs from the database for long-term storage. However, the main purpose of the file is to easily integrate ConfluenceData Center to a third-party logging platform.

On this page:

- Event coverage and log retention
- Log file details
- Integrating with logging agents

Event coverage and log retention

The Audit log settings menu controls the coverage of audit logs in both database and log file. However, this menu does not control the log file's retention period.

The log file's retention is ultimately controlled bylog rotation. We use basic log rotation to manage the volume of logs. We automatically archive the audit logfile when:

- the node's time reaches 12:00 midnight, or
- the audit log file reaches100MB.

Once a node reaches the log file retention limit, the oldest one is deleted. By default the limit is 100 log files (the current audit log file + 99 archives). Make sure you allocate enough disk space for these log files on each application node. For the default setting of 100 files, you should allow 10GB.

Log file details

ConfluenceData Center writes audit logs in real time to the home directory. Specifically, these logs are written to the audit logfile. On clusteredConfluenceData Center deployments, each application node will produce its own logfile in itslocalhome directory.

Location

To integrate the audit logfile with a third-party logging platform, you'll need to know its exact location. This may vary, depending on how you configured your home directory. For more information about the local home directory, seeConfluence Home and other important directories).

On a clusteredConfluenceData Center deployment, the audit logfile's directory should be the same on all nodes.

SeeCloudWatch Logs Agent Referencefor more information. If you want to see how we automate this via Ansible, check out our deployment playbooks onhttps://bitbucket.org/atlassian/dc-deployments-automation/src/master/.

File name

The audit log file name uses the following naming convention:

YYYYMMDD-XXXXX.audit.log

Thexxxxxportion is a 5-digit number (starting with 00000) tracking the number of audit log files archived in the same day (YYYMMDD). For example, if there are 5 archived log files today (January 1, 2020), then:

- the oldest archived log file is 20200101.00000.audit.log
- the current audit log file is 20200101.00005.audit.log

Format

Each audit log is written as a JSON entry to the audit log file. Every line in the filerepresents a single event, allowing you to useregular expressions to do simple searches if needed.

Integrating with logging agents

Most enterprise environments use a third-party logging platform to aggregate, store, and otherwise manage logs from all hosts. Logging platforms like AWS CloudWatch and Splunkuse agents to collect logs from every host in the environment. These agents are installed on each host, collecting local logs and sending them back to a centralized location to be aggregated, analyzed, audited, and/or stored.

If your logging platform uses agents this way, you can configure each node's agent to monitor the audit log file directly. Logging agents from most major platforms (including AWS CloudWatch, Splunk, ELK, and Sumo Logic) are compatible with the audit logfile.

Amazon CloudWatch Agent

We provideQuick Starts for ConfluenceData Centerfor easy deployments on AWS. This Quick Start lets you deploy ConfluenceData Center along with an Amazon CloudWatch instance to monitor it.

To set up Amazon CloudWatch, use the **Enable CloudWatch Integration** parameter's default setting (namely, Metrics and Logs). The Quick Start will then configure the Amazon CloudWatch Agentto collect the logs from each node's audit log files. The agent will send these logs to a separate log group namedconfluence-<aws-stack-id>-audit.

Our Quick Start also sets up a default dashboard to help you read the collected data, including logs from each audit logfile.Refer toWorking With Log Groups and Log Streamsfor related information.

Manual configuration

If needed, you can also manually configure the Amazon CloudWatch agent to collect the audit log files. To do this, set the followingparameters in the Agent Configuration File:

- file: set this toto<local home directory>/log/audit/*.Don't forget to set the absolute path to thehome directory.
- log_group_nameandlog_stream_name:use these to sendConfluence Data Center's audit logs to a specific log group or stream.

Splunk Universal Forwarder

For Splunk Enterprise or Splunk Cloud, you can use the Splunk Universal Forwarderas your logging agent. This will involve installing the universal forwarder on each application node.

You'll also need to define each node's audit log directoryas one of the forwarder's inputs. This will set the forwarder to send all logs from the audit log directoryto a pre-configuredreceiver. One way to define the forwarder's inputs is through the Splunk CLI. For Linux systems, use the following command on each application node:

./splunk add monitor <local home directory>/log/audit/*audit.log

Refer to the following links for detailed instructions on configuring the Splunk Universal Forwarder on each node:

- How to forward data to Splunk Enterprise
- How to forward data to Splunk Cloud

Filebeat (for the ELK stack)

Within the ELK stack, you can use the Filebeat plugin to collect logs from each node's audit log files. Each time a log is written to the current audit log file, Filebeat will forward that log to Elasticsearch or Logstash.

To set this up,install Filebeatfirston each application node. Then, set the audit log filedirectory as aFilebeat input. To do that, add its directory as apathinthefilebeat.inputssection of each node'sfilebeat.ymlc onfiguration file. For example:

```
filebeat.inputs:
    type: log
    enabled:true
    paths:
    <local home directory>/log/audit/
```

Sumo Logic installed collectors

If you have a Sumo Logic instance, you canuseinstalled collectorsto collect logs from each node's audit log files. To do this, install a collectoron each node first. Then, add<local home directory>/log/audit/*as aLocal File Sourceto each node's collector.

Set retention rules to delete unwanted data

This feature is available with a Confluence Data Center license.

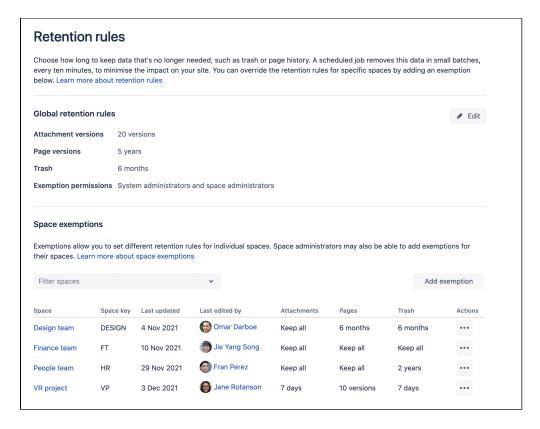
When you edit a page or attached file, Confluence stores the previous content so you can restore it if you need to. Over time these historical versions start to add up, increasing the size of your database and attachments directory. It's not uncommon for some pages to have hundreds of historical versions, or for a space to have hundreds of items in the trash.

On this page:

- Retention rule criteria
- How versions are deleted
- Define your retention strategy
- Change the global retention rules
- Add a space exemption
- Remove a space exemption
- Allow space administrators to manage exemptions
- View the retention rules in a space
- Change the retention rules in a space
- Revert back to global retention rules in a space
- Considerations for administrators

Retention rules allow you to automatically delete historical versions of pages and attachments, and purge deleted items from the trash. You can:

- set global rules that will apply to all spaces
- define exemptions for spaces that have special requirements, and need different rules
- allow space administrators to set rules for their spaces.



Screenshot showing global retention rules administration screen

Retention rule criteria

You can set a retention rule for:

- historical page versions (pages only, you can't set a retention rule for blog posts)
- historical attachment versions, and
- · items in the trash.

The criteria you can use to determine what should be deleted is outlined below.

Rule criteria	Page and Attachment versions	Trash
Keep all (default)	Historical versions will not be automatically deleted.	Items will not be automatically purged from the trash
Keep by age	Keep versions for a specific amount of time. For example you could set it to automatically delete any version older than 2 years.	Keep deleted items in the trash for a specific amount of time. For example you could set it to purge any item deleted more than 3 months ago.
Keep by number	Keep a maximum number of versions. For example you could set it to keep the 5 most recent versions, and automatically delete any earlier versions.	Not applicable for trash.

It's important to note that for page and attachment versions, the latest version is never deleted, only the history. Retention rules never prevent people from creating new versions.

How versions are deleted

A scheduled job will permanently delete any items that don't the meet retention rules. This "soft" job runs every 10 minutes and deletes items in small batches (of about 3000 items) to ensure there's no performance impact to your site. When you first set a rule, the job may need to run quite a few times before all items that don't meet that rule are deleted.

If you need to delete versions more quickly, you can manually run the "hard" job, which will delete all items that don't meet the rules in one cycle. This can have a performance impact however, so you might want to only run thehard job when Confluence is less busy.

See Scheduled jobs to learn more about these jobs and how to run or disable them.

Each version deleted or item purged from the trash is written to the audit log if the End user activity coverage area is set to **Advanced** or higher.

Define your retention strategy

Before you set any rules, it's important to define your retention strategy. Historical versions and trash that don't meet the retention rule criteria will be permanently deleted, and can't be restored, so it's essential you get it right.

Global rules vs space exemptions

There are two approaches you can take:

- Set your global rule to keep all, and use space exemptions to target individual spaces that can be cleaned up more aggressively.
- Add space exemptions for the spaces where you need to retain history and trash, and set these to keep all, then set a global rule to clean up all remaining spaces.

The order that you set your rules is important - once you set a global rule, it will start deleting almost immediately. Make sure all your exemptions are in place first.

Here's some example scenarios to help you think through your strategy.

Mia is the administrator of a large Confluence site that has been active for about 8 years. Backups and upgrades have become increasingly difficult due to the database and attachment directory now weighing in at over 10 terabytes.

To reduce Confluence's footprint, the Mia:

- Discusses various options with stakeholders, and decides to put in place an aggressive global retention rule.
- Communicates the plan to stakeholders in their organisation, including the date the global rule will be set.
- Makes a note of the current size of the database and attachments directory
- Sets the following global retention rules
 - Page versions keep by age 2 years
 - Attachment versions keep by age 1 year
 - Trash keep by age 6 months

Within a few days, Mia observes a reduction in the size of the database, and a huge 1tb reduction in the size of the attachments directory, both of which contribute to a noticeable decrease in backup time.

Omar administers the Confluence site for an insurance company. Omar would like to clean up the site, but knows that some teams need to keep detailed records of their work, for auditing and compliance purposes.

To clean up unnecessary data Omar:

- Discusses various retention rule options with the compliance team. They decide that the best approach is to keep everything by default, but empower individual teams to make decisions about their own spaces.
- Leaves the global retention rules as
 - Page versions keep all
 - Attachment versions keep all
 - Trash keep all
- Sets the exemption permissions to allow system administrators and space administrators to manage retention rules.
- Communicates to team leads that they can set their own rules in Space Tools, and points them to a page prepared by the compliance team which outlines what data must be kept, and for how long.

Within a few days, Omar sees exemptions appearing in the space exemptions list, as several teams add their own retention rules.

Fran administers several different Confluence instances that have sprung up across the organisation as a result of acquisitions over the years. They're in the process of consolidating all their sites, and plan to move some of them to Confluence cloud over the next year or so.

Fran wants this process to be as smooth as possible, and doesn't really want to migrate a huge amount of unnecessary historical data. Fran:

- Uses Analytics to make a short list of spaces that have not been viewed or edited in the last 6
 months.
- Informs space owners that much of the version history will soon be removed.
- Leaves the global retention rules as:
 - Page versions keep all
 - Attachment versions keep all
 - Trash keep all
- Adds an exemption for each space that can be cleaned up, and sets a very agressive retention rule:
 - Page versions keep by number 5 versions
 - Attachment versions keep by number 2 versions
 - Trash keep by age 1 month

When the time comes to start migrating spaces to their new cloud instance, Fran prioritises the cleaned up spaces, and finds the migration happens more quickly as there is significantly less data to transfer.

Delegate responsibility to space administrators

You'll need to decide whether to allow space administrators to manage retention rules for their spaces. This allows the administrators most familiar with the content to make decisions about how long to keep historical versions and items in the trash.

See Allow space administrators to manage exemptions.

Change the global retention rules



Check the considerations for administrators before changing retention rules.

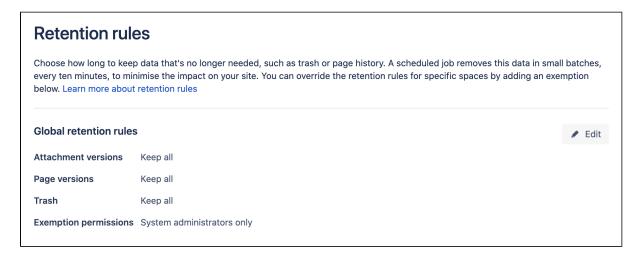
You need system administrator global permissions to do this.

To change a global retention rule:

- 1. Go to General Configuration Retention rules.
- 2. Select Edit under Global retention rules.
- Under Page versions selectKeep all, Keep by number, or Keep by age, and then enter a value (if required).
- 4. Under **Attachment versions**select**Keep all**, **Keep by number**, or **Keep by age**, and then enter a value (if required).
- 5. Under **Trash**, select **Keep all**, or **Keep by deleted date**, and then enter a value (if required).
- 6. Choose whether to allow space administrators to set retention rules for their space.

7. Save your changes.

These rules will apply to all spaces in your site, unless an exemption has been added for a particular space. This includes personal spaces and archived spaces.



Screenshot showing global retention rules

Add a space exemption

For the situations where you don't want a global retention rule to apply, you can add an exemption to define different rules for a space. For example you may set the global retention rule to delete all page versions older than 2 years, but add an exemption for your Finance and HR spaces, so that all historical versions are kept. Alternatively you could set the global retention rule to keep all, and use exemptions to clean up particular spaces.

You need system administrator global permissions to do this.

To add an exemption for a space:

- 1. Go to O Seneral Configuration Retention rules.
- 2. Select Add exemption.
- 3. Select a space.
- 4. Under Page versions selectKeep all, Keep by number, or Keep by age, and then enter a value (if required).
- Under Attachment versionsselectKeep all, Keep by number, or Keep by age, and then enter a value (if required).
- 6. Under Trash, select Keep all, or Keep by deleted date, and then enter a value (if required).
- 7. Choose whether retention rules for this space can be managed by space administrators.
- 8. Save your changes.

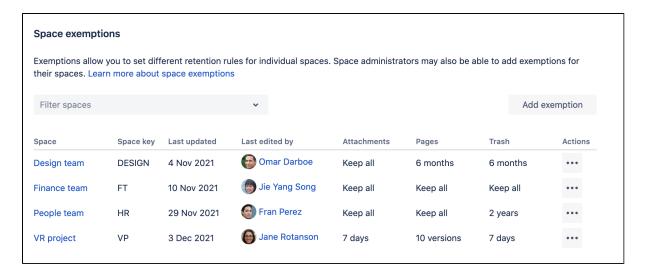
The global retention rules will no longer apply to this space.

Remove a space exemption

To remove an exemption:

- 1. Go to O-> General Configuration Retention rules.
- 2. In the space exemptions list, locate your space and choose ... > Remove.

The global retention rules will now apply to this space.



Screenshot showing the global retention rules screen with exemptions listed for particular spaces.

Allow space administrators to manage exemptions

Often system administrators don't have detailed knowledge of the type of content that is stored in each space, so in big sites, it can be useful to allow space administrators to manage the retention rule exemptions for their spaces.

To allow space administrators to manage retention rules:

- 1. Go to O Seneral Configuration Retention rules.
- 2. Select Edit under Global retention rules.
- 3. Under Exemption permissions, select System administrators and space administrators.
- 4. Save your change.

This applies to all spaces, unless there is an exemption which sets different exemption permissions.

To allow space administrators to manage retention rules for a specific space:

- 1. Go to O-> General Configuration Retention rules.
- 2. Add or edit an existing Exemption
- 3. Under Exemption permissions, select either System administrators and space administrators.
- 4. Save your change.

Space administrators will be able to change the retention rules, as described below.

View the retention rules in a space

Space administrators can always see the retention rules that apply to their space, even if they don't have permission to edit the rules.

To view the retention rules in a particular space:

- 1. Go to the space and choose Space tools>Content Tools from the bottom of the sidebar
- Select the Retention rules tab.

Space administrators will only be able to edit the retention rules for their space if a system administrator has allowed space administrators to manage exemptions.

Change the retention rules in a space

You need **space admin** space permissions to do this. The edit button will be disabled if your system administrator has not allowed space administrators to manage exemptions.

To change the retention rules in a particular space:

- 1. Go to the space and choose Space tools > Content Tools from the bottom of the sidebar
- 2. Select the Retention rules tab
- 3. Select Edit.
- 4. Select Use retention rules defined in this space from the Exemption permissions drop down.
- Under Page versions selectKeep all, Keep by number, or Keep by age, and then enter a value (if required).
- Under Attachment versionsselectKeep all, Keep by number, or Keep by age, and then enter a value (if required).
- 7. Under Trash, select Keep all, or Keep by deleted date, and then enter a value (if required).
- 8. Save your change.

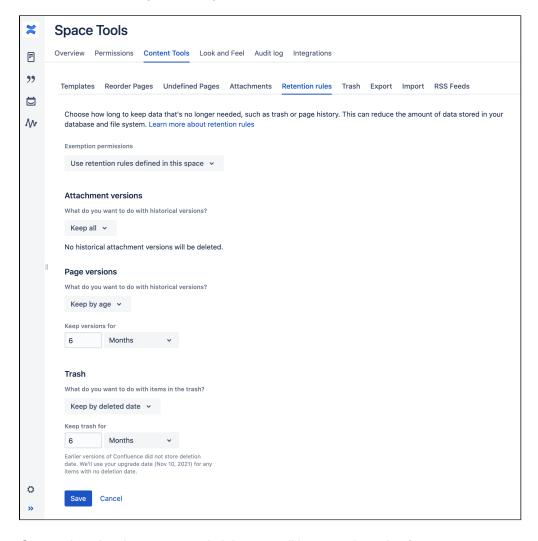
This will add an exemption for this space, or update the existing exemption, if one already exists.

Revert back to global retention rules in a space

To revert back to the global retention rules in a particular space:

- 1. Go to the space and choose Space tools>Content Tools from the bottom of the sidebar
- 2. Select the Retention rules tab.
- 3. Select Edit.
- 4. Select Inherit global retention rules from the Exemption permissions drop down.
- 5. Save your change.

This will remove the space exemption.



Screenshot showing a space administrator editing retention rules for a space

Considerations for administrators

There are a few things you need to consider before changing the retention rules.

Global retention rules apply to all spaces, including archived and personal spaces

The global retention rules apply to all spaces, including archived spaces and personal spaces. If you want to avoid automatically deleting historical versions in archived spaces or personal spaces, you will need to use exemptions either to protect spaces that need to be preserved, or to target spaces that can be cleaned up.

Once you set a global rule, Confluence will start deleting items that don't meet that rule almost immediately, so make sure any exemptions are in place before setting a global rule.

Deleted versions cannot be restored

Deleted versions do not go to the trash, and cannot be restored once deleted. If you need to retain data for regulatory or compliance reasons, you may want to only allow retention rules to be added and updated in the global administration, by system administrators. This is the default.

As always, we recommend you have a robust backup strategy, and a plan for how you will restore data from your backups if required.

Versions can increment very quickly

It's not unusual for many page or file versions to be created within a short space of time.

- A new version of a page is created each someone clicks Publish (or Save if you have collaborative editing disabled).
- A new version of a file is created each time a file with the same name is attached to the page, or a new version uploaded via the **Upload** button in the preview.

You should factor this in when determining the rule criteria to use, so versions aren't deleted too aggressively . Deleting versions by date rather than number may be more appropriate.

Versions are not renumbered and people who contributed to deleted versions will not be listed in page history

The page history screen will show all remaining versions. Page versions that have been deleted, either by a user, or automatically, will not appear. This is intentional, to indicate that any earlier collaboration (including by the original page creator) is no longer available to view or restore.

In this example version 6 was deleted manually by a user (1), and versions 3 and earlier (2) were deleted because they didn't meet the retention rule criteria.

The page creator (shown in the byline) will still be the original page creator, and all previous contributors will be recognised if you search by contributor, but we don't list the people who contributed to versions that were deleted on the page history or attachments versions page.

Data pipeline

(i)

This feature is available with a Confluence Data Center license.

Data pipeline provides an easy way to export data from Jira, Confluence, or Bitbucket, and feed it into your existing data platform (likeTableau orPowerBI). This allows you to:

- generate richer reports and visualizations of site activity
- better understand how your teams are using your application
- make better decisions on optimizing the use of Jira or Confluence in your organization

You can triggera data export in your applications admin console or through the RESTAPI. Data will be exported in CSV format. You can only perform one data export at a time.

For a detailed reference of the exported data's schema, seeData pipeline export schema.

Data pipeline is available inData Centereditions of:

- Jira 8.14 and later
- Confluence 7.12 and later
- Bitbucket 7.13 and later

On this page:

- Requirements
- Considerations
- Access the data pipeline
- Schedule regular exports
- Check the status of an export
- Cancel an export
- Exclude projects from the export
- Configuring the dataexport
- Use the data pipeline REST API
- Output files
- Troubleshooting issues with data exports

Requirements

To trigger data exports through the REST API, youll need:

- A valid Confluence Data Center license
- Systems Administratorglobal permissions

Considerations

There are a number of security and performance impacts youll need to consider before getting started.

Security

The export will include all data, including PII (Personally Identifiable Information) and restricted content. This is to provide you with as much data as possible, so you can filter and transform to generate the insights youre after.

If you need to filter out data based on security and confidentiality, this must be done after the data is exported.

Exported files are saved in your shared home directory, so youll also want to check this is secured appropriately.

Export performance

Exporting data can take a long time in large instances. We intentionally export data at a limited rate to keep any performance impact to your site under a 5% threshold. Its important to note that there is no impact to performance unless an export is in progress.

When scheduling your exports, we recommend that you:

- Limit the amount of data exported using the from Dateparameter, as a date further in the past will export more data, resulting in a longer data export.
- Schedule exports during hours of low activity, or on a node with no activity, if you do observe any performance degradation during the export.

	Number	Approximate export duration
Users	100,000	8 minutes
Spaces	15,000	12 minutes
Pages	25 million	12 hours
Comments	15 million	1 hour
Analytics events	20 million	2 hours

The total export time was around 16 hours.

(i) Test performance VS production

The data presented here is based on our own internal testing. The actual duration and impact of dataexport on your own environment will likely differ depending on your infrastructure, configuration, and load.

Our tests were conducted on a single node Data Center instance in AWS:

- EC2 instance type: c5.4xlarge
- RDS instance type: db.m5.4xlarge

Access the data pipeline

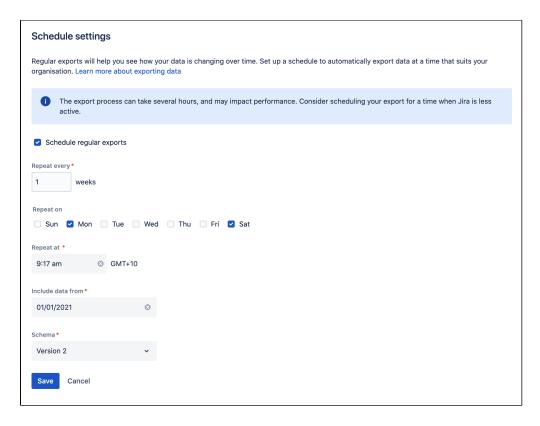
To access the data pipeline go to General Configuration>Data pipeline.

Schedule regular exports

The way to get the most value out of the data pipeline is to schedule regular exports. The data pipeline performs a full export every time, so if you have a large site, you may want to only export once a week.

To set the export schedule:

- 1. From the Data pipeline screen, select **Schedule settings**.
- 2. Select the **Schedule regular exports** checkbox.
- 3. Select the date to include data from. Data from before this date wont be included. This is usually set to 12 months or less.
- 4. Choose how often to repeat the export.
- 5. Select a time to start the export. You may want to schedule the export to happen outside working hours.
- 6. Select the **Schema version** to use (if more than one schema is available).
- 7. Save your schedule.



Timezones and recurring exports

We use your server timezone to schedule exports (or system timezone if youve overridden the server time in the application). The export schedule isnt updated if you change your timezone. If you do need to change the timezone, youll need to edit the schedule and re-enter the export time.

You can schedule exports to happen as often as you need. If you choose to export on multiple days, the first export will occur on the nearest day after you save the schedule. Using the example in the screenshot above, if you set up your schedule on Thursday, the first export would occur on Saturday, and the second export on Monday. We dont wait for the start of the week.

Export schema

The export schema defines the structure of the export. We version the schema so that you know your export will have the same structure as previous exports. This helps you avoid problems if youve built dashboards or reports based on this data.

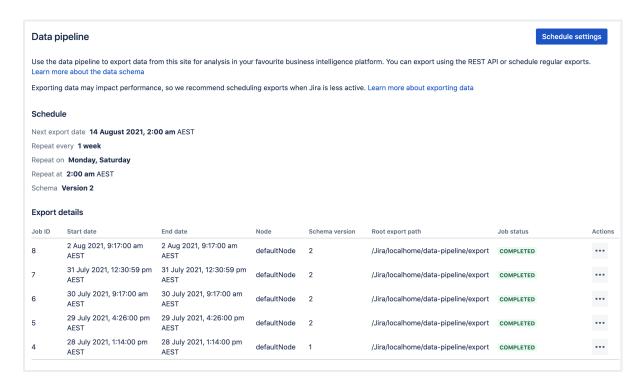
We only introduce new schema versions for breaking changes, such as removing a field, or if the way the data is structured changes. New fields are simply added to the latest schema version.

Older schema versions will be marked as deprecated, and may be removed in future versions. You can still export using these versions, just be aware we wont update them with any new fields.

Check the status of an export

You can check the status of an export and view when your last export ran from the data pipeline screen.

The **Export details** table will show the most recent exports, and the current status.



Select ····>View details to see the full details of the export in JSON format. Details include the export parameters, status, and any errors returned if the export failed.

For help resolving failed or cancelled exports, seeData pipeline troubleshooting.

Cancel an export

To cancel an export while it is in progress:

- Go to the **Data pipeline** screen.
- Select --- next to the export, and choose Cancel export.
- Confirm you want to cancel the export.

It can take a few minutes for the processes to be terminated. Any files already written will remain in the export directory. You can delete these files if you dont need them.

Exclude projects from the export

You can exclude spaces from the export by adding them to an opt-out list. This is useful if you dont need to report on that particular space, or if it contains sensitive content that youd prefer not to export.

To add spaces to the opt-out list, make aPOSTrequest to base-url>/rest/datapipeline/1.0/config
/optout and pass the space keys as follows.

```
{
  "type": "SPACE",
  "keys": ["HR","TEST"]
}
```

These spaces will be excluded from all future exports.

For full details, including how to remove spaces from the opt-out list, refer to the Data pipeline REST API reference.

Automatic data export cancellations

If you shut down a node running a data export, the export will be cancelled. However, if the JVM is not notified after a crash or hardware-level failure, the export process may get locked. This means you'll need to manually mark the export as cancelled (through the UI, or via the REST API by making adeleterequest). This releases the process lock, allowing you to perform another data export.

Configuring the dataexport

You can configure the format of the export data using the following system properties.

Default value	Description		
plugin.data	plugin.data.pipeline.embedded.line.break.preserve		
Specifies whether embedded line breaks should be preserved in the output files. can be problematic for some tools such as Hadoop.			
	This property is set toFalseby default, which means that line breaks are escaped.		
plugin.data	plugin.data.pipeline.embedded.line.break.escape.char		
\\n	Escaping character for embedded line breaks. By default, we'll print\nfor every embedded line break.		
plugin.data	plugin.data.pipeline.minimum.usable.disk.space.after.export		
To prevent you from running out of disk space, the data pipeline will check before an export that there is at least 5GB free disk space.			
	Set this property, in gigabytes, to increase or decrease the limit. To disable this check, set this property to -1(not recommended).		

Use the data pipeline REST API

You can use the data pipeline REST API to export data.

To start a data pipeline export, make a POST request to base-url>/rest/datapipeline/latest/export.

Here is an example request, using cURL and a personal access token for authentication:

```
curl -H "Authorization:Bearer ABCD1234" -H "X-Atlassian-Token: no-check"
-X POST https://myexamplesite.com/rest/datapipeline/latest/
export?fromDate=2020-10-22T01:30:11Z
```

You can also use the API to check the status, change the export location, and schedule or cancel an export.

For full details, refer to the Data pipeline REST API reference.

Output files

Each time you perform a data export, we assign a numerical job ID to the task (starting with1for your first ever data export). This job ID is used in the file name, and location of the files containing your exported data.

Location of exported files

Exported data is saved as separate CSV files. The files are saved to the following directory:

- <shared-home>/data-pipeline/export/<job-id> if you run Confluence in a cluster
- <local-home>/data-pipeline/export/<job-id> you are using non-clustered Confluence

Within the <job-id> directory you will see the following files:

```
• users_job<job_id>_<schema_version>_<timestamp>.csv
```

- spaces_job<job_id>_<schema_version>_<timestamp>.csvpages_job<job_id>_<schema_version>_<timestamp>.csv
- comments_job<job_id>_<schema_version>_<timestamp>.csv
- analytics_events_job<job_id>_<schema_version>_<timestamp>.csv

To load and transform the data in these files, you'll need to understand the schema. SeeData pipeline export schema.

Set a custom export path

By default, the data pipeline exports the files to the home directory, but you can use the REST API to set a custom export path.

To change the root export path, make a PUT request to <base-url>/rest/datapipeline/1.0/config/export-path.

In the body of the request pass the absolute path to your preferred directory.

For full details, including how to revert back to the default path, refer to the Data pipeline REST API reference.

Sample Spark and Hadoop import configurations

If you have an existing Spark or Hadoop instance, use the following references to configure how to import your data for further transformation.

```
%python
# File location
file_location = "/FileStore/**/export_2020_09_24T03_32_18Z.csv"
# Automatically set data type for columns
infer_schema = "true"
# Skip first row as it's a header
first_row_is_header = "true"
# Ignore multiline within double quotes
multiline_support = "true"
# The applied options are for CSV files. For other file types, these will be ignored. Note escape &
quote options for RFC-4801 compliant files
df = spark.read.format("csv") \
  .option("inferSchema", infer_schema) \
  .option("header", first_row_is_header)
  .option("multiLine", multiline_support) \
  .option("quote", "\"") \
  .option("escape", "\"") \
  .option("encoding", "UTF-8").load(file_location)
display(df)
```

```
CREATE EXTERNAL TABLE IF NOT EXISTS some_db.datapipeline_export (
  `page_id` string,
  `instance_url` string,
  `space_key` string,
  `page_url` string,
  `page_type` string,
  `page_title` string,
  `page_status` string,
  `page_content` string,
  `page_parent_id` string,
  `labels` string,
  `page_version` string,
  `creator_id` string,
  `last_modifier_id` string,
  `created_date` string, `updated_date` string,
  `last_update_description` string
ROW FORMAT SERDE 'org.apache.hadoop.hive.serde2.OpenCSVSerde'
WITH SERDEPROPERTIES (
  "escapeChar" = "\\",
  'quoteChar' = '"',
  'separatorChar' = ','
) LOCATION 's3://my-data-pipeline-bucket/test-exports/'
TBLPROPERTIES ('has_encrypted_data'='false');
```

Troubleshooting issues with data exports

Exports can fail for a number of reasons, for example if your search index isnt up to date. For guidance on common failures, and how to resolve them, seeData pipeline troubleshooting in our knowledge base.

Data pipeline export schema

This page describes the structure and data schema of the Confluence data export files.

To learn more about how the set up and configure your data pipeline, seeDa ta pipeline.

Output file format and structure

The output files are written in CSV format and are RFC4180 compliant. They have the following characteristics:

- Each file has a header. This includes files from exports that resulted in no data.
- New lines are separated by CRLF characters \r\n.
- Fields containing line breaks (CRLF), double quotes, and commas are enclosed in double quote.
- If double-quotes are present inside fields, then a double-quote appearing inside a field are escaped by preceding it with another double quote. For example: "aaa", "b" "bb", "ccc".
- Fields with no data (null values) are represented in the CSV export by two consecutive delimiters(as in, , ,).
- Embedded break lines are escaped by default and printed asn.

Fields are available in all schema versions, unless specifically noted below.

Users file

Field	Description		
instance_url	Type: URL		
	Description : Base URL of the current instance.		
	Example: https://yoursitename.com		
user_id	Type: String		
	Description: ID of the user		
	Example : ff8080817572401e01757240b3520000		
user_name	Type: String		
	Description : User name of the user.		
	Example: jsmith		
user_fullname	Type: String		
	Description: Full name of the user.		
	Example: John Smith		
user_email	Type: Email		
	Description: Email address of the user		
	Example: jsmith@example.com		

On this page:

- Output file format and structure
- Users file
- Spaces file
- Pages file
- Comments file
- Analytics events file

Spaces file

Field	Description
space_key	Type: String
	Description : Unique identifier that forms part of the URL for thatspace
	Example: AMF
instance_url	Type: String
	Description : Base URL of the current instance
	Example: https://example.com
space_url	Type: URL
	Description : The space URL
	Example: https://example.com/display/SPACEKEY
homepage_url	Type: URL
	Description : The space's home page URL
	Example: https://example/display/SPACEKEY/Page+name
space_name	Type: String
	Description: Title of the space
	Example: Design Team Space
space_type	Type: String
	Description : Whether the space is a global or personal space
	Example: global
space_status	Type: String
	Description : Whether the status of the space is current or archived
	Example: CURRENT
creator_id	Type: User
	Description : ID of the user who created the space
	Example: ff8080817572401e01757240b3520000
last_modifier_id	Type: User
	Description : ID of the user who last modified the space
	Example: ff8080817572401e01757240b3520000
created_date	Type: Time
	Description: Space creation timestamp
	Example: 2021-02-26T04:14:38Z

updated_date	Type: Time
	Description: Last modification timestamp
	Example: 2021-02-26T04:14:38Z

Pages file

Field	Description
page_id	Type: Number
	Description: Unique ID of the page
instance_url	Type: String
	Description: Base URL of the current instance
	Example: https://example.com
space_key	Type: String
	Description: Space key of the space the page exists in
page_url	Type: String
	Description: URL of the page
	Example: https:/example/display/SPACEKEY/Page+name
page_type	Type: String
	Description : Whether the entity is a page or a blog post
	Example: page
page_title	Type: String
	Description: Title of the page
page_status	Type: String
	Description : Status of the page (the only value is current, this does not indicate that a page is in a space that has been archived)
page_content	Type: String
	Description : Content of the page in Confluence storage format (limited to 10,000 characters)
	Example:
	<pre><ac:layout><ac:layout-section ac:type="two_equal"><ac:layout-cell> This is sample content in a layout</ac:layout-cell><ac:layout-cell> With two columns</ac:layout-cell></ac:layout-section></ac:layout></pre>
page_parent_id	Type: Number
	Description: ID of the current page's direct parent

labels Type: String			
	Description: Comma separated list of labels of the page		
	Example: ["personal", "expense]		
page_version	rsion Type: String		
	Description : Version number of the latest version page		
	Example: 3		
creator_id	Type: Number		
	Description : ID of the user who created the page		
	Example: ff8080817572401e01757240b3520000		
last_modifier_id	Type: User		
	Description: ID of the user who last updated the page		
	Example: ff8080817572401e01757240b3520000		
created_date	Type: Time		
	Description: Creation timestamp		
	Example: 2021-02-26T04:14:38Z		
updated_date	Type: Time		
	Description: Last modification timestamp		
	Example: 2021-02-26T04:14:38Z		
last_update_de scription	Type: String		
σοπραστ	Description : Version comment entered when the page was last updated (limited to 2,000 characters)		

Comments file

Field	Description
comment_id	Type: Number
	Description: Unique ID of the comment
instance_url	Type: String
	Description : Base URL of the current instance
	Example: https://example.com
comment_url	Type: String
	Description: Full URL of the comment

page_id	Type: Number
	Description : Unique ID of the page which contains the comment
parent_comme nt_id	Type: Number
_	Description : If the comment is a reply, this is the ID of the parent comment (empty for top level comments)
comment_conte	Type: String
	Description : Content of the comment in Confluence storage format (limited to 2,000 characters)
	Example:
	Sample comment on a page
creator_id	Type: User
	Description: ID of the user who created the comment
	Example: ff8080817572401e01757240b3520000
last_modifier_id	Type: User
	Description: ID of the user who last modified the comment
	Example: ff8080817572401e01757240b3520000
created_date	Type: Time
	Description : Creation timestamp
	Example: 2021-02-26T04:14:38Z
updated_date	Type: Time
	Description: Last modification timestamp
	Example: 2021-02-26T04:14:38Z

Analytics events file

Field	Description
instanc	Type: String
e_url	Description: Base URL of the current instance
	Example: https://example.com
event_id	Type: Number
	Description: Unique ID of the analytics event

event_ name	Type: String Description: Name of the analytics event. Events include page_viewed, page_created, page_updated, blog_viewed, blog_created, blog_updated, comment_created, attachment_viewed, attachment_created.
created _date	Type: Time Description: Creation timestamp Example: 2021-02-26T04:14:38Z
event_ author_ id	Type: User Description: ID of the user who performed the action that triggered the event Example: ff8080817572401e01757240b3520000
event_ space_ key	Type: String Description: Space key of the space the event was triggered in or affects (affected object) Example: SPACEKEY
event_ contain er_id	Type: Number Description: ID of the containing entity. For pages this is the page ID, for attachments and comments, its the page ID of the page the attachment or comment appears on.
event_ content _id	Type: Number Description: ID of the entity. For pages this is the page ID, for attachments, it is the attachment ID, and for comments it's the comment ID.

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
- Scheduled Jobs
- Configuring the Allowlist
- Configuring the Time Interval at which Drafts are Saved

Related pages:

- Customizing your Confluence Site
- Confluence administrator's guide

Viewing System Information

The System Information screen provides information about Confluence's configuration, which plugins are in use, and the environment in which Confluence has been deployed.

To view your system information go to General Configuration System Information.

Notes:

- The handy **memory graph** helps you keep track of Confluence's memory usage.
- Your system configuration information is helpful to Atlassian Support when diagnosing errors you may face using Confluence. When logging a support request or bug report, please provide as much detail as possible about your installation and environment.

Related pages:

- Cache Statistics
- Live Monitoring Using the JMX Interface
- Tracking Customizations Made to your Confluence Installation

Live Monitoring Using the JMX Interface

JMX (Java Management ExtensionsAPI) allows you to monitor the status of your Confluence instance in real time. JMX uses objects called MBeans (Managed Beans) to expose data and resources from your application, providing useful data such as the resource usage of your instance and its database latency, allowing you to diagnose problems or performance issues.

In this page we'll guide you through how to useJConsoleto monitor Confluence locally and remotely. JConsole is included in the Java Development Kit (JDK), but you can use any JMX client.



⚠ This guide provides a basic introduction to the JMX interface and is provided as is. Our support team can help you troubleshoot a specific Confluence problem, but aren't able to help you set up your monitoring system or interpret the results.

Write JMX metrics to a log file

You can also choose to write the following JMX metrics to a log file. This is useful when you are troubleshooting a problem.

- CacheStatistics
- IndexingStatistics
- MailTaskQueue
- RequestMetrics
- SystemInformation
- ThreadPool
- OS
- GC
- Threading
- TomcatManager
- RequestProcessor

To write JMX metrics to a log file:

- 1. Go to O-> General Configuration > Scheduled jobs
- 2. Enable the Log JMX Metrics job.

This job runs once per minute by default, and writes metrics to the <local-home/logs/atlassianconfluence-jmx.logfile.

Monitor Confluence locally using JConsole

If you are troubleshooting a particular issue, or only need to monitor Confluence for a short time, you can use local monitoring. Local monitoring can have a performance impact on your server, so its not recommended for long term monitoring of your production system.

To monitor locally:

- 1. Start JConsole (you'll find it in thebindirectory of the JDK installation directory)
- 2. SelectLocal Process.
- 3. Select the Confluence process. It will be called something likeorg.apache.catalina.startup. Bootstrap start

SeeUsing JConsolefor more information on local monitoring.

Monitor Confluence remotely using JConsole

Remote monitoring is recommended for production systems, as it does not consume resources on your Confluence server.

To monitor remotely:

1. Add the following properties to yoursetenv.sh/setenv.batfile. The port can be any port that is not in use.

```
set CATALINA_OPTS=-Dcom.sun.management.jmxremote %CATALINA_OPTS%
set CATALINA_OPTS=-Dcom.sun.management.jmxremote.port=8099 %CATALINA_OPTS%
```

- Decide how you will secure your remote connection. SeeRemote Monitoring and Management for more information.
 - Although it is possible to disable authentication, we do not recommend doing this on a production system.
- 3. Start JConsole (you'll find it in thebindirectory of the JDK installation directory).
- 4. SelectRemote Process.
- 5. Enter your hostname and port (this is the port you specified earlier, not the Confluence port).
- 6. ClickConnect.

SeeUsing JConsolefor more information on remote monitoring.

Confluence MBeans

You can use the following Confluence MBeans to see live information about your Confluence instance.

CacheStatistics

This MBean shows information about Confluence caches. This info can also be found on the Cache Statisticspag e.

IndexingStatistics

This MBean shows information related to search indexing. Here's some useful attributes.

Property name	Function	Values
Flushing	Indicate whether the cache is currently flushing	True/False
LastElapsedMilliseconds	Time taken during last indexing	Milliseconds
TaskQueueLength	Shows number of tasks in the queue	Integer
ReIndexing	Indicates whether Confluence is currently reindexing	True/False

SystemInformation

This MBean shows information such as the Confluence version and uptime. This info can also be found on the Sy stem Information page.

Property name	Function	Values
DatabaseExampleLaten cy	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
AverageExecutionTimeForLastTenRequ ests	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	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.

Property name	Function	Values
AllJobNames	Shows information on current scheduled jobs including the time they were last run	String
CurrentlyRunningJobNa mes	Lists the scheduled jobs that are currently running	List

Additional MBeans

To also monitor Hibernate and Hazelcast (Confluence Data Center only) you will need to add the following properties to yoursetenv.sh/setenv.batfile first.

```
set CATALINA_OPTS=-Dconfluence.hazelcast.jmx.enable=true %CATALINA_OPTS%
set CATALINA_OPTS=-Dconfluence.hibernate.jmx.enable=true %CATALINA_OPTS%
```

This will make the Hibernate and Hazelcast MBeans available in your JMX client.

Monitoring high CPU consuming threads

The Top Threads Pluginfor JConsole is useful for monitoring whether the CPU is spiking. Use the following command to start JConsole with this plugin:

```
JConsole -pluginpath /pathto/topthreads.jar
```

Tracking Customizations 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 customizations after the upgrade.

To see the modifications made to files in your Confluence installation:

- 1. Choose the cog icon , then choose General Configuration
- 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	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. jaror other JAR files. If you modify classes, the Confluence modification detection does not report the modification.

View Space Activity

The space and site statistics features described on this page are dis abled by default. We don't recommend this feature for large sites as it can significantly slow down your site.

If you have Confluence Data Center, check outAnalyticswhich is designed for enterprise scale, and can be used in big, busy sites.

Space activity information is disabled by default, and the 'Activity' tab won't be visible unless the Confluence Usage Stats system app is enabled. See notes below.

If enabled, the space activity screen displays statistics on the activity in each space. These include:

- How many pages and blog posts have been viewed, added or updated over a given period.
- Which content is the most popular (most frequently viewed).
- Which content is the most active (most frequently edited).
- Which people are the most active contributors/editors of content.

To view the activity in a space:

- 1. Go to the space and choose **Space Tools** at the bottom of the sidebar
- 2. Choose Activity

You'll see a graphic display of the number of pages and blog posts that have been viewed, added, and edited, showing trends over a period of time.

Screenshot: The Space Activity tab

Related pages:

- Page History and Page Comparison Views
- · Watch Pages, Spaces and Blogs
- How do I get more statistics from Confluence?

Space Activity					
Overview	Permissions	Content Tools	Look and Feel	Integrations	Activity
Activit	у				
Activity for week starting 10 February 2013 Period: month week day Week: << Previous Now Next >> See also: Global Confluence activity					

In addition to the graphical representation of Views and Edits, the top ten most popular and most active pages and/or blog posts will be listed, with a link to each.

Screenshot: Popular content, active content, and active contributors.

Most popular content (Views)

- 6. 🖹 Ideas for Angry Nerds 2 revenge 5. 🖺 Status update (1) of the ... (2)
- 7. Status update (1)
- 8. Games as a communication medium (1)

Most active content (Edits)

- 1. Angry Nerds Home (6)
 2. Which nerd are you? (3)
 3. Step 1 download the game (2)
 4. Why I loved Angry Nerds (2)
 5. Step 2 Play (2)
 1. Josh User (5)
 2. Why I loved Angry Nerds (1)
 3. Which nerd are you? (1)
 4. Step 2 Play (2)
 4. Step 2 Play (1)

 - 6. Games as a communication medium (1)

Most active contributors (Edits)

Enable Confluence Usage Stats

To enable the Confluence Usage Stats system app:

- 1. Go to > Manage apps.
- 2. Select System from the dropdown then search for Confluence Usage Stats.
- 3. Expand the listing and choose **Enable**.

Notes

- To view Space Activity the Confluence Usage Stats system plugin must be enabled. This plugin is known to cause performance problems on large installations and is disabled by default. System administrators can enable this system app.
- The app only collects data when it's activated.
- If you're using Confluence Data Center or Confluence Cloud, space activity information isn't available.
- Page hits aren't unique the graph on the Space Activity screen includes all page hits, including multiple visits by the same user.

Viewing Site Statistics

The space and site statistics features described on this page are **disabled by default**. We don't recommend this feature for large sites as it can significantly slow down your site.

If you have Confluence Data Center, check outAnalyticswhich is designed for enterprise scale, and can be used in big, busy sites.

If enabled, the global activity screen displays statistics on the activity in your Confluence site. These include:

- How many pages and blog posts have been viewed, added or updated over a given period.
- Which spaces are the most popular (most frequently viewed).
- Which spaces are the most active (most frequently edited).
- Which people are the most active contributors /editors of content.

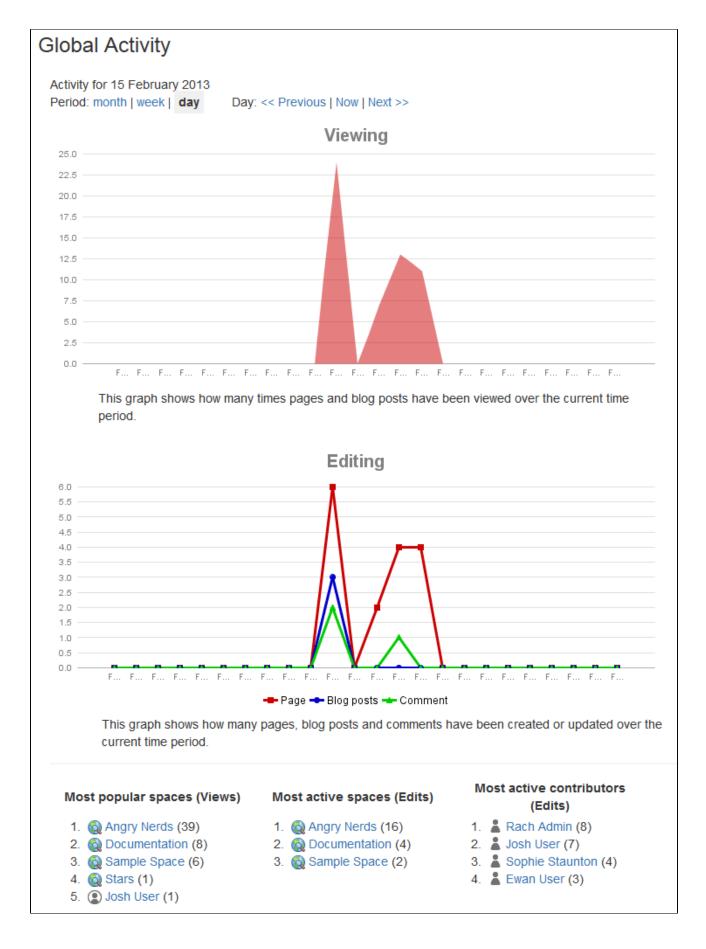
Related pages:

- How do I get more statistics from Confluence?
- Cache Statistics
- View Space Activity
- Live Monitoring Using the JMX Interface

To view the activity on your site:

- 1. Choose the cog icon , then choose General Configuration
- 2. Choose 'Global Activity' in the 'Administration' section of the left-hand panel (only appears if enabled seebelow).

Screenshot: Global Activity



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 Improve activity plugin performance NEW
- 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 recognized by your Confluence installation:

- 1. Choose the cog icon , then choose General Configuration
- 2. Choose **System Information** in the left-hand panel.
- 3. Scroll down to the section titled System Properties.

Configuring the Server Base URL

The **Server Base URL** is the URL via which users access Confluence. The base URL**must** be set to the same URL by which browsers will be viewing your Confluence site.

Confluence will automatically detect the base URL during setup, but you may need to set it manually if your site's URL changes or if you set up Confluence from a different URL to the one that will be used to access it publicly.

You need to have System Administrator permissions in order to perform this function.

To change the Server Base URL:

- 1. Go to Configuration.
- SelectEdit.
- 3. Enter the new URL in the Server Base URL field.
- 4. Save your changes.

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 also need to make these changes:
 - 1. Stop Confluence.
 - 2. Go to the Confluence installation directory and edit <installation-directory>\conf\server.xml.
 - 3. 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 yourpath attribute should look like this:

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

In this example we've used/confluenceas the context path. Note that you can't use/resources as your context path, as this is used by Confluence, and will cause problems later on.

- 4. Save the file.
- 5. Go to the Confluence home directory and edit<confluence-home>/confluence.cfg.xml
- 6. Changeconfluence.webapp.context.pathto reflect the new context path. For example:

confluence.webapp.context.path">/confluence/property>

7. Restart Confluence and check you can access it at http://www.foobar.com/confluence.

You may also want toclear the Confluence plugins cachebefore restarting.

- 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. For more information on proxing Atlassian applications, seeProx ying Atlassian Server applications.
- This information needs to be added in the Connector element at {CONFLUENCE_INSTALLATION} \conf\server.xml.

Configuring the Confluence Search and Index

Confluence administrators can adjust the behavior of the Confluence search, and manage the index used by the search.

- Configuring Indexing Language
- Configuring Search
- Content Index Administration
- Enabling OpenSearch
- Rebuilding the Ancestor Table
- Setting Up Confluence to Index External Sites
- Setting Up an External Search Tool to Index Confluence

Related pages:

- Search
- Confluence Administrator's Guide

Configuring Indexing Language

Changing the indexing language to be used in your Confluence site may improve the accuracy of Confluence search results, if the majority of the content in your site is in a language other than English.

Confluence supports content indexing in:

- Arabic
- Brazilian
- Chinese
- CJK
- Custom Japanese
- Czech
- Danish
- Dutch
- English (default)
- Finish
- French
- German
- Greek
- Hungarian
- Italian
- Norwegian
- Persian
- Romanian
- Russian
- Spanish
- Swedish

To configure the indexing language:

- 1. Go to Configuration then choose Edit.
- 2. Select the **Indexing Language** from the dropdown list in the **Formatt ing and International Settings** section.
- 3. Choose Save.

Related pages:

- Choosing a Default Language
- Installing a Language Pack
- Content Index Administration
- How to Rebuild the Content Indexes From Scratch on Confluence Server

Configuring Search

There are a few ways to search for content in Confluence:

- Using the search panel, which allows you to quickly search and filter results.
- Using the advanced search page.
- Using a search macro embedded on a Confluence page (for example, the Livesearch Macro or QuickNav Gadget).

Read more about the differentsearch options in Confluence.

By default, the search panel feature is enabled, with the maximum number of simultaneous requests set to 40. These options can be modified as described below.

Related pages:

Search

Set the number of simultaneous search requests

Confluence admins can set the maximum number of simultaneous searches users can perform using the search panel. By default, the maximum is set to 40. This limit applies to a single Confluence node. If you're running Confluence Data Center with multiple nodes, this number will increase.

If your Confluence server serves a large number of individuals who use this feature regularly, some of whom are being denied access to it, you may wish to increase this value.

To change the maximum number of simultaneous search requests in Confluence:

- 1. Choose the cog icon , then choose General Configuration
- 2. Choose Further Configuration in the left-hand panel.
- 3. Choose Edit.
- 4. Enter the appropriate number in the field beside Max Simultaneous Requests.
- 5. Choose **Save**.

Disable quick search options

The search panel feature offers a quick way for users to search and filter content in Confluence. We recommend keeping this feature enabled, unless it's causing significant performance issues on your site.

If you disable the quick search option:

- The search panel will no longer appear when users click the search field. When you enter a search query, we'll take you to the advanced search page.
- The Confluence QuickNav Gadget will no longer drop down a list of search results. When you enter a search query, we'll take you to the advanced search page.

Other search macros, including the Livesearch Macro and the Page Tree Search Macro, won't be affected if you disable the quick search option.

To disable quick search options from your Confluence site:

- 1. Choose the cog icon , then choose General Configuration
- 2. ChooseFurther Configuration in the left-hand panel.
- 3. ChooseEdit.
- 4. Deselect the Quick Search checkbox.
- 5. ChooseSave.

Content Index Administration

The search index is used by search, the dashboard, some macros, and all the other places where we show information about content in your Confluence site. The search index is made up of:

- a content index which contains content such as the text of pages, blog posts, and comments
- a change index which contains data about each change, such as when a page was last edited.

On this page:

- View the index queues
- Rebuild the search index
- Location of search indexes
- Index recovery in a cluster
- Check the size of your index
- Troubleshooting

Related pages:

- Scheduled Jobs
- Search
- Configuring the Confluence Search and Index

These indexes are updated automatically as people get work done on your site. Changes, such as a new page, comment, or edit to an existing page, aren't updated in each index immediately. They're placed into queues and regularly processed in batches (as often as every 5 seconds).

View the index queues

It can take a while for the queues to process if there are thousands of changes to your site within a short period.

To check the contents of the queue:

- 1. Go to O-Seneral Configuration-Content Indexing.
- 2. Select the Content queue or Change queuetab.

Here you can see the number of items in the queue, the last time the queue was processed, and how long it took to process. This information is useful for troubleshooting if your users report issues with search, dashboard activity feeds.

Rebuild the search index

There are some circumstances where you may need to reindex your site, for example, if your users report issues with search, dashboard activity feeds, or if directed to as part of an upgrade.

To rebuild the search index:

- 1. Go to O-> General Configuration > Content Indexing.
- 2. Select **Rebuild** and follow the prompts to confirm you want to rebuild the index.

This will rebuild content index and the change index, and can take some time for very large sites.

Screenshot: The Content Indexing screen showing information about the last time the index was rebuilt.

29/07/2020

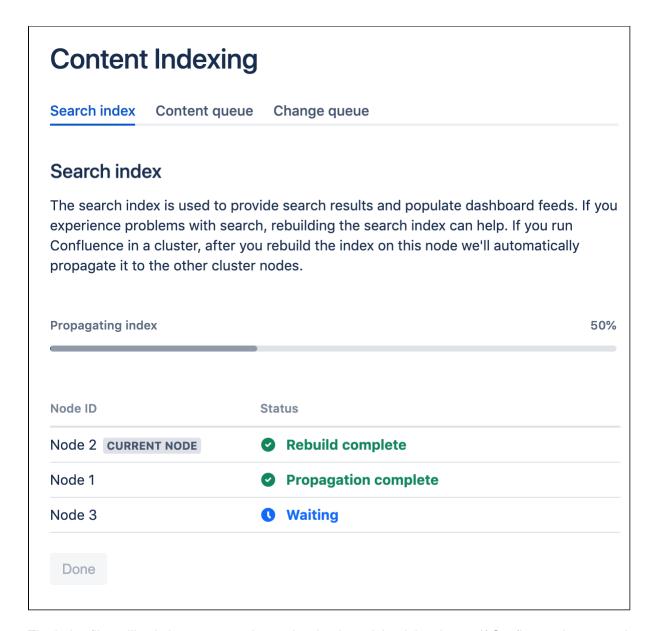
Content Indexing				
Search index	Content queue	Change queue		
Search index The search index is used to provide search results and populate dashboard feeds. If you experience problems with search, rebuilding the search index can help.				
Rebuild				
Last rebuilt		Duration		

Propagate the search index to your cluster

If you're running Confluence in a cluster, once the index is rebuilt on the current node, we'll automatically propagate the index files to all the other nodes in the cluster.

43 minutes

Screenshot: The Content Indexing screen showing progress as the index is propagated to each node in the cluster.



The index files will only be propagated to nodes that have joined the cluster. If Confluence is not running on a node, we won't be able to propagate the index to that node.

If there's a problem, you'll see an error with information about what went wrong, for example, if a node becomes unavailable, or there's insufficient disk space to copy the index.

Impact on end users

Confluence will continue to use the existing index until the new index has been rebuilt successfully. Your users can continue to search and use Confluence, but may experience some performance degradation. This is because rebuilding the index increases the load on your server significantly.

Rebuilding the index can take several hours. The amount of time depends on the number, type, and size of pages and attachments in your site, the amount of memory allocated, and disk throughput.

If you have a large site, there are some things you can do to reduce the impact on your users:

- If you're running Confluence on a single node, kick off the rebuild on a weekend, or during a scheduled maintenance window.
- If you're running Confluence in a cluster, remove the node rebuilding the index from your load balancer. Once propagation is complete, you can add the node back into the pool.

Disk space requirements

You'll need enough free disk space in your home directory (or local home directories if you're running Confluence in a cluster) to temporarily accommodate two full copies of your index. This is because the existing index is not removed until the new index is ready.

If you're running Confluence in a cluster, you'll also need enough free space in your shared home directory to accommodatean additional index snapshot.

Location of search indexes

You can find the Confluence index in the <home-directory>/indexdirectory.

If you're running Confluence in a cluster, a full copy of the Confluence indexes are stored in the <local-home>/indexdirectory on each Confluence node. A journal service keeps each index in sync.

If you need to see the contents of the search index for any reason, there is a tool you can use to browse the index directly. SeeHow to view the contents of the search index in Confluence Server and Data Center.

Index recovery in a cluster

If you run Confluence in a cluster, a snapshot of the search index is also stored in the shared home directory. These snapshots are created by the Clean Journal Entries scheduled job which, by default, runs once per day.

When you start a Confluence node it will check whether its index is current, and if not, it will request a recovery snapshot from the shared home directory. If a snapshot is not available, it will generate a snapshot from a running node (with a matching build number). Once the recovery snapshot is extracted into the index directory, Confluence will continue the startup process. The journal service will then make any further updates required to bring the index up to date.

If the snapshot can't be generated, or is not received in time, existing index files will be removed and Confluence will perform a reindex on that node. If your index is very large or your file system slow, you may need to increase the time Confluence waits for the snapshot to be generated using theconfluence. cluster.index.recovery.generation.timeoutsystem property.

Index recovery only happens on node startup, so if you suspect a problem with a particular cluster node's index, restart that node to trigger index recovery.

The index recovery snapshot isn't used when you manually rebuild your index from the UI. The rebuild process generates a brand new snapshot, before propagating it to other nodes in the cluster.

Check the size of your index

You can measure the index size in two ways, size on disk, or you can use the number of pages and blogs as a rough indication of the amount of content in the index.

To check the size on disk:

- 1. Go to <local-home>/index
- Check the size of that directory. The way you do this will depend on your operating system.

To check the number of pages and blogs in the index:

- 1. Go to Co Seneral Configuration System information
- 2. Scroll down to the Confluence usage section and check the Content (Current Versions) value.

Troubleshooting

If you have problems rebuilding the search index, the following may help.

Can't rebuild the index

If you're unable to rebuild the index from the Confluence UI, or if you still have problems with search after rebuilding the index, you may need to rebuild the index from scratch. The way you do this depends on whether Confluence is running in a cluster:

- How to Rebuild the Content Indexes From Scratch on Confluence Server
- How to Rebuild the Content Indexes From Scratch

Can't access content Indexing page

If the content indexing page does not load properly, and you see a "We can't check the status of your index, you may have lost your connection, refresh the page to try again" error, try updating your browser to the latest version.

Poor performance while rebuilding the index

If you experience stability problems while the index is being rebuilt, you canreduce the number of threads Confluence should use to rebuild the index. Set the reindex. thread.count system property to define the maximum number of threads that can be used.

If bothreindex.thread.countandindex.queue.thread.countare unset, we default to 4 threads.

Out of memory errors while rebuilding the index

If you experience out of memory errors while rebuilding the index, increasing the heap memory available to Confluence may help. See How to fix out of memory errors by increasing available memory.

Rebuild index failed to propagate to other nodes in the cluster

This generally happens when there is not enough free disk space for the local home directory on each node to accommodate two copies of the index. SeeFailed to propagate index in Confluence Data Center 7.7 and laterto find out how to re-try the propagation.

Enabling OpenSearch

With OpenSearch autodiscovery, you can add Confluence search to your Firefox or Internet Explorer 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
- 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.

Information about OpenSearch

- Confluence supports the autodiscovery part of the OpenSearch standard, by supplying an OpenSearch description document. This is an XML file that describes the web interface provided by Confluence's search function.
- Any client applications that support OpenSearch will be able to add Confluence to their list of search engines.

Related pages:

- Search
- Confluence Administrator's Guide

Rebuilding the Ancestor Table



The way you fix problems with the ancestor table changed significantly in Confluence 6.14.

If you're using Confluence 6.13 or earlier, seeRebuilding the Ancestor Table in Confluence 6.13.5 or earlierto find out how to rebuild the ancestor table.

The ancestortable records the parent and descendant (child) relationship between pages. It is also used when determining whether a page will inherit view restrictions from a parent page.

Occasionally records in the ancestor table can become corrupted. The Repair the Ancestors Table scheduled job finds and automatically fixes problems in the ancestors table in all current spaces. The job runs daily.

The job ignores archived spaces, so if you suspect there is a problem with the ancestor table in a particular archived space, you will need to change the status of the space to 'current', then run the job manually.

Repair the ancestor table manually

If you suspect there is a problem, you can also run this job manually.

- 1. Go to General Configuration > Scheduled Jobs.
- 2. Locate the Repair the Ancestors Table job and choose Run.

The job should complete quickly, and there won't be any impact on users.

Viewing the job results

If you want to see the result of the job each time it runs, you can change the logging level of com.atlassian. confluence.pages.ancestorsto INFO.

You'll then see a message similar to the one below in the Confluence application log, each time the job runs:

```
Ancestors have been repaired. Found and fixed 3 broken pages.
It took 71 sec for 6407 spaces, average space processing time 0 sec.
```

Rebuilding the ancestor table in earlier Confluence versions

If you're using Confluence 6.13 or earlier the way you rebuild the ancestor table is different. SeeRebuilding the Ancestor Table in Confluence 6.13.5 or earlierfor more information.

Setting Up Confluence to Index External Sites

Confluence cannot easily index external sites, due to the way Lucenesearch works in Confluence, butth ere are two alternatives:

- 1. Embed External Pages Into Confluence
- 2. Replace Confluence Search

Related pages:

- Setting Up an External Search Tool to Index Confluence
- Configuring the Confluence Search and Index

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 Macroa nd use it embed the external content inside normal Confluence pages.

The actual content of the external site won't be indexed.

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 Confluence 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. 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 APIs.

Related pages:

- Setting Up Confluence to Index External Sites
- Configuring the Confluence Search and Index

Configuring Mail

- Configuring a Server for Outgoing Mail
 Setting Up a Mail Session for the Confluence Distribution
 Configuring the Recommended Updates Email Notification
 The Mail Queue
- Customizing 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 personalize 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. Go to General Configuration Mail Servers. This will list all currently configured SMTP servers.
- 2. Click Add New SMTP Server (or edit an existing server).
- 3. 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. This must be an ordinary email address, you can'tenter variables in this field
 - From Name: Enter the name that will be displayed in the 'from' field for email messages originating from this server. This is the text which appears before the user's registered email address (in square brackets).

This field accepts the following variables, which reference specific details defined in the relevant Confluence user's profile:

Variable	Description
\${fullname}	The user's full name.
\${email}	The user's email address.
\${email.hostname}	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.
- 4. Enter your Hostname, Port, User name and Password details.

If your SMTP host uses the Transport Layer Security (TLS) protocol select **Use TLS**.

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, seeSetting 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 serveras described above.

- Click Send Test Email to check that the server is working. Check that you get the test email in your inhox
- You can flush the email queue to send the email message immediately. Go toMail Queue, and click Fl ush Mail Queue. See The Mail Queue.

A user can test that notifications are working as follows:

- Go to your user profile (using the **Settings** link) and edit your email preferences. See Email Notifications.
- 2. Enable **Notify On My Actions**. (By default, Confluence does not send you notifications for your own changes.)
- 3. Go to a page you wish to get notifications about.
- 4. Choose Watch at the top-right of the page. See Watch Pages, Spaces and Blogs.
- 5. Edit the page, make a change, and save the page.
- 6. Check your email inbox. You may need to wait a while for the email message to arrive.

Setting Up a Mail Session for the Confluence Distribution

The simplest way to set up a mail server is through the Confluence Administration console. SeeConfiguring a Server for Outgoing Mail.

If youwant to add different options or parameters you can also set up a mail session for the Confluence distribution. In the example below we'll set up Gmail.

To set up a mail session for the Confluence distribution:

- 1. Stop Confluence.
- Move (don't copy) the following files from <confluence-install>\confluence\WEB-INF\lib to <confluence-install>\lib

```
jakarta.mail-x.x.x.jar
javax.activation-x.x.x.jar
javax.activation-api-1.2.0.jar
```

(x.x.x. represents the version numbers on the jar files in your installation)

Don't leave a renamed backup of the jar files in\confluence\WEB-INF\lib. Even with a different file name, the files will still be loaded as long as it remains in the directory.

3. Edit the <confluence-install>\conf\server.xml file and add the following at the end of the Confluence <context> tag, just before </Context>.

Note: you're editing the <context> tag that contains the *Confluence* context path, not the one that contains the *Synchrony* context path.

```
<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. Go to Some Servers Servers.
- 6. Select either Edit an existing configuration, or Add a new SMTP mail server.
- 7. 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.

8. Save your changes and send a test email.

Configuring the Recommended Updates Email Notification

Confluence sends a regular email report to subscribers, containing the top content that is relevant to the person receiving the message, from spaces they have permission to view. This is called the 'Recommended Updates' notification.

If you have Confluence Administrator or System Administrator permissions, you can configure the default settings that determine how often the Recommended Updates notification is sent. When new users are added to Confluence, the default settings will be applied to their user profiles.

Confluence users can choose their personal settings, which will override the defaults. See Email Notifications.

Initial settings of the defaults

When you install Confluence, the initial values of the default settings are as follows:

- The default frequency is weekly.
- If your Confluence site has public signup enabled, the Recommended Updates notification is disabled by default. If public signup is not enabled, the notification is enabled by default.

You can change the above settings, specifying a different default value for the site.

Notes:

- The Recommended Updates notification is sent only to people who have a user profile in Confluence. If your Confluence site uses external user management, such as LDAP, then people will receive the report only after they have logged in for the first time. (The first login creates their user profile.)
- The daily email message is sent at 1 p.m. in the user's configured time zone.
- The weekly email message is sent at 1 p.m. on Thursdays in the user's configured time zone.

Configuring the Recommended Updates notification

You can set thedefault 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 Q, then choose General Configuration
- 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' system app. SeeDisabling and enabling apps.

On this page:

- Initial settings of the defaults
- Configuring the Recommended Updates notification
- Disabling the Recommended Updates notification for the entire site

Related pages:

Email Notifications

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
- 2. Choose Mail Queue in the left-hand panel. This will display the email messages currently in the queue.
- 3. Choose Flush Mail Queue to send all email messages immediately.
- 4. Choose **Error Queue** to view failed email messages. You can try to **Resend** the messages, which will flush the mails back to the mail queue, or you can **Delete** them from here.

Related pages:

- Configuring a Server for Outgoing Mail
- Setting Up a Mail Session for the Confluence Distribution

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

Configuring Character Encoding

Confluence and your database must be configured to use the same character encoding. To avoid problems with character encodingalways set all character encodings to UTF-8 (or the equivalent for your database, for example, AL32UTF8 for Oracle databases).

On this page:

- Configuring Confluence character encoding
- Database character encoding
- Problems with character encodings

Related pages:

Configuring Database Character Encoding

Configuring Confluence character encoding

By default, Confluence uses UTF-8 character encoding. Confluence has a number of checks in place to make sure your database is also using UTF-8 (or equivalent for your database).

While it is possible to change the character encoding, it is **not recommended**. Changing the Confluence character encoding will change your HTTP request and response encoding and your filesystem encoding as used by exports and Velocity templates. You may also be prevented from restarting or upgrading Confluence, depending on your database.

To change the Confluence character encoding (not recommended):

- 1. Go to Configuration and choose Edit
- 2. Enter the new character encoding of your choice in the text box next to **Encoding** then **Save**.

Database character encoding

Your database, and the JDBC or datasource connection to it, must be configured to use UTF-8(or the equivalent for your database, for example, AL32UTF8 for Oracle databases). There are a number of checks in place to warn you if your database character encoding is incorrect.

SeeConfiguring Database Character Encodingfor more information.

Problems with character encodings

SeeTroubleshooting Character Encodingsto find out how to test your character encoding.

Troubleshooting Character Encodings

If character encoding is not configured correctly in your Confluence site, you may experience problems like:

- 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

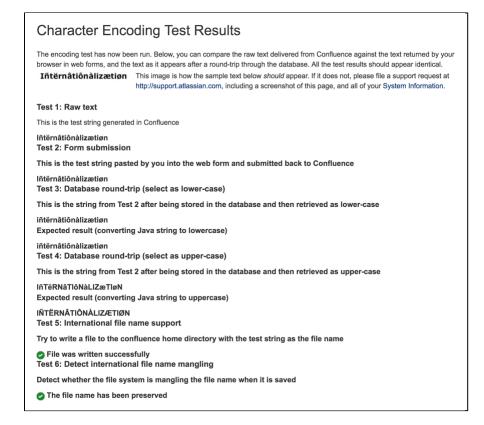
To diagnose the problem, follow these steps.

1. Run the encoding test

Confluence includes an encoding test that can reveal problems with your configuration. You'll need to be a Confluence admin to do this.

- 1. Head to<your-confluence-url>/admin/encodingtest.action
- 2. Follow the prompts to paste a line of text and start the test. You can also paste text in a specific language, for example Japanese, if you're experiencing a particular problem with that language.

If the text displayed in the encoding test is different to what you entered, then there are problems with your character encoding settings. Here's what a successful test looks like.



2. Use the same encoding for your database

Your database and Confluence must use the same character encoding. SeeConfiguring Database Character Encodingfor more information.

3. Get help

If you're still having problems with character encoding, create a support request, and our support team will help you solve the problem.

Include the following details to help us identify your problem:

screenshots of the problem occurring

- results of the encoding test
 information about your database (including version)
 A copy of the information on yourSystem Information page.

"" 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 behavior.

I write a page with a Euro sign in it (). All is well, the Euro sign shows up in the wiki markup textbox, 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 recognize the symbol.

There are two potential reasons for this behavior:

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 notaffect already existing data.

Database is not using utf-8. Confluence and your connection are.

If your Database encoding is not set to UTF-8, yet is using some other encoding such as *latin1*, it could be one of the potential reasons why you lose the "" characters at some stage. It could be occurring due to **caching**. When Confluence saves data to the database, it may also keep a local cached copy. If the database encoding is set incorrectly, the Euro character may not be correctly recorded in the database, but Confluence will continue to use its cached copy of that data (which is encoded correctly). The encoding error will only be noticed when the cache expires, and the incorrectly encoded data is fetched from the database.

For instance the *latin1* encoding would store and display all 2-byte UTF8 characters correctly except for the euro character which is replaced by '?' before being stored. As Confluence's encoding was set to UTF-8, the 2-byte UTF-8 characters were stored in *latin1* database assuming that they were two *latin1* different characters, instead of one utf8 character. Nevertheless, this is not the case for 3-byte utf8 characters, such as the Euro symbol.

Please ensure that you set the character encoding to UTF-8 for all the entities of your system as advised in this guide.

MySQL 3.x Character Encoding Problems

MySQL 3.x is known to have some problems upper- and lower-casing certain (non-ASCII) characters.

Diagnosing the problem

- 1. Follow the instructions for Troubleshooting Character Encodings.
- 2. If the upper- and lower-cased strings displayed on the Encoding Test are different, then your database is probably affected.

An example (faulty) output of the Encoding Test is shown below:

Screenshot: Encoding Test Output (excerpt)

Test 4: Database round-trip (select as upper-case)

This is the string from Test 2 after being stored in the database and then retrieved as upper-case

IñTËRNATIÔNÀLIZæTIØN <



IÑTËRNÂTIÔNÀLIZÆTIØN

Solution

Upgrade to a newer version of MySQL. (4.1 is confirmed to work.)

Other Settings

- Configuring a WebDAV client for Confluence
- Configuring a WebDAV client for Cor
 Configuring HTTP Timeout Settings
 Configuring Number Formats
 Configuring Shortcut Links
 Configuring Time and Date Formats
 Enabling the Remote API
 Enabling Threaded Comments
 Installing a Language Pack
 Installing Patched Class Files

Configuring a WebDAV client for Confluence

WebDAV allows users to access Confluence content via a WebDAV client, such as 'My Network Places' in Microsoft Windows. Provided that the user has permission, they will be able to read and write to spaces, pages and attachments in Confluence. Users will be asked to log in and the standard Confluence content access permissions will apply to the equivalent content available through the WebDAV client.

Mapping a Confluence WebDAV network drive requires a set of specific criteria to be met. For specific information, please see Wi ndows Network Drive Requirements.

Introduction to Confluence's WebDAV Client Integration

By default, all WebDAV clients have permission to write to Confluence. Write permissions include the ability for a WebDAV client to create, edit, move or delete content associated with spaces, pages and attachments in a Confluence installation.

On this page:

- Introduction to Confluence's WebDAV Client Integration
- Using a WebDAV Client to Work with
- Restricting WebDAV Client Write Access to Confluence
- Disabling Strict Path Checking
- Virtual Files and Folders

Related pages:

- Global Permissions Overview
- Disabling and enabling apps
- Attachment Storage Configuration

On the 'WebDAV Configuration' screen in the Confluence Administration Console, you can:

- Deny a WebDAV client write permissions to a Confluence installationusing a regular expression (regex)
- Disable or enable strict path checking
- Enable or disable access to specific virtual files/folders

Note:

- The 'WebDav Configuration' page is only available if the WebDAV plugin has been enabled. This plugin is bundled with Confluence, and can be enabled or disabled by the System Administrator.
- The settings on the 'WebDay Configuration' pagedonotapplytoexternal attachment storage configuration.

Using a WebDAV Client to Work with Pages

The following sections tell you how to set up a WebDAV client natively for a range of different operating systems. WebDAV clients typically appear as drives in your operating system's file browser application, such as Windows Explorer in Microsoft Windows, or Kongueror in Linux.

Accessing Confluence in Finder on Mac OSX



You can successfully connect but you can't see content when using HTTPS, so this technique won't work for Confluence Cloud. Use a third-party WebDAV client instead.

To use Finder to view and manage Confluence content:

- 1. In Finder chooseGo>Connect to Server
- 2. Enter your Confluence URL in this format:

http://<confluence base URL>/plugins/servlet/confluence/default

For example if your Confluence URL ishttp://ourconfluence.sample.com/wikiyou would enter:

http://ourconfluence.sample.com/wiki/plugins/servlet/confluence/default

3. Enter your Confluence username and password and click Connect



Use your username (jsmith), not your email address, unless your email address is your username.

Confluence will appear as a shared drive in Finder. You can use the same URL to connect using a third party WebDav client, like CyberDuck.

Accessing Confluence in Explorer in Microsoft Windows

This section covers the two methods for configuring a WebDAV client natively in Microsoft Windows:

- As a network drive
- As a web folder

If possible, use the network drive method as this will enable more comprehensive WebDAV client interaction with Confluence than that provided by a web folder. However, your Confluence instance must meet several environmental constraints if you use this method. If you cannot configure your instance to meet these requirements, then use the web folder method or third-party WebDAV client software.

If you're using SSL you may need to add @SSLto the end of your server URL, for example:

http://<confluence server url>@SSL/confluence/plugins/servlet/confluence/default

If you run into any problems with the procedures in this section, please refer to the WebDAV Troubleshooting page.

Windows Network Drive

To map a Confluence WebDAV client network drive, your Confluence instance must be configured so that *all*o f the following criteria is met:

- Has no context root
- There's an issuethat 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'sserver discovery issue.

To map a Confluence WebDAV client network drive in Microsoft Windows:

- In Windows go to Map Network Drive.
 SeeMap a network drive in the Windows documentation to find out how to get to this in your version of Windows.
- 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. ClickFinish

When prompted for login credentials, specify your Confluence username and password.

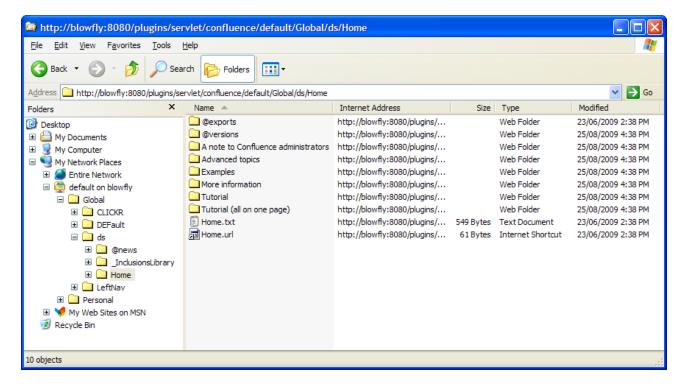
Windows Web Folder

To map a Confluence WebDAV client web folder:

1. Go to My Network Places and choose Add a network place and follow the prompts.

- 2. 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/defaultOrhttp://<confluence server url>/plugins/servlet/confluence/default) and clickNext
- 3. Enter your Confluence username and password
- 4. Provide a meaningful name for your web folder and proceed with the wizard
- 5. ClickFinish

Screenshot: A Confluence WebDAV Client Web Folder in Windows XP



Setting up a WebDAV client in Linux

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 managerKonqueror, which is part of the LinuxK 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/defaultOrwebdavs: //<confluence server url>/plugins/servlet /confluence/default) and pressEnter.
- 3. Enter your Confluence username and password if prompted

You should be able to click to load many, but not all files. In practice, you would normally save a modified file locally, then drag it to the Konqueror window to upload it to Confluence.

Restricting WebDAV Client Write Access to Confluence

In earlier versions of the WebDAV plugin, separate options for restricting a WebDAV client's write permissions (that is, create/move, edit and delete actions), were available. However, in the current version of this plugin, they have been simplified and combined into a general write permission restriction that covers all of these actions.

WebDAV clients are now denied write permission to your Confluence installation by setting a regex that matches specific content within the WebDAV client's user agent header. Upon setting a regex, it will be added to a list of restricted WebDAV clients. Any WebDAV clients whose user agent header matches a regex in this list will be denied write permission to your Confluence installation.

Example: A PROPFIND method header generated by a Microsoft Web Folder WebDAV client, showing the user agent header field:

```
PROPFIND /plugins/servlet/confluence/default HTTP/1.1
Content-Language: en-us
Accept-Language: en-us
Content-Type: text/xml
Translate: f
Depth: 1
Content-Length: 489
User-Agent: Microsoft Data Access Internet Publishing Provider DAV
Host: 127.0.0.1:8082
Connection: Keep-Alive
```



Unlike earlier versions of the WebDAV plugin, which could only restrict write permissions forallWebD AV clients, the current version of this plugin allows you to restrict write permissions to specific WebDAV clients.

To restrict a WebDAV client's write access permissions to your Confluence installation:

- 1. Choose thecog icon , then choose General Configuration
- 2. Choose 'WebDav Configuration' in the left panel
- 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 Repeat steps 3 and 4 to add a regex for each additional WebDAV client you want to restrict.
- 5. HitSave

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
- 2. Click**WebDav Configuration**under 'Configuration' in the left panel
- 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
- 5. HitSave

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 :	
Save	

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 minimize these problems.

To disable the WebDAV plugin's strict path checking option:

- 1. Choose the cog icon . then choose General Configuration
- 2. Click**WebDav Configuration**under 'Configuration' in the left panel
- 3. Clear the 'Disable strict path check' check box
- 4. HitSave

Virtual Files and Folders

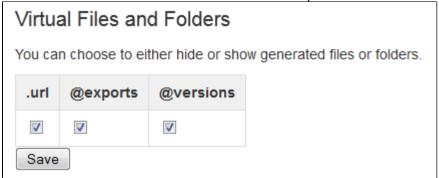
In the unlikely event that you have problems with the WebDAV client's performance or stability, you can enable access to automatically generated (that is, virtual) files and folders.

Note:

By default, these options are hidden on the 'WebDAV Configuration' page. To make them visible, append the parameter?hiddenOptionsEnabled=trueto 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



To enable or disable access to virtual files and folders:

- Choose thecog icon , then chooseGeneral Configuration
 ClickWebDav Configurationunder 'Configuration' in the left panel
- 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. HitSave

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
- 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:###############

Confluence uses the guidelines in this Java document from Oracle: Class NumberFormat.

To change the number formats in Confluence:

- 2 Choose Edit
- 3. Update the Long Number Format and Decimal Number Format to suit your requirements
- 4. Choose **Save**

Configuring Shortcut Links

Shortcut links provide a quick way of linking to resources that are frequently referenced from Confluence. When you create a shortcut link, you assign a key to an URL so that, when editing, a user can type just the key instead of the complete URL.

Example: Creating a shortcut to Google

Most Google searches look like this: http://www.google.com/search?q=. If you create a shortcut for this search with the key 'google', every time a user needs to use http://www.google.com/search?q=searchterms, they can just type [searchterms@google] instead.

On this page:

- Creating shortcut links
- Using shortcut links
- Deleting shortcut links

Here is a screenshot showing the shortcuts currently defined onhttp://confluence.atlassian.com:

Key	Expanded Value	Default Alias	Operations
cache	http://www.google.com/search?q=cache:		Remove
imdb	http://us.imdb.com/Title?		<u>Remove</u>
jira	http://jira.atlassian.com/secure/QuickSearch.jspa?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=		<u>Remove</u>

Shortcut links are added and maintained by Confluence administrators from the Administration Console.

Creating shortcut links

To create a shortcut link:

- 1. Choose the cog icon O, then choose General Configuration
- 2. Choose **Shortcut Links** in the left-hand panel.
- 3. Enter a Key for your shortcut. This is the shortcut name a user will use to reference the URL.
- 4. Enter the **Expanded Value**. This is the URL for the link. You can use '%s' in the URL to specify where the user's input is inserted. If there is no '%s' in the URL, the user's input will be put at the end.
- 5. Enter a **Default Alias**. This is the text of the link which will be displayed on the page where the shortcut is used, with the user's text being substituted for '%s'.
- 6. Choose Submit.

Using shortcut links

Enter a shortcut link on the **Advanced** tab of the Insert Link dialog. See Links for details.

Specify in the link what should be appended to the end of the shortcut URL, followed by an at-sign (@) and the key of the shortcut. Shortcut names are case-insensitive. So, for example, using the keys shown in the above screenshot:

To link to	Type this	Resulting URL	Demonstration
a 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 change how times and dates appear throughout your Confluence site to suit your organization's preferred date format.

On this page:

- Site date and time
- Relative dates
- Date lozenges

Related pages:

- Choosing a Default Language
- Installing a Language Pack

Site date and time

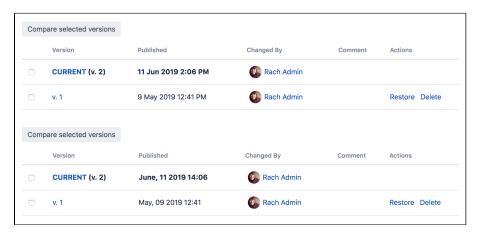
To change the time and date formats for your entire site:

- 1. Go to General Configuration.
- 2. Choose Edit.
- 3. Enter your preferredTime Format, Date Time Format and Date Format.
- 4. Choose Save.

Confluence uses the Java SimpleDateFormat class. Head to Java SimpleDateFormatto see all allowed values, or use one of the common format examples below.

Format	Example
dd MMMM yyyy	05 June 2019
MMMM d, yyyy	June 5, 2019
d MMM yy	5 Jun 19
dd-MM-yy	05-06-19
d-M-yy	5-6-19
h:mm a	3:25 PM
HH:mm	15:25

Here's what the page history for the same page might look like with different date and time formats.



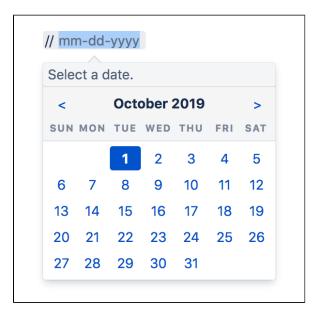
Relative dates

Some parts of Confluence use relative dates when the change happened recently. For example a comment might have been added "yesterday", or a page modified "about 2 hours ago".

It's not possible to customize this format. Full dates are displayed in your preferred format once the change is more than 1 day old.

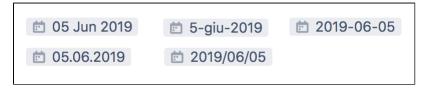
Date lozenges

To insert a date lozenge, in the editor type I/O r choose + > **Date** from the toolbar.



When you insert a date lozenge, the date format displayed depends on the language settings for the current user. This means the lozenge will appear differently for different users.

The example below shows English, Italian, Polish, German, and Japanese date formats. The same format is used for UK and US English.



Enabling the Remote API



(i) XML-RPC and SOAP remote APIs were deprecated in Confluence 5.5. We recommend using the fully supported Confluence Server REST API wherever possible.

To use the XML-RPC and SOAP remote APIs younced to enable the APIs from the **Administration Console**. You'll need System Administrator permissions to do this.

To enable the remote API:

- Choose the cog icon , then choose General Configuration
 Click Further Configuration in the left-hand panel.

- 4. Click the check box next to Remote API (XML-RPC & SOAP).
- 5. Click Save.

Enabling Threaded Comments

Comments on pages or blog posts are displayed in one of two views:

- **Threaded**: Shows the comments in a hierarchy of responses. Each reply to a comment is indented to indicate the relationships between the comments.
- **Flat**: Displays all the comments in one single list and does not indicate the relationships between comments.

By default, comments are displayed in **threaded** mode. A Confluence Administrator (see Global Permissions Overview) can enable or disable the threaded view for the entire Confluence site.

To enable or disable the threaded view:

- 1. Choose the cog icon , then choose General Configuration
- 2. Select Further Configuration in the left-hand panel
- 3. Choose**Edit**
- Select or unselect the Threaded Comments checkbox to enable or disable threaded mode
- 5. ChooseSave

Related pages:

- Comment on pages and blog posts
- Confluence administrator's guide

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.

Confluence is available in these languages right out of the box:

- etina (esk republika | Czech Republic)
- Dansk (Danmark | Denmark)
- Deutsch (Deutschland | Germany)
- Eesti (Eesti | Estonia)
- English (UK)
- English (US)
- Espaol (Espaa | Spain)
- Franais (France)
- slenska (sland | Iceland)
- Italiano (Italia | Italy)
- Magyar (Magyarorszg | Hungary)
- Nederlands (Nederland | The Netherlands)
- Norsk (Norge | Norway)
- Polski (Polska | Poland)
- Portugus (Brasil | Brazil)
- Romn (Romnia | Romania)
- Slovenina (Slovensk republika | Slovak Republic)
- Suomi (Suomi | Finland)
- Svenska (Sverige | Sweden)
- (| Russia)
- (| China)
- (| Japan)
- (| Republic of Korea)

You can make additional languages available by installing language pack apps. You'll need to be a Confluence administrator to install a language pack.

Installing additional language packs

We no longer provide community translations for Confluence. You can find a small number of third-party language packs on the Atlassian Marketplace.

Showing User Interface Key Names for Translation

This feature is useful if you are troubleshooting translations of the Confluence user interface. After opening the Confluence dashboard, you can add the following action to the end of your Confluence URL:

?i18ntranslate=on

For example http://myconfluencesite.com?i18ntranslate=on

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.

The key names are displayed with a 'lightning bolt' graphic. Here's an example from a space sidebar:

Related pages:

- Choosing a Default Language
- Configuring Indexing Language
- Installing Marketplace apps



To turn off the translation view, add the following to the end of the Confluence URL:

?i18ntranslate=off

Installing Patched Class Files

Atlassian support or the Atlassian bug-fixing team may occasionally provide patches for critical issues that have been resolved but have not yet made it into a release. Those patches will be class files which are attached to the relevant issue in our Jira bug-tracking system.

Installation Instructions for the Confluence Distribution

Follow these steps to install a patched class file:

- 1. Shut down your confluence instance.
- 2. Copy the supplied class files to <installation-directory>/confluence/WEB-INF/classes /<subdirectories>, where:
 - <installation-directory> must be replaced with your Confluence Installation directory. (If
 you need more information, read about the Confluence Installation Directory.)
 - <subdirectories> must be replaced by the value specified in the relevant Jira issue. This value will be different for different issues. In some cases, the subdirectories will not exist and you will need to create them before copying the class files. Some issues will contain the patch in the form of a ZIP file which will contain the desired directory structure.
- 3. Restart your Confluence instance for the changes to become effective.
- ① 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.

① Once the issue that the patch relates to is resolved, you should upgrade to the version of Confluence that contains the fix, and revert the patch. Patches are often naive and untested and may not solve the problem in the most efficient way. As such, an official fix should be preferred in all cases.

Configuring System Properties

This page describes how to set Java properties and options on startup for Confluence.



See How to fix out of memory errors by increasing available memory for specific instructions for OutOfMemory Errors.

On this page:

- Linux
- Windows (starting from .bat file)
- Windows service
- Confluence Data Center deployed in AWS
- Verifying your settings
- Recognized system properties

Related pages:

- Recognized System Properties
- How to fix out of memory errors by increasing available memory

Linux

To configure System Properties in Linux installations:

- 1. Edit the <installation-directory > /bin/setenv.shfile.
- Find the sectionCATALINA_OPTS=
 (this is JAVA_OPTS= in Confluence 5.5 and earlier)
- 3. Refer to the list of parameters in Recognized System Properties.
- Add all parameters in a space-separated list, inside the quotations. Make sure to keep the string \${CATALINA_OPTS}" in place.

Windows (starting from .bat file)

To Configure System Properties in Windows Installations When Starting from the .bat File:

- 1. Edit the <installation-directory > /bin/setenv.batfile.
- 2. Find the sectionset CATALINA_OPTS=%CATALINA_OPTS% (this is JAVA_OPTS=%JAVA_OPTS%in Confluence 5.5 and earlier)
- 3. Refer to the list of parameters in Recognized System Properties.
- Add all parameters in a space-separated list. Make sure to keep the string %CATALINA OPTS% in place.

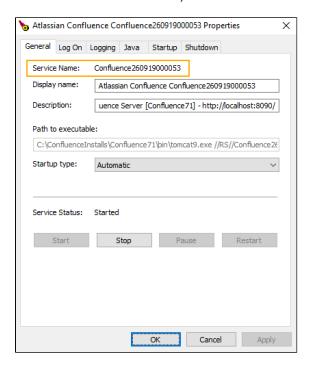
Windows service

There are two ways to configure system properties when youStart Confluence Automatically on Windows as a Service, either viacommand lineorin the Windows Registry

Setting properties for Windows services via command line

To set properties for Windows services via a command line:

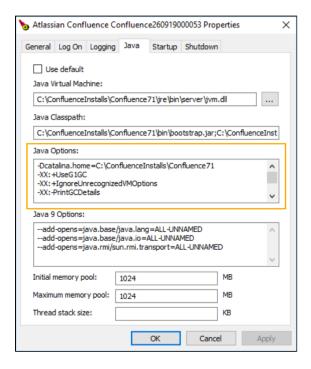
1. Identify the name of the service that Confluence is installed as in Windows (Go to**Control Panel>Adm** inistrative Tools>Services):



- In the above example, the service name is Confluence 260919000053.
- 2. Open the command window (Choose**Start>cmd.exe**)
- 3. cdto thebindirectory of your Confluence instance andrun the following command:

tomcat9w //ES//<SERVICENAME>

- In the above example, it would betomcat9w //ES//Confluence260919000053
 The Tomcat version number may be different if you are using an earlier version of Confluence.
- 4. Click on the Javatab to see the list of current start-up options:
- 5. Append any new option on its own new line by adding to the end of the existing Java Options. Refer to the list of parametersinRecognized System Properties.



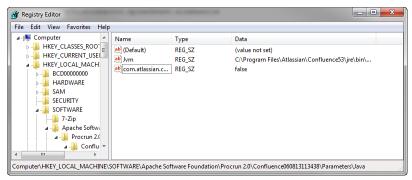
Setting properties for Windows services via the Windows registry

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:

64bit: HKEY_LOCAL_MACHINE >> SOFTWARE >> WOW6432Node >> Apache Software Foundation >> Procrun 2.0 >> Confluence service name

32bit: HKEY_LOCAL_MACHINE >> SOFTWARE >> 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 parametersinRecognized System Properties. Enter each on a separate line.

Confluence Data Center deployed in AWS

If you've used the Quick Start or CloudFormation template to deploy Confluence Data Center in AWS, you will pass system properties via the Cloud Formation Template, and not using the methods described above.

- 1. In the AWS Console, choose Update Stack
- 2. Under Advanced, enter system properties in the Catalina Properties field as follows:

```
-Xms1024m -Xmx1024m -Dsystemproperty=value
```

3. Changes are applied when a new nodes are provisioned.

Verifying your settings

To see what Confluence is using, check Viewing System Properties.

Recognized system properties

SeeRecognized System Properties for the full list of system properties available to your Confluence version.

Recognized 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 – Dflag to the Java virtual machine in which Confluence is running. See thefull instructions: Configuring System Properties.

Since	Default Value	Effect
atlassian	.forceSchem	aUpdate
1.0	false	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.
confluenc	ce.home	
1.0	Any filesyste m path	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.
confluenc	ce.dev.mode	
1.0	false	Enables additional debugging options that may be of use to Confluence developers (additionally it changes spring bean creation to use lazyinitializationby default to decrease startup time). Do not enable this flag on a production system.
confluenc	ce.disable.ma	ailpolling
2.4	false	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
confluenc	ce.i18n.reloa	dbundles
1.0	true	Setting this property will cause Confluence to reload its i18n resource bundles every time an internationalized string is looked up. This can be useful when testing translations, but will make Confluence run <i>insanely slowly</i> .
confluenc	ce.ignore.dek	bug.logging
1.0	true	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.
confluenc	ce.jmx.disabl	ed
3.0	false	If set to "true", will disable Confluence's JMX monitoring. This has the same effect as setting the "enabled" property to false in WEB-INF/classes/jmxContext.xml
confluenc	ce.optimize.ir	ndex.modulo
2.2	20	Number of index queue flushes before the index is optimized.
		This property was removed in Confluence 5.2 when optimize index was removed.
confluenc	ce.plugins.bu	indled.disable
2.9	false	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.
atlassian	indexing.cor	ntentbody.maxsize

3.0	1048576	When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename. If the amount of content extracted from the file exceeds the limit set by this property (def ault is 1MB, in bytes), the file's contents will still be indexed and searchable, but will not appear when the file is returned in search results. Increasing this limit may make
		displaying search results slower. See Configuring Attachment Size for more info.
atlassia	n.mail.fetchdis	abled
3.5	false	Disables mail fetching services for IMAP and POP
atlassia	n.mail.senddis	rabled
3.5	false	Disables sending of mail
atlassia	n.disable.cach	ies
2.4	true	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 devlopment if you are frequently changing static resources and don't want to continually flush your browser cache.
confluer	nce.html.enco	de.automatic
2.9		Setting this property forces the antixss encoding on or off, overriding the behavior dictated by settings. The default behavior differs between Confluence versions.
org.osgi	i.framework.bo	potdelegation
2.10	empty	Comma-separated list of package names to provide from application for OSGi plugins. Typically required when profiling Confluence. For example: "com.jprofiler.,com.yourkit."
confluer	nce.diff.pool.si	ze
3.1	20	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).
confluer	nce.diff.timeou	t
3.1	1000	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.
confluer	nce.html.diff.tir	neout
4.0	10000	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.
atlassia	n.user.experin	nentalMapping
2.10	false	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 CC NF-12319, fixed in Confluence 3.1.1.
confluer	nce.import.use	e-experimental-importer
3.2	false	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.

3.3	false	Disables automatic minification of JavaScript and CSS resources served by Confluence.
index.qu	eue.thread.c	ount
3.3	See "Effect"	Sets the number of threads to be used for the reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. There is no default value, i.e. • If you don't set index.queue.thread.count, the number of threads to be used are calculated based on the number of objects that need to be reindexed and the number of processors available (a maximum of 50 threads will be used). • If you set index.queue.thread.count=2, 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 index.queue.thread.count=200, then ten threads (the maximum allowed) will be used to reindex the content. Note: For Confluence versions from 3.3 to 5.6 the maximum threads is 10.
index.qu	eue.batch.siz	ze
3.3	1500	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.
password	d.confirmatio	n.disabled
3.4	false	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.
confluen	ce.browser.la	anguage.enabled
3.5	true	Setting this property to "false" disables the detection of browser language headers, effectively restoring Confluence behavior 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	
Univer sal Plugin Manag er 1.5	false	When this property is set to true, then UPM will not try to access the The Atlassian Marketplace. This is useful for application servers that do not have access to the Internet. See the UPM documentation.
confluen	ce.reindex.do	ocuments.to.pop
3.5.9	20	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.
confluen	ce.reindex.at	ttachments.to .pop
	10	Indicates how many attachments each indexing thread should process at a time during

3.5.11	false	Forces Confluence to treat any LDAP directories it migrates as Active Directory, rather than relying on looking for sAMAccountName in the username attribute. This is necessary if you are upgrading from before Confluence 3.5, and need to use an attribute other than sAMAccountName to identify your users and are seeing LDAP: error code 4 - Sizelimit Exceeded exceptions in your logs. For more details, see Unable to Log In with Confluence 3.5 or Later Due to 'LDAP error code 4 - Sizelimit Exceeded'
confluer	nce.context.ba	tching.disable
4.0	false	Disables batching for web resources in contexts (e.g. editor, main, admin). Useful for diagnosing the source of javascript or CSS errors.
com.atla	assian.logout.d	disable.session.invalidation
4.0	false	Disables the session invalidation on log out. As of 4.0 the default behavior 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.
officeco	nnector.sprea	dsheet.xlsxmaxsize
4.0.5	2097152	Indicates the maximum size in bytes of an Excel file that can be viewed using theviewx 1s macro. If empty, the maximum size defaults to 2Mb. See CONF-21043 for more details.
com.atla	assian.conflue	nce.extra.calendar3.display.events.calendar.maxpercalendar
	200	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.atla	assian.conflue	nce.allow.downgrade
4.3.2	false	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.
confluer	nce.skip.reinde	ex
	false	Generally a full reindex is not required when upgrading Confluence, but there may be some occasions where an upgrade task will kick-off the reindex process. Set this property to true, to skip rebuilding the search index when Confluence is upgraded. This can be useful if you have a very large site and wish to delay rebuilding the index until a time after the upgrade is complete.
reindex.	thread.count	
5.2		Sets the number of threads to be used for a one-off reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. This system property does not affect the incremental indexing that Confluence does. From Confluence 7.14 the default value of this property is calculated based on either
		the number of available CPUs or free memory, whichever is lower.
reindex.	attachments.t	hread.count
5.2	4	Sets the number of concurrent threads to be used when reindexing attachments, and allows you to reduce the concurrency for these more memory intensive index items. From Confluence 7.14 the default value of this property is calculated based on either the number of available CPUs or free memory, whichever is lower. Minimum 4 threads.

		xport.word.max.embedded.images
5.2	50	This property limits the number of images that are included when you export a Confluence page to Word. When you export a page with many large images images to Word, all the images are loaded into memory, which can then cause out of memory errors affecting your entire Confluence site. You can temporarily increase this limit, using this system property, if you need to export a page with lots of images.
confluer	nce.mbox.direc	etory
5.4.1		Setting this property defines the directory on your Confluence Server where mailboxes can be imported from (for use with the Confluence Mail Archiving system app). 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.
confluer	nce.search.ma	x.results
5.5	1000	Setting this property changes the maximum number of items Confluence Search will return.
confluer	nce.upgrade.re	covery.file.enabled
5.5	true	By default, Confluence creates an upgrade recovery file before and after an upgrade. The operation can take long time on large databases and can be safely turned off if there is a process to back up database and verify the backup before performing an upgrade. Setting this property to false will disable upgrade recovery file creation.
confluer	nce.junit.report	directory
5.5		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.
officeco	nnector.textext	tract.word.docxmaxsize
5.5.3	16777216	When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename. Confluence will only extract content from a Word document up to the limit set by this property (defaults to 16MB, in bytes), before proceeding to index it. This will mean that only part of file's contents are searchable. The check uses the uncompressed file size, not the compressed size on disk in the case of .docx files. See Configuring Attachment Size for more info.
cluster.l	ogin.remembe	rme.enabled
5.6	False	In a cluster, setting this property to True will enable the 'Remember Me' checkbox on the login page. This is not recommended in a cluster, and is disabled by default (i.e. 'Remember me' is always on and users can move seamlessly between nodes).
		This system property has no effect in standalone Confluence.
confluer	nce.cluster.haz	elcast.listenPort
5.6	5801	In a cluster, this property can be used to override the default port that Hazelcast will bind to, for example, if the port is unavailable, or you need to run more than one node on the same host (not recommended). It defaults to 5801.

5.7		The number of threads allocated to the file conversion service is calculated dynamically based on the amount of memory assigned to the instance and the number of CPU cores (usually 4 to 6 threads). This property can be used to change the number of threads. Decrease threads to resolve OOME issues, increase threads to resolve problems with documents spending too long in the queue.
confluer	nce.documen	t.conversion.threads.wait
5.7	1000	Set this property to change the maximum number of items that can be queued for conversion. Any file conversion requests that are made when this maximum limit has been reached are aborted.
confluer	nce.cluster.nc	ode.name
5.7		Set this property to give each node in your Data Center cluster a human readable name (displayed in email notifications and in the footer). If left unset, Confluence will assign a node identifier to each node.
confluer	nce.documen	t.conversion.fontpath
5.8.7		Set this property to define a directory where you can add additional fonts to be used when rendering files (in previews and thumbnails).
		This is useful if you need to support previewing files with specific fonts, or fonts with multibtye characters (such as Japanese).
confluer	nce.documen	t.conversion.words.defaultfontname
5.8.7		Set this property to change the default font for rendering Word (.doc and .docx) files in Confluence.
		Specify just the name of the font (not the path) .
confluer	nce.documen	t.conversion.slides.defaultfontname.regular
5.8.7		Set this property to change the default font for rendering regular fonts in Powerpoint (. ppt and . $pptx$) files in Confluence.
		Specify just the name of the font (not the path) .
confluer	nce.documen	t.conversion.slides.defaultfontname.asian
5.8.7	TakaoP Gothic	Set this property to change the default font for rendering asian fonts in Powerpoint (. ppt and .pptx) files in Confluence.
		Specify just the name of the font (not the path) .
confluer	nce.documen	t.conversion.slides.defaultfontname.symbol
5.8.7		Set this property to change the default font for rendering symbols in Powerpoint ($.ppt$ a nd $.pptx$) files in Confluence.
		This is the font that will be used for bullets and other symbols when the font Symbol is not found.
		Specify just the name of the font (not the path) .
confluer	nce.clickjacki	ng.protection.disable
5.8.15	false	Security features prevent Confluence from being embedded in an <iframe> . To disable this protection, set this property to true which will allow Confluence to be embedded in an <iframe> .</iframe></iframe>
	aco clustor in	dex.recovery.query.timeout

5.9.1	10	In Confluence Data Center, t he amount of time, in seconds, that a confluence node needing index recovery will wait for another node to offer an index snapshot, before it gives up and stops attempting to recover the index.
confluen	ce.cluster.ind	dex.recovery.generation.timeout
5.9.1	120	In Confluence Data Center, the amount of time, in seconds, that a confluence node needing index recovery will wait for an index snapshot to be created by another node, before it gives up and and stops attempting to recover the index.
confluen	ce.cluster.ind	dex.recovery.num.attempts
5.9.1	1	In Confluence Data Center, the number of times that a node needing index recovery will attempt to recover its index. Set this property to 0 to disable index recovery on that node (for example, when you want to force a node to automatically rebuild its own index on startup).
com.atla	ssian.conflue	ence.officeconnector.canary.memory_value
5.9.1	1024	Sets the memory (in megabytes) available to the Office Connector Canary process, which is a workaround for a known issue with the Import from Word option. See JVM crashes during Import from Word in Confluence for more information.
com.atla	ssian.conflue	ence.officeconnector.canary.timeout
5.9.1	120000	Sets the maximum timeout (in milliseconds) for the Office Connector Canary process, which is a workaround for a known issue with the Import from Word option. See JVM crashes during Import from Word in Confluence for more information.
atlassian	.plugins.ena	ble.wait
5.9.5	300	Set this property to increase the default time to wait for plugins to start up. This is useful if you have problems with plugins not starting up in time, causing Confluence to fail to start.
confluen	ce.cluster.ha	izelcast.max.no.heartbeat.seconds
5.9.7	30	In Confluence Data Center, this sets how long (in seconds) a node can be out of communication with the cluster before it's removed from the cluster. See balancing uptime and data integrity for more info on when you may want to change this setting.
confluen	ce.startup.re	migration.disable
5.10.8	False	Set this property to true if you repeatedly experience issues with Confluence creating a new page version as it attempts to migrate pages containing unmigrated wiki-markup macros each time a plugin is install or updated. See
		CONFSERVER-37710 CLOSED for more details.
cluster.s	afety.time.to.	live.split.ms
6.0.0	60000	In Confluence Data Center, this is the amount of time (in milliseconds) that the cluster safety job will wait to allow the nodes to rejoin after a split brain is detected. If the node still can't find the cluster safety number in the cache after this time, the node will panic.
confluen	ce.cph.max.	entries
6.0.0	2000	This is the maximum number of pages that can be copied when you copy a page and its child pages. Increase this if you need to copy a page with more than the default number of child pages.
confluen	ce.cph.batch	ı.size
6.0.0	10	This is the number of pages copied in each batch when you copy a page and its children. Increase or reduce this number if you experience problems copying a page with many child pages.

C O O	0004	This is the part that Completeness the complete that passesses called a rational additional responses
6.0.0	8091	This is the port that Synchrony, the service that powers collaborative editing, runs on. You should only need to change this if port 8091 is not available. From 6.0.4, Confluence Server will accept either reza.port or synchrony.port.
synchro	ny.memory.ma	ax (formerly reza.memory.max)
6.0.0	2g	This is the maximum heap size (Xmx) allocated to Synchrony, the service that powers collaborative editing. Change this value if you need to increase or decrease the heap size. From 6.0.4, Confluence Server will accept either reza.memory.max or synchrony.memory.max.
		The default value of this property was increased to 2 gigabytes in 7.10.0.
synchro	ny.stack.spac	Э
6.0.0	2048k	This sets the stack size (Xss) of the Synchrony JVM. Increase if you experience stack overflow errors, or decrease if you experience out of memory errors from Synchrony.
		This property only applies when Synchrony is managed by Confluence.
synchro	ny.enable.xhr.	fallback
6.0.0	True	XML HTTP Request (XHR) fallback allows a user, who cannot connect to Synchrony via a WebSocket, to use the Confluence Editor. From Confluence 6.1 this is enabled by default, and only used when a WebSocket connection is not available. You should not need to disable this, unless directed by our support team.
synchro	ny.database.te	est.connection.on.checkin
6.0.0	True	Verifies the connection to the database is valid at every connection checkin. Atlassian Support may suggest you set this property to False if you experience performance issues in sites that have very frequent page edits.
synchro	ny.proxy.enab	led
6.0.0	True	By default, Confluence uses an internal proxy to communicate between the Confluence JVM and Synchrony JVM. See Administering Collaborative Editing for more information. In Confluence 6.0, set this property to true to manually enable the internal proxy (useful if you have configured a reverse proxy and want to also use the internal Synchrony proxy). In Confluence 6.1 or later it should not be necessary to use this system property, as Confluence intelligently determines when to use the proxy. This property only applies when Synchrony is managed by Confluence. It has no effect
		on a Synchrony standalone cluster.
synchro	ny.bind (forme	erly known as reza.bind)
6.0.0	0.0.0.0	This is the specific network interface that Synchrony listens on. It is unlikely that you will need to change this when Synchrony is managed by Confluence. In Confluence Data Center, when running a Synchrony standalone cluster this should be set to the same value assynchrony.cluster.bind. From 6.0.4, Confluence Server will accept either reza.bind or synchrony.bind.
synchro	ny.context.pat	h
6.0.0	/synchro ny	This is the context path for Synchrony. There should be no need to change this in Confluence Server, or when Synchrony is managed by Confluence.
	1	permits.size

6.0.0	(number of cores)	This property sets the number of concurrent PDF exports that can be performed. It defaults to the number of cores on your server or cluster node.
confluen	ce.pdfexport.t	imeout.seconds
6.0.0	30	This property sets the amount of time (in seconds) a new PDF export request should wait before failing, if the maximum number of concurrent PDF exports (set inconfluence. pdfexport.permits.size) has already been reached.
confluen	ce.flyingpdf.d	efault.characters.per.line
6.0.3	80	If the total characters in a table column heading exceeds the value of this property, Confluence will automatically adjust the width of the other table columns so that all columns will fit within one page when the page is exported to PDF.
synchro	ny.host	
6.0.4	127.0.0.1	This is the IP that Confluence uses to connect to Synchrony. It defaults to localhost. Cha nge this if you need to allow Confluence to contact Synchrony via a custom hostname or IP address.
		This property only applies when Synchrony is managed by Confluence. It has no effect on a Synchrony standalone cluster.
synchroi	ny.proxy.healt	hcheck.disabled
6.1.0	false	The Synchrony proxy health check is used to check whether the Synchrony proxy is running and responding to requests. It requires a http connection. If a http connector is not present in your <code>server.xml</code> (for example you're using a https or AJP connector) the health check will fail even if the Synchrony proxy is operational. You can use this system property to disable the health check if necessary.
page.inc	lex.macro.ma	x.pages
6.1.4	1000	Sets the maximum number of pages that the Page Index macro can display. The Page Index macro can significantly slow down your Confluence instance and cause out of memory errors when used in a space with a large number of pages. If the number of pages in the space exceeds this limit, the Page Index macro will show a page count, and a message that there are too many pages to display.
		The default value of this property was reduced to 1000 in Confluence 7.11.
atlassiar	n.indexing.atta	uchment.maxsize
6.2.2	1048576 00	When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename.
		If the uploaded file is larger than the limit set by this property (default is 100MB, in bytes), text extraction and indexing will be skipped. See Configuring Attachment Size for more info.
officecor	nnector.excel.	extractor.maxlength
6.2.2	1048576	When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename.
		Confluence will only extract content from an Excel spreadsheet up to the limit set by this property (default is 1MB, in bytes), before proceeding to index it. This will mean that only part of file's contents are searchable. See Configuring Attachment Size for more info.
atlassiar	n.image_filter.	transform.max_data_size

6.2.2	16000000	Applying image effects to large images can cause out of memory errors. We prevent users from applying image effects to images with a data size greater than 16MB.
		Set this property, in bytes, to reduce the maximum data size. Note: this is the data size not the size of the file on disk.
atlassia	 n.image_filter.t	ransform.max.pixel
6.2.2	4000	Applying image effects to large images can cause out of memory errors. We prevent users from applying image effects to images larger than 4000x4000 pixels.
		Set this property, in pixels, to change the maximum image dimensions.
confluer	nce.collab.edit.	user.limit
6.3.0	12	When collaborative editing is enabled, this sets the maximum number of users that can simultaneously edit a page. Reduce this number if you experience degraded performance when many people are editing.
jobs.limi	it.per.purge	
6.4.3	2000	The Purge Old Job Run Details scheduled job deletes details of old scheduled jobs from the database in batches.
		Set this property to change the number of records to remove in each batch.
all.jobs.t	ttl.hours	
6.4.3	2160	By default, the Purge Old Job Run Details scheduled job deletes details of successful scheduled jobs older than 90 days (or 2160 hours).
		Set this property, to change number of hours to keep details of successful jobs in the database.
unsucce	essful.jobs.ttl.ho	ours
6.4.3	168	By default, the Purge Old Job Run Details scheduled job deletes details of unsuccessful (failed or aborted) scheduled jobs older than 7 days (or 168 hours).
		Set this property, to change number of hours to keep details of unsuccessful jobs in the database.
hide.sys	stem.error.deta	ils
6.5.0	False	Set this property to true if you want to hide details on the error screen that appears in the browser when Confluence can't start up. This can be useful for public-facing sites, where you may not want to display details of the problem.
atlassia	n.recovery.pas	sword
6.6.1		Allows an administrator to start Confluence in recovery mode and specify a temporary administrator password. This is useful if the administrator is locked out of the instance after a site import, or cannot reset their password by other methods. See Restore Passwords To Recover Admin User Rights for more information.
		This system property must be removed immediatley after the admin account has been recovered or a new admin account created.
confluer	nce.extra.userli	ster.limit
6.6.3, 6.7.1	10000	Set this property to change the maximum number of people the User List macro can display. This macro is known to cause out of memory errors when attempting to display a very large number of users.

	document.co	pool.size priversion.sandbox.pool.size)
6.10.0	2	Use this property to increase the number of processes (sandboxes) in the external process pool. More processes means more tasks can be executed in parallel, but will consume more memory and CPU resources on each node. This property only applies to Data Center. This property was renamed in Confluence 6.12.
conversi	an eandhay e	tartup.time.limit.secs
		onversion.sandbox.startup.time.limit.secs)
6.10.0	30	When a document file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. In Confluence Data Center this is handled by the external process pool .
		This property sets the amount of time (in seconds) that a process will wait for document conversion to start, before terminating the process.
		This property only applies to Data Center. This property was renamed in Confluence 6.12.
documer	nt.conversion	sandbox.request.time.limit.secs
6.10.0	30	When a document file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. In Confluence Data Center this is handled by the external process pool .
		This property sets the amount of time (in seconds) that a process will wait for document conversion to complete, before terminating the process, and marking thumbnail generation for that file as failed.
		This property only applies to Data Center.
		From Confluence 7.8.0 this property also applies to the HTML conversion that occurs when you insert a file using the Office Word or Office Excel macro.
sandbox	termination.t	olerance
6.10.0	5	This property specifies how often a process in the external process pool will check if the request time limit (set in the request time.limit.secs property above) has been exceeded. It's calculated by dividing the request time limit by the value of this property. For example, if the request time limit is 30 seconds, and the tolerance set in this property is 5, the process will check if the request time limit has been exceeded every 6 seconds.
		This property only applies to Data Center.
		nemory.limit.megabytes onversion.sandbox.memory.limit.megabytes)
6.10.0	512	When a document file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. In Confluence Data Center this is handled by the external process pool .
		This property limits the amount of heap memory each process in the external process pool can consume.
		This property only applies to Data Center. This property was renamed in Confluence 6.12.
		sandbox.log.level

6.10.0	INFO	Use this property to change the logging level of document conversion in the external process pool to WARNING, INFO, or FINE. This is useful if you need to troubleshoot a problem with the sandbox. This property only applies to Data Center.
sandbox.	error.delay	.millis
6.10.0	50	This property sets how often (in milliseconds) document conversion errors are captured for diagnostic purposes.
		This property only applies to Data Center.
documer	t.conversio	n.sandbox.disable
6.10.0	false	Set this property to true if you don't want to handle document conversion (thumbnail generation) in the external process pool .
		When disabled, document conversion will be handled in the Confluence JVM, as is the case in Confluence Server.
		This property only applies to Data Center.
conversion	on.sandbox	.java.options
6.10.0		When a document or image file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. In Confluence Data Center this is handled by theexternal process pool.
		Use this property to override the default value of any of the following Confluence Server system properties (introduced in 7.0.1), and pass new values directly to the JVMs in the external process pool:
		 confluence.document.conversion.imaging.convert.timeout confluence.document.conversion.imaging.convert.timeout confluence.document.conversion.imaging.enabled.tif confluence.document.conversion.imaging.enabled.psd
diagnosti	cs.os.chec	k.period.secs
6.11.0	120	Set this property to change how often operating system diagnostics checks should be performed (in seconds).
		This property only applies to the Low free memory (OS-1001) and Low free disk space (OS-1002) alerts.
diagnosti	cs.os.min.f	ree.memory.megabytes
6.11.0	256	This property applies to the free memory diagnostic alert (OS-1001).
		Set this property to change the threshold at which the amount of free memory (in megabytes) should trigger this alert.
diagnosti	cs.os.min.f	ree.disk.space.megabytes
6.11.0	8192	This property applies to the free disk space diagnostic alert (OS-1002).
		Set this property to change the threshold at which the amount of free disk space (in megabytes) in the local or shared home directory should trigger this alert.
diagnosti	cs.slow.htti	p.request.secs

6.11.0	60	This property applies to the HTTP request diagnostic alert (HTTP-1001). This alert is disabled by default.
		Set this property to change the threshold (in seconds) at which a slow HTTP request should trigger this alert.
diagnosti	cs.slow.lon	ng.running.task.secs
6.11.0	300	This property applies to the long running task diagnostic alert (JOB-1001).
		Set this property to change the threshold (in seconds) at which a long running task should trigger this alert.
diagnosti	cs.slow.ma	acro.rendering.secs
6.11.0	30	This property applies to the macro rendering diagnostic alert (MACRO-1001). This alert is disabled by default. Set this property to change the threshold (in seconds) at which rendering a macro on a
		page should trigger this alert.
diagnosti	cs.jvm.mer	mory.check.period.secs
6.11.0	10	Set this property to change how often JVM diagnostics checks should be performed (in seconds).
		This property only applies to the Thread memory allocation rate (JVM-1001) and Garbage collection (JVM-1002)alerts.
diagnosti	cs.jvm.mer	mory.allocation.rate.percent
6.11.0	5	This property applies to the thread memory allocation rate diagnostic alert (JVM-1001). This alert is disabled by default.
		Set this property to change the threshold (as a percentage) at which the memory allocation to a particular thread, during the monitoring period (set in diagnostics.jvm. memory.allocation.monitoring.period.secs), should trigger this alert.
diagnosti	cs.jvm.mer	mory.allocation.monitoring.period.secs
6.11.0	20	This property applies to the thread memory allocation rate diagnostic alert (JVM-1001). This alert is disabled by default.
		Set this property to change the monitoring period (in seconds) for this alert.
diagnosti	cs.jvm.garl	page.collector.percent
6.11.0	10	This property applies to the garbage collection diagnostic alert (JVM-1002). This alert checks how much memory has been allocated to garbage collection during the monitoring period (set in diagnostics.jvm.garbage.collector.monitoring.period.secs).
		Set this property to change the threshold (as a percentage) at which the memory allocated to garbage collection should trigger this alert.
diagnosti	cs.jvm.garl	page.collector.monitoring.period.secs
6.11.0	20	This property applies to the garbage collection diagnostic alert (JVM-1002).
		Set this property to change the monitoring period (in seconds) for this alert.
diagnosti	cs.alert.ret	ention.period.days
6.11.0	30	Set this property to change how often diagnostic alert data should be retained in the database (in days).

6.11.0	30	Set this property to change how often we check for, and remove diagnostic alert data that is older than 30 days (the limit set by diagnostics.alert.retention.period.days)
pdf.expo	rt.sandbox.	disable
6.12.0	false	Set this property to true if you don't want to handle PDF exports in the external process pool.
		When disabled, PDF exports will be handled in the Confluence JVM, as is the case in Confluence Server.
		This property only applies to Data Center.
pdf.expo	rt.sandbox.	request.time.limit.secs
6.12.0	180	When you export a space to PDF, Confluence exports the content of each page to HTML, converts that HTML to PDF, and then finally merges all the pages together into a single PDF file.In Confluence Data Center this is handled by the external process pool.
		This property sets the amount of time (in seconds) that a process should wait to complete, before being terminated. This time limit applies both to the time to convert the content from HTML to PDF, and the time to merge the final PDF file.
		This property only applies to Data Center.
pdf.expo	rt.sandbox.	memory.requirement.megabytes
6.12.0	64	In Confluence Data Center PDF exports are handled by the external process pool .
		This property sets the minimum memory (in megabytes) that a sandbox process must have available to start a PDF export. Ifconversion.sandbox.memory.limit.megabytesis set to less than the value of this property, PDF export will not start. We don't recommend setting this property to less than 64MB.
		This property only applies to Data Center.
synchron	y.eviction.s	soft.job.threshold.hours
7.0.1	72	This property changes the behaviour of the Synchrony softeviction scheduled job.
		It sets the minimum time, in hours, since a Synchrony change log was last modified, to make it eligible to be cleaned up. By default, only Synchrony change logs last modified more than 3 days ago, for pages/blogs that do not have an active editing session, will be evicted.
synchron	y.eviction.l	nard.job.threshold.hours
7.0.1	360	This property changes the behaviour of the Synchrony hardeviction scheduled job
		It sets the minimum age, in hours, of content eligible to be evicted by the hard eviction scheduled job. By default, all Synchrony data for any content that is 15 days or older will be evicted by this job, regardlessof whether it has been modified more recently.

7.0.4		
7.0.1	30	When a complex imagefile (such as ICO, EMF, WMF, plus TIF and PSD if enabled) is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. This process is known to cause out of memory errors for large or complex files. This property sets the amount of time (in seconds) that Confluence will wait for conversion to complete for an image file, before terminating the process. This property applies to Confluence Server and Data Center. For Data Center also see Double Contents.
		cument conversion for Confluence Data Center
confluen	ce.document.	.conversion.slides.convert.timeout
7.0.1	30	When a presentationfile (such as PPT, PPTX, POT)is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. This process is known to cause out of memory errors for large or complex files.
		This property sets the amount of time (in seconds) that Confluence will wait for conversion to complete for a presentation file, before terminating the process.
		This property applies to Confluence Serve and Data Center.For Data Center seeDocum ent conversion for Confluence Data Center.
confluen	ce.document.	.conversion.imaging.enabled.tif
7.0.1	false	When a file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. By default thumbnails are not generated for TIFF / TIF as they're known to cause out of memory errors.
		Set this property to 'true' to turn on thumbnail generation for TIFF files.
		If enabled, a timeout will apply to this type of file. This is set by theconfluence.document. conversion.imaging.convert.timeout system property.
		This property applies to Confluence Server and Data Center.For Data Center seeDocument conversion for Confluence Data Center
confluen	ce.document.	.conversion.imaging.enabled.psd
7.0.1	false	When a file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. By default thumbnails are not generated for Photoshop PSD files as they're known to cause out of memory errors.
		Set this property to 'true' to turn on thumbnail generation for PSD files.
		If enabled, a timeout will apply to this type of file. This is set by theconfluence.document. conversion.imaging.convert.timeout system property.
		This property applies to Confluence Server and Data Center.For Data Center seeDocument conversion for Confluence Data Center.
officecor	nector.power	rpoint.extractor.maxlength
7.0.1	1048576	When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename.
		Confluence will only extract content from a PowerPoint presentation up to the limit set by this property (default is 1MB, in bytes), before proceeding to index it. This will mean that only part of file's contents are searchable. See Configuring Attachment Size for

7.2.0	3000	Maximum width, in pixels, that a chart macro can be set to display on a page. If a higher value is entered, the chart will automatically be reduced to this default.
confluen	ice.chart.ma	cro.height.max
7.2.0	3000	Maximum height, in pixels, that a chart macro can be set to display on a page. If a higher value is entered, the chart will automatically be reduced to this default.
confluen	ice.search.lu	ucene.termFilterBitSetThreshold
7.2.0	20	Set this property to change the behaviour of the term filter in Confluence's Lucene based implementation of search. A bitset is only created when the number of matched documents times the threshold set in this property exceeds total number of documents. This will reduce memory usage and improve Confluence performance. You shouldn't need to change this threshold.
page-tre	e.partial-loa	ding-batch-size
7.3.0	200	Confluence limits the number of pages that initially display at each level of the page tree. This helps safeguard the performance of your site. Set this property to increase or decrease the number of pages to initially display at each level of the page hierarchy. At least one child page is always displayed, so if you set the value to 10 for example, you may see 11 or 12 pages.
page-tre	e.partial-loa	ding.disable
7.3.0	false	Confluence limits the number of pages that initially display at each level of the page tree. This helps safeguard the performance of your site. Set this property totrue if you always want all pages to be displayed by default in the page tree.
confluen	ice.word.imp	port.maxsize
7.3.3	20	Sets the maximum uncompressed size, in megabtyes, of a Microsoft Word document that can be imported into Confluence using the Import from Word feature. This is to prevent very large files from causing out of memory errors.
gatekee	per.request-	timeout.seconds
7.4.0	60	Some open-ended queries can take a long time to display when you Inspect Permissions in a Confluence Data Center site with a lot of spaces and users. To avoid this having an impact on your site, we have a timeout.
		Set this property to change the length of the timeout, in seconds. You can't set this lower than 30 seconds, or higher than 900 seconds (15 minutes).
		v.sysadmin.only iew.sysadmin.only)
7.5.0	false	Set this property to true to restrict the Audit Log page in Confluence Administration to people with System Administrator global permissions. By default people with System Administrator or Confluence Administrator global permission can access this feature. This property applies to Confluence Server and Data Center.
		This property was renamed in 7.7.0.
plugin.a	udit.file.max	file.count
7.5.0	100	In Confluence Data Center, audit events are written to an audit log file in the local home directory. These log files are rotated once they reach the file size set by plugin.audit.file. max.file.size.
		Set this property to change the total number of audit log files to keep in the file system.
nlugin a	udit.file.max	file.size

7.5.0	100	In Confluence Data Center, audit events are written to an audit log file in the local home directory.
		Set this property to change the maximum file size, in megabytes, that a log file can reach before the log is rotated.
legacy.a	udit.migrato	r.num.threads
7.5.0		Sets the number of threads to be used for migrating existing audit log events to the new audit log format when upgrading to Confluence 7.5 or later.
		There is no default value. When not set, Confluence will use 2 times the number of processors available.
		To increase throughput, you can set this property to use a specific number of threads. For example, if you set this property to 16, then 16 threads will be used to migrate the audit log events, regardless of the number of processors available. Note that more threads will result in higher database utilisation.
legacy.a	udit.migrato	r.batch.size
7.5.0	1000	Sets the number of audit log events to migrate to the new audit log format, when upgrading to Confluence 7.5 or later. This migration can take a long time in large sites.
		When testing your upgrade, if you find the migration is causing performance issue on your Confluence and/or database server, you can set this property to decrease the number of records processed in each batch.
Number	ItemPerPage	eOfPaginatedListAction
7.5.0	30	Sets the number of items to be listed on the Undefined Pages and Restricted Pages tabs in Space Tools before Confluence paginates the results.
		Set this property to change the number of items to display per page.
Caching	EnablingIter	nNumber
7.5.0	1000	The Undefined Pages tab in Space Tools provides a list undefined links to pages that do not yet exist. As this list can be quite memory intensive to generate, the results are cached.
		Set this property to change the maximum number of undefined link references to cache.
Caching	EnablingIter	nTimeout
7.5.0	5	The Undefined Pages tab in Space Tools provides a list undefined links to pages that do not yet exist. As this list can be quite memory intensive to generate, the results are cached.
		Set this property to change the amount of time, in minutes, the results should be cached (time to live).
confluer	ice.event.du	ration_checker.threshold_in_seconds
7.6.2	60	When asynchronous events are generated faster than Confluence can process them, we write a message to the application log to advise that we'll process each task synchronously until the queue is cleared.
		Set this property to change how often this message appears, in seconds.
		This property is also available in the 7.4 Long Term Support release from 7.4.4.
	_	·

7.6.2	false	Set this property to true if you need to disable TLS hostname verification for any SMTP mail servers you've configured in Confluence. This is not recommended as it can increase the risk of a Man In The Middle attack.
		This property is also available in Long Term Support releases from 6.13.17 and 7.4.5.
atlassia	n.image_filte	er.transform.max.pixel.drop_shadow
7.7.0	2000	Applying a drop shadow image effect to large images can cause out of memory errors. We prevent users from applying a drop shadow to images larger than 2000x2000 pixels, or the value set byatlassian.image_filter.transform.max.pixel(whichever is smallest). Set this property, in pixels, to change the maximum image dimensions for drop shadow.
		This property is also available in the 7.4 Long Term Support release from 7.4.4.
pagePro	opertiesRep	ortContentRetrieverMaxResult
7.7.0	3000	The Page Properties Report macro displays a maximum of 3000 pages. Set this property to decrease the maximum number of pages the macro can display. We dont recommend you increase this limit, as it can have a performance impact on your site.
		This property is also available in the 7.4 Long Term Support release from 7.4.6.
conflue	nce.webhool	ks.allow.all.hosts
7.7.0	false	Set this property to true allow administrators to configure localhost URLs as a webhook endpoint. This is disabled by default for security reasons because all ports on the same network can be pinged by the UI.
audit.lo	g.search.disa	abled
7.8.0	false	Set this property to true if you do not want to log the Search performed event in the audit log. This event is included by default when the End User Activity coverage area is set to Full, and records the search terms entered in search, advanced search, and in search macros (such as Livesearch and Page Tree Search).
		This property only applies to Data Center.
docume	ent.conversion	on.sandbox.memory.requirement.megabytes
7.8.0	128	When a Word or Excel document file is inserted into a page using the Office Word or Office Excel macro, its contents are converted to a format that can be viewed inline and previewed. In Confluence Data Center this is handled by the external process pool. This property limits the amount of heap memory, in megabytes, this process in the
		external process pool can consume.
		This property only applies to Data Center.
atlassia	n.pats.enabl	led
7.9.0	true	Personal access tokens are used to authenticate REST API requests. By default, any user can create a personal access token.
		Set this property to false to prevent users from creating tokens. Any existing tokens can't be used to authenticate while this property is set to false.
atlassia	n.pats.etern	al.tokens.enabled
		Personal access tokens are used to authenticate REST API requests. Set this property f
7.9.0	true	alse to prevent users from creating tokens that never expire.

7.9.0	365	Personal access tokens are used to authenticate REST API requests. Set this property, in days, to define the maximum days to until expiry.
atlassian	.pats.max.tol	kens.per.user
7.9.0	10	Personal access tokens are used to authenticate REST API requests. Set this property to limit the number of tokens an individual user can create.
atlassian	.allow.insecu	re.url.parameter.login
7.10.0	False	Set this property to true to use the os_username+os_password query parameter to log in to Confluence.
		This login method is blocked by default, and we strongly recommend you use a more secure method, such as personal access tokens.
com.atla	ıssian.conflu	uence.extra.calendar3.concurrent.task.max
7.11.0	20	The max number of worker threads used by Team Calendars.
		This property only applies to Data Center.
com.atla	ıssian.conflu	uence.extra.calendar3.greenhopper.sprint.enabled
7.11.0	true	If set to false integrating calendars with Jira sprints is disabled.
		This property only applies to Data Center.
com.atla	ıssian.conflu	uence.extra.calendar3.jira.timeout.socket
7.11.0	10000	Timeout in milliseconds for calendars to connect to Jira. Increasing this may help if you experience timeouts on Jira calendars.
		This property only applies to Data Center.
com.atla	ıssian.conflu	uence.extra.calendar3.jira.issues.max
7.11.0	1000	The maximum number of events that will be loaded from a Jira calendar. Increasing this may cause performance issues.
		This property only applies to Data Center.
com.atla	ıssian.conflu	uence.extra.calendar3.display.events.dashboard.maxperday
7.11.0	10	The maximum number of calendar events that will be shown per day on the upcoming events view on the Confluence dashboard.
		This property only applies to Data Center.
com.atla	ıssian.conflu	uence.extra.calendar3.display.events.calendar.maxpercalendar
7.11.0	200	The maximum number of events that can be displayed from a single calendar.
		This property only applies to Data Center.
com.atla	ıssian.conflu	uence.extra.calendar3.display.events.calendar.maxperdaysummary
7.11.0	3	The maximum number of events that will be shown per day in the summary email.
com.atla	ıssian.conflu	uence.extra.calendar3.display.events.calendar.maxdailysummary
7.11.0	4	The maximum number of calendar events that will be shown in total in the daily summary email.
		This property only applies to Data Center.

7.11.0	8	The maximum number of calendar events that will be shown in total in the weekly summary email.
		This property only applies to Data Center.
com.atla	assian.con	fluence.extra.calendar3.display.timeline.calendar.maxmonth
7.11.0	6	The maximum number of months that will be shown in the timeline view of a calendar.
		This property only applies to Data Center.
plugin.da	ata.pipeline	embedded.line.break.preserve
7.12.0	false	Specifies whether embedded line breaks should be preserved in the output files. Line breaks can be problematic for some tools such as Hadoop.
		This property is set toFalseby default, which means that line breaks are escaped.
plugin.da	ata.pipeline.	embedded.line.break.escape.char
7.12.0	\\n	Escaping character for embedded line breaks. By default, we'll print\nfor every embedded line break.
reindex.a	attachments	s.freeMbPerThread
7.14.0	500	Confluence calculates the number of threads to be used when reindexing attachments based on either the number of available CPUs or free memory, whichever is lower. Seen eindex.attachments.thread.count for more details.
		Set this property to change the amount of free memory per thread (in megabytes) used in this calculation. Reducing this value may result in more threads being used during reindexing, which allows for greater concurrency.

Working with Confluence Logs

Confluence uses Apache's log4j logging service. This allows administrators to control the logging behavior and the log output file. There are sixlog4j logging levels.

If you request help from Atlassian Support, we will almost always ask for the Confluence application logs. The easiest way to get these logs is to go to Seneral Configuration > Troubleshooting and support toolsand follow the prompts to create a Support Zip.

On this page:

- Application log files
- Change the log file configuration
- Change the destination of the log files
- Change the size and number of log files
- Change the logging levels
- Specific Confluence logging options
- Scanning log files for known problems
- Mark logs when troubleshooting issues
- Tomcat logs

Related pages:

- Enabling Detailed SQL Logging
- Enabling user access logging
- Generating a Thread Dump

Application log files

By default, the application log files can be found in the <local-home>/logsdirectory. This location is configurable, so you may need to check the log config for your location.

To make troubleshooting problems easier, the application log is split into several distinct log files:

confluence-application.log

This is the main application log file, most entries will be written here. When you start ConfluenceAny log entries written to the console will also be repeated in this log.

atlassian-confluence-index.log

This file contains entries related to the search index.

atlassian-confluence-outgoing-mail.log

This file contains entries related to outgoing mail, such as notifications.

atlassian-confluence-security.log

This file contains entries related to your users and user directories.

atlassian-synchrony.log

This file contains entries related to Synchrony, which powers collaborative editing.

atlassian-diagnostics.log

This file contains entries for an experimental diagnostics feature which provides alerts for things like low disk space and memory.

In our documentation, when we refer to the "application log", we are referring to any of these files.

You can check the exact classes or packages that are logged to each file in the log4j.properties file underLOGGING LOCATION AND APPENDER.

Change the log file configuration

The logging behavior for Confluence and Synchrony 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.

Change the destination of the log files

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 log4 j.properties file.

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. 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.

The Synchrony log destination can also be changed in the same way in file.

Confluence ships with the full suite of appenders offered by log4j. Read more about appenders in the log4j documentation.

For more detailed information seeConfiguring log4j in Confluence to send specific entries to a different log file in our Knowledge Base.

Note:If you change the location of your log files, they will no longer be included when you generate a support zip. This means you'll need to attach your logs to any support requests manually.

Change the size and number of log files

By default, Confluence keeps 5 application log files, which are overwritten as they reach 20 MB.

You can change the default log size and the number of log files to keep by editing the following values in <CO NFLUENCE-INSTALL>/confluence/WEB-INF/classes/log4j.propertiesfile.

```
log4j.appender.confluencelog.MaxFileSize=20480KB
log4j.appender.confluencelog.MaxBackupIndex=5
```

Change the logging levels

This can be done in the Confluence UI. See Configuring Logging for instructions on how to change the logging configuration of Confluence.

Specific Confluence logging options

Here's some specific log configurations you may need when troubleshooting.

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. SeeEnabling Detailed SQL Logging.

Log the details of users accessing each Confluence page

Access logging using Tomcat Valve is enabled by default from Confluence 7.11. These logs are not part of the application log, and can be found at<install directory>/logs/conf_access_log.<date>.log.

You can however configure the application log to show which users are accessing which pages in Confluence. SeeHow to Enable User Access Loggingin our Knowledge Base.

Scanning log files for known problems

Atlassian Troubleshooting and support tools includes a log analyzer that can check for you Confluence logs for errors and match them against known problems in our knowledge base and issue tracker.

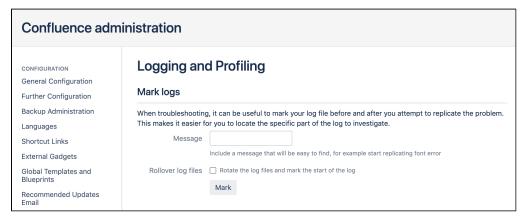
See Troubleshooting Problems and Requesting Technical Supportto find out how to set up a periodic scan of your log files.

Mark logs when troubleshooting issues

When troubleshooting it can be useful to mark your log files before and after you attempt to replicate the problem. This makes it easier for you to locate the specific part of the log to investigate.

To mark the application log files:

- 1. Go to General Configuration>Logging and Profiling.
- 2. Enter a message, for example "Reproduce directory sync issue".
- 3. Select**Rollover log files**if you want to start a new log file with your mark (this will delete the oldest log file).
- 4. SelectMark.



Screenshot: The logging and profiling screen

Your message will be added to all the application log files (such as atlassian-confluence.log, and atlassian-confluence-security.log). You can mark your logs as often as you need to.

Here's an example:

Tomcat logs

There are some additional logs in your Confluence install directory that can be useful when troubleshooting issues with your Confluence site.

- <install directory>/logs/catalina-<date>.log
 This log records Tomcat operations, such as starting and stopping application server.
- <install directory>/logs/conf_access_log.<date>.log
 This is where you find Confluence's access logs. These logs are configured in the server.xml. SeeT omcat Access Log Valve documentation for further configuration options.
- <install directory>/logs/gc.<date>.log
 This is where you find the garbage collection logs. These logs provide useful information if you're experiencing long GC pauses.

Configuring Logging

There are many situations where you may want to change what is written to the Confluence application logs, particularly when troubleshooting a specific problem.

You can temporarily change the logging behaviourin the Logging and Profiling screen, while Confluence is running. Your changes will be discarded when you restart Confluence.

Alternatively, you can permanently change the logging behaviour in the log4j properties file.

On this page:

- Temporarily change logging behaviour while Confluence is running
- Permanently change logging level in the properties file
- Configuring Levels for java.util.logging in logging.properties

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.

Temporarily change logging behaviour while Confluence is running

When troubleshooting a problem, you can temporarily change the logging behaviour while Confluence is running. These changes aren't written to the log4j.properties file so are discarded when you next stop Confluence.

Change the logging level of an existing class or package

You need **System Administrator** global permissions to do this. Seelog4j Logging Levels for information on the specific levels.

To change the logging level of an existing class or package:

- 1. Go to O-Seneral Configuration Logging and Profiling.
- 2. Locate the relevant class or package, and select a value from the **New Level** dropdown.
- 3. Save your changes.

Alternatively, choose **Remove** if you want to stop logging a particular class or package.

Remember, your changes will not be written to the log4j.properties file so will be discarded when you next stop Confluence.

Add logging for an additional class or package

Sometimes our support team may ask you to enable some additional logging when troubleshooting a specific problem. You need **System Administrator** global permissions to do this.

To set the logging level for a new class or package:

- 1. Go to Configuration Logging and Profiling.
- 2. Enter the name in the Class/Package name field. It will be something likecom.atlassian. confluence.example.
- 3. Select Add entry to add the new class or package to the list.
- 4. Locate the relevant class or package, and select a value from the New Level dropdown.
- 5. Saveyour changes.

Turn page profiling on or off

SeeTroubleshooting Slow Performance Using Page Request Profiling for more information on when and how to use page profiling.

Turn detailed SQL logging on or off

SeeEnabling Detailed SQL Loggingfor more information on when and how to use detailed SQL logging.

Preset logging configurations

Confluence provides two preset log configurations:

Production

This is the recommended default configuration, which aims to provide the most important information without flooding your logs.

Diagnostic

This changes the logging level of most packages to debug. It can be useful when troubleshooting, but results in slower performance and will fill up your log files more quickly.

If you want to reset your logging configuration to the default, select**Production**.

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** Current Class/Package Name **New Level** Level **INFO** atlassian.plugin INFO ▼ Remove INFO com.atlassian.confluence.admin.actions INFO ▼ Remove

Permanently change logging level in the properties file

You need access to the installation directory to do this.

Thelog4j.propertiesfile lists all the default classes and packages, plus some additional packages that you can choose to enable.

To permanently change the logging level of a class or package:

1. Stop Confluence.

- 2. Edit the < CONFLUENCE INSTALL > / confluence / WEB INF / classes / log4j.properties file.
- 3. Under **Logging Levels**, locate the package you want to change. Remove the # symbol to uncomment the package if necessary.
- 4. Change the logging level to DEBUG, INFO, WARNING, ERROR, or FATAL. Seelog4j Logging Levels for accepted values. For example:

```
log4j.logger.com.atlassian.confluence.cache=DEBUG
```

5. Save the file, and restart Confluence for the changes to take effect.

If you're running Confluence in a cluster, you will need to repeat your change on every node. You can take each node down one by one, there's no need to stop the whole cluster.

If you want to log a class or package that's not already listed, simply add it to the file. Seethelog4j documentation for more information.

If you want to change the location, number, or size of the logs, seeWorking with Confluence Logs.

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

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 <CON FLUENCE-INSTALL>/confluence/WEB-INF/classes/. The logging levels are different from log4j. Seej ava.util.loggingin the Java 11 documentation.

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 suboptimally
- 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
- (i) 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** (changes made here will not persist across restarts).
 - 2. Manually modifying the <Confluence-Install>\confluence\WEB-INF\classes\log4j. properties file.

Default Log Level

The standard Confluence log level **WARN** is a way for Confluence to communicate with the server administrator. Logging at WARN level and higher should be reserved for situations that require some kind of attention from the server administrator, and for which corrective action is possible.

Seelog4j manualfor more information.

Troubleshooting SQL Exceptions

If you get an exception similar to those shown below, it is a good idea to increase the logging levels of your Confluence instance. If you request Atlassian support, this additional logging will help us work out the cause of the error.

Increased logging levels will enable us to diagnose errors like these:

```
org.springframework.dao.DataIntegrityViolationException: (HibernateTemplate): data integrity violated by SQL ''; nested exception is java.sql.BatchUpdateException: Duplicate entry '1234' for key 1 at org.springframework.jdbc.support.SQLStateSQLExceptionTranslator.translate (SQLStateSQLExceptionTranslator.java:88) caused by: java.sql.BatchUpdateException: Duplicate entry '1234' for key 1 at com.mysql.jdbc.ServerPreparedStatement.executeBatch(ServerPreparedStatement.java:647)
```

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.

(i) 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
```

- 1 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 *ta ble* 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

Configure access logs

Access logs record every request made to your site which can be useful for auditing purposes and when troubleshooting a problem with an integration, app, or feature

View access logs

The log file is located in <install-directory>/logs/conf_access_log<date>.log.

Here's an example of the log output:

Access log pattern

The Confluence access log uses the Apace Tomcat access log valve. The information recorded about each request is configurable.

The default pattern is:

```
%t %{X-AUSERNAME}o %I %h %r %s %Dms %b %{Referer}i %{User-Agent}i
```

This will log the date and time, username, remote logical username, remote host name (or IP), the first line of the request, the HTTP status code of the response, time taken to process the request (in milliseconds), bytes sent (excluding headers), the referer, and user agent.

See Access Log valve attributes in the Tomcat 9 documentation for more information on each of the attributes.

Change the log retention

The access log is configured to keep logs for a maximum of 30 days. You can choose to increase or decrease this limit, but be aware you'll need to allow enough disk space to accommodate the log files.

To change how long access logs are kept:

- 1. Stop Confluence.
- 2. Edit the <install-directory>/conf/server.xml file.
- 3. Locate the access log valve, and change the value of maxDays.

```
<Valve className="org.apache.catalina.valves.AccessLogValve"
    directory="logs"
    maxDays="30"
    pattern="%t %{X-AUSERNAME}o %I %h %r %s %Dms %b %{Referer}i %{User-Agent}i"
    prefix="conf_access_log"
    requestAttributesEnabled="true"
    rotatable="true"
    suffix=".log"
    />
```

4. Save the file, and restart Confluence.

If you're running Confluence in a cluster, you will need to repeat this process on each node. You don't need to stop the whole cluster, you can update each node in turn.

Disable access logging

To disable access logging:

- 1. Stop Confluence.
- 2. Edit the <install-directory>/conf/server.xml file.
- 3. Remove the entire access log valve, shown here.

4. Save the file, and restart Confluence.

If you're running Confluence in a cluster, you will need to repeat this process on each node. You don't need to stop the whole cluster, you can update each node in turn.

Other access log options

For a lightweight alternative access log solution, you can also choose to enable access logging in the application logs. SeeHow to Enable User Access Loggingto find out how to do this.

This option is best suited to smaller sites, as the additional log entries may cause the application log to fill up and rotate too quickly.

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 optimization jobs to clear Confluence's temporary files and caches
- Index optimization jobs to ensure Confluence's search index is up to date
- Mail queue optimization jobs to ensure Confluence's mail queue is maintained and notifications have been sent.

• You'll needSystem Administrator permissions in order to edit and manually run jobs.

On this page:

- Accessing Confluence's scheduled jobs configuration
- Running a job manually
- Changing a job's schedule
- Disabling or re-enabling a job
- Viewing a job's execution history
- Types of jobs
- Cron expressions

Related pages:

- Trigger Module
- Configuring Backups

Accessing Confluence's scheduled jobs configuration

To access Confluence's Scheduled Jobs configuration page:

- 1. Go to Seneral Configuration > Scheduled Jobs
- 2. All scheduled jobs are listed with:
 - Status- the job's status, which is either 'Scheduled' (it it is currently enabled) or 'Disabled'.
 - Last Execution- the date and time when the job was last executed. This field will be empty of the job was never executed.
 - **Next Execution**-the date and time when the job is next scheduled to be executed. This field will contain dash symbol ('-') if the job is disabled.
 - Avg. Duration- the length of time (in milliseconds) that it took to complete the job (the last time it ran).
 - Actions- Options to edit the job's schedule, run it manually, view the history or disable the job.

Screenshot: Scheduled Jobs

Scheduled Jobs Status **Last Execution Next Execution** Avg. Duration Actions Back Up Confluence Scheduled Nov 20, 2016 02:00 Nov 19, 2016 02:00 592755 History · **Check Cluster Safety** Scheduled Nov 21, 2016 01:56 Nov 21, 2016 01:57 3 History · Clean Journal Entries Scheduled Nov 20, 2016 02:00 Nov 19, 2016 02:00 82 History Clean Old Audit Records Scheduled Nov 21, 2016 00:00 Nov 22, 2016 00:00 History · Clean Temporary Directory Scheduled Nov 20, 2016 04:00 Nov 19, 2016 04:00 199 History · Clear Expired Mail Errors Scheduled Nov 20, 2016 03:00 Nov 19, 2016 03:00 History · Clear Expired Remember Me Tokens Scheduled Nov 20, 2016 00:00 Nov 20, 2016 00:00 72 History

Running a job manually

To run a job manually head to the Scheduled Jobs list and choose **Run** next to the job. It will runimmediately.

Not all jobs can be run manually.

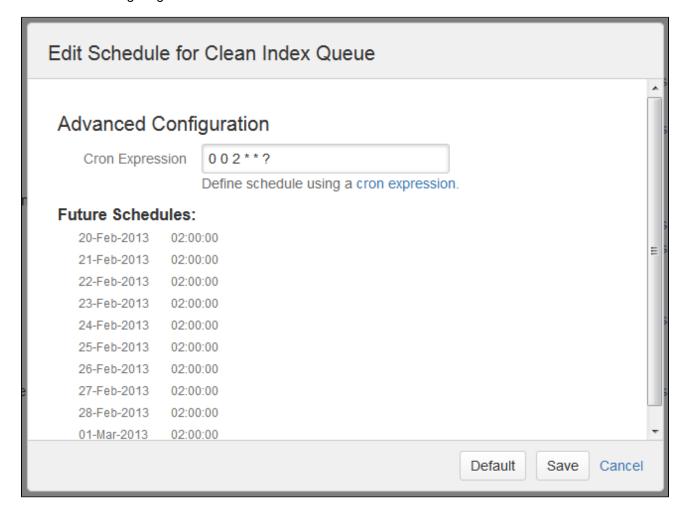
Changing a job's schedule

To change a job's schedule:

- 1. Choose **Edit** next to the job you want to change.
- Enter the new day or time to run the job as a cron expression there's more info about cron expressionsbelow.
- 3. Save your changes to the job's schedule, or Revert back to the default setting.

Not all jobs' schedules are configurable.

Screenshot: Configuring a Job Schedule



Disabling or re-enabling a job

By default, all jobs in Confluence are enabled.

Use the Disable / Enable links in the action column to disable and re-enable each job.

Not all jobs in Confluence can be disabled.

Viewing a job's execution history

To see when a job was last run, and how long the job took to run, click the **History** link beside the job.

If a job has not run at least once the History link won't appear.

Screenshot: Job Execution History

History for Clean Temporary Directory			
Started	Ended	Duration (ms)	A
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	

[♠]Execution history is not available in Confluence Data Center.

Types of jobs

Here's a summary of some of the scheduled jobs that you may want to adjust.

Job Name	Description	Execution Behavior	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.		Every 30 seconds
Clean Journal Entries	Periodically clears journal entries that have already been processed to ensure that its size does not grow indefinitely.	Per cluster	At 2am every day
Clean Temporary Directory	Cleans up temporary files generated in the <confluence-home>/temp directory. This temp directory is created by exports etc. 1 This doesn't include the temp directory located in the confluence install directory.</confluence-home>	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 while keeping the job's frequency to 24 hours.	Per cluster	At 12am every day
Flush Change Index Queue	Flushes the Change Index Queue so Confluence's search results stay up to date.	Per node	Every minute

Flush Content Index Queue	Flushes the Content Index Queue so Confluence's search results stay up to date.	Per node	Every minute
Flush Edge Index Queue	Flushes the Edge Index Queue so Confluence's search results stay up to date.	Per node	Every 30 seconds
Flush Local Task Queue	Flushes the local task queue. (These are internal Confluence tasks that are typically flushed at a high frequency.)	Per node	Every minute
Flush Mail Queue	Sends notifications that have been queued up in the mail queue. This doesn't include batched notifications. Edit the Send batched notifications job if you also want to change how often notifications are sent for changes to a page or blog post.	Per node	Every minute
Send batched notifications	Sends email notifications containing all changes to a page or blog post since the last time the job ran. Increase the time for fewer emails or reduce the time if more immediate notifications are important in your site.	Per cluster	Every 10 minutes
Flush Task Queue	Flushes the task queue. (These are internal Confluence tasks that are typically flushed at a high frequency.)	Per node	Every minute
Send Recomme nded Updates Email	Triggers sending recommended update emails to users. The job runs hourly, but users will receive the notification weekly or daily, depending on the setting in their profile, at a time that matches their timezone.	Per cluster	Hourly
Purge Old Job Run Details	Confluence stores the details of each scheduled job that is run in the scheduler_run_details table in your database. In order to keep this table small for troubleshooting and debugging, the Purge Old Job Run Details job regularly removes the details of:	Per cluster	at 11pm every day
	 successful jobs run more than 90 days ago unsuccessful jobs run more than 7 days ago 		
	You can override these settings using the following system properties; jobs.limit.per.purge, all.jobs.ttl.hours and unsuccessful.jobs.ttl.hours.		
Property Entry Gardening	When a page is created from a blueprint, some data is left behind in the os_property table after the page is published. This job cleans up leftover data, which could contain personally identifiable information.	Per cluster	At 12am every day
Clean up unpublishe d Blueprint Page Entities	This job cleans up metadata stored about draft pages created from blueprints, which could contain personally identifiable information.	Per cluster	At 2:23am every day
Synchrony data eviction (soft)	Evicts Synchrony data for any content that has not been modified in the last 3 days, and does not have an active editor session. See How to remove Synchrony data for more information.	Per cluster	Every 10 minutes

Synchrony data eviction (hard)	Evicts Synchrony data for any content that is 15 days or older, regardlessof whether it has been modified more recently. See Ho w to remove Synchrony data for more information.	Per cluster	Disabled by default
Versions Removal (Soft)	Deletes any historical page or attachment versions that don't meet the retention rules. Deletion happens in batches to minimise performance impact. After changing a retention rule, the job may need to run multiple times before all historical versions are removed.	Per cluster	Every 10 minutes
Versions Removal (Hard)	Deletes any historical page or attachment versions that don't meet the retention rules in one go, rather than in batches. This job can be run manually when required, but may have an impact on your site performance.	Per cluster	Disabled by default

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 is 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 Allowlist

Confluence administrators can choose to allow incoming and outgoing connections and content from specified sources for use in the:

- RSS Feed Macro
- HTML Include macro(disabled by default)
- gadgets
- Shared Links Blueprint
- Widget Connector Macro

by adding URLs to the allowlist.

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 allowlist.

Application links are automatically added to the allowlist. You don't need to manually add them.

Add allowed URLs to theallowlist

To add a URL to the allowlist:

- 1. Go to General Configuration > Allowlist.
- 2. Enter the URL or expression you want to allow.
- 3. Choose the **Type** of expression (see below for examples of the types available).
- 4. Choose Allow Incoming if you need to allow CORS requests (see below).
- 5. Choose Allow anonymous users if you need to allow unauthenticated users.
- 6. Choose Add.

Your URL or expression appears in the allowlist.

To test that your allowlistedURL 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 allowlist, you can choose from a number of expression types.

When deciding the best expression type to use, aimfor a more restrictive URL, rather than less restrictive, to best protect your site.

Туре	Description	Example
Domain name	Allows all URLs from the specified domain.	https://www.example.
Exact match	Allows only the specified URL.	https://www.example.com/thispage
Wildcard Expression	Allows all matching URLs. Use the wildcard * character to replace one or more characters.	https://*example.com
Regular Expression	Allows all URLs matching the regular expression.	http(s)?://www\. example\.com

Allow Incoming

Allow Incoming enables CORS requests from the specified origin. The URL must match the formatscheme: // host[:port], with no trailing slashes (:portis optional). Sohttp://example.com/would not allowCORS requests from the domain example.com.

Allow anonymous users

You can use the Allow anonymous users option to allow outbound requests on behalf of unauthenticated users.

This isn't recommended for URLs that may contain private data, such as URLs from application links. If you do need to provide anonymous access, consider using an exact URL or wildcard based rule to limit access to just the required resources.

Change default settings for application links

When you create an application link, the URL is automatically added to the Confluence allowlist. By default, outbound requests from these URLs is only allowed for authenticated users.

To change the default behaviour for all application links, including new application links:

- 1. Go to O-> General Configuration > Allowlist.
- 2. Select Configure Settings.
- 3. Select either:
 - Allow all users to allow outbound requests for all users, including anonymous users
 - Allow authenticated users to deny outbound requests for anonymous users
 - Restrict by default to deny outbound requests for all users (the applink will not be added to the allowlist at all)
- 4. Save your changes.

All existing application links, and any new application links added to the allowlist, will use this setting.

Disable theallowlist

The allowlistis enabled by default. You can choose to disable the allowlisthowever this will allow all URLs. including malicious content.



We strongly discourage you from disabling the allowlist, as it will leave you vulnerable to Server-Side Request Forgery (SSRF) attacks, such as the one disclosed in

© CONFSERVER-61399 PUBLISHED

To disable the allowlist:

- 1. Go to Seneral Configuration Allowlist.
- 2. ChooseTurn off allowlist.
- 3. ChooseConfirm.

All URLs will now be allowed. This is not recommended.

Configuring the Time Interval at which Drafts are Saved

Confluence saves a draft of your page once every thirty seconds by default. Confluence administrators can configure how often drafts are saved.

To set the time interval at which drafts are saved:

- 1. Choose the cog icon , then choose General Configuration
- ClickFurther Configuration in the left-hand panel.
- 3. Edit the setting for Draft Save Interval.

This setting applies regardless of whether collaborative editing is on or off. When collaborative editing is on, every keystroke is also aved by Synchrony in real time.

Related pages:

Drafts

Configuring Confluence Security

This section gives guidelines on configuring the security of your Confluence site:

- Confluence Security Overview and Advisories
- Proxy and HTTPS setup for Confluence
- Configuring Secure Administrator Sessions
- Confluence Cookies
- Using Fail2Ban to limit login attempts
- Securing Confluence with Apache
- 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
- Configuring RSS Feeds
- Preventing and Cleaning Up Spam

Related pages:

- Permissions and restrictions
- 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, since version 3.5 of Confluence passwords are hashed through the salted PKCS5S2 implementation provided by Embedded Crowdbefore 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 linkis generated and mailed to the user's registered address.

When external user management is enabled, password storage is delegated to the external system.

On this page:

- Application Security Overview
- Finding and Reporting a Security Vulnerability
- Publication of Confluence Security Advisories
- Severity Levels
- Our Security Bugfix Policy
- Published Security Advisories

Buffer Overflows

Confluence is a 100% pure Java application with no native components. As such it is highly resistant to buffer overflow vulnerabilities possible buffer overruns are limited to those that are bugs in the Java Runtime Environment itself.

SQL Injection

Confluence interacts with the database through the Hibernate Object-Relational mapper. Database queries are generated using standard APIs for parameter replacement rather than string concatenation. As such, Confluence is highly resistant to SQL injection attacks.

Script Injection

Confluence is a self-contained Java application and does not launch external processes. As such, it is highly resistant to script injection attacks.

Cross-Site Scripting

As a content-management system that allows user-generated content to be posted on the web, precautions have been taken within the application to prevent cross-site scripting attacks:

- The wiki markup language in Confluence does not support dangerous HTML markup
- Macros allowing the insertion of raw HTML are disabled by default
- HTML uploaded as a file attachment is served with a content-type requesting the file be downloaded, rather than being displayed inline
- Only system administrators can make HTML-level customizations of the application

When cross-site scripting vulnerabilities are found in the Confluence web application, we endeavor to fix them as quickly as possible.

Transport Layer Security

Confluence does not directly support SSL/TLS. Administrators who are concerned about transport-layer security should set up SSL/TLS at the level of the Java web application server, or the HTTP proxy in front of the Confluence application.

For more information on configuring Confluence for SSL, see: Running Confluence Over SSL or HTTPS

Session Management

Confluence delegates session management to the Java application server in which it is deployed. We are not aware of any viable session-hijacking attacks against the Tomcat application server shipped with Confluence. If you are deploying Confluence in some other application server, you should ensure that it is not vulnerable to session hijacking.

Plugin Security

Administrators install third party plugins **at their own risk**. Plugins run in the same virtual machine as the Confluence server, and have access to the Java runtime environment, and the Confluence server API.

Administrators should always be aware of the source of the plugins they are installing, and whether they trust those plugins.

Administrator Trust Model

Confluence is written under the assumption that anyone given System Administrator privileges is trusted. System administrators are able, either directly or by installing plugins, to perform any operation that the Confluence application is capable of.

As with any application, you should not run Confluence as the root/Administrator user. If you want Confluence to listen on a privileged network port, you should set up port forwarding or proxying rather than run Confluence with additional privileges. The extra-careful may consider running Confluence inside a chroot jail.

Stack Traces

To help when debugging a problem, 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.

This includes information such as operating system and version and Java version. With proper network security, this is not enough information to be considered dangerous. The username of the current user may be included.

Thread dumps include usernames and URLs by default. If you don't want to include this additional diagnostic information, you can disable Thread diagnostics.

Finding and Reporting a Security Vulnerability

Atlassian's approach to reporting security vulnerabilities is detailed in How to Report a Security Issue.

Publication of Confluence Security Advisories

Atlassian's approach to releasing security advisories is detailed in Security Advisory Publishing Policy.

Severity Levels

Atlassian's approach to ranking security issues is detailed in Severity Levels for Security Issues.

Our Security Bugfix Policy

Our approach to releasing patches for security issues is detailed in ourSecurity Bugfix Policy.

Published Security Advisories

Proxy and HTTPS setup for Confluence

Many customers choose to run Confluence behind a reverse proxy, often with HTTPS enabled. Getting your proxy configuration right is essential, to avoid problems later when using Confluence.

Proxy and HTTPS access are both configured in Tomcat, Confluence's application server.

Sample connectors

To make setup and configuration as straightforward as possible, we've provided a number of sample connectors in the Tomcat<install-directory>/conf/server.xmlfile.

Sample connector	Description
DEFAULT - Direct connector with no proxy, for unproxied HTTP access to Confluence	This is the default option. Use this option when you don't have a reverse proxy and are <i>not</i> enabling HTTPS.
HTTP - Proxying Confluence via Apache or Nginx over HTTP	Choose this option if you have a reverse proxy, but are <i>not</i> enabling HTTPS.
HTTPS - Direct connector with no proxy, for unproxied HTTPS access to Confluence.	Choose this option if you want to use HTTPS without a reverse proxy.HTTPS will be terminated at Tomcat.
HTTPS - Proxying Confluence via Apache or Nginx over HTTPS	Use this option when you want to use a reverse proxy and enable HTTPS. This is the most common configuration.

We only provide HTTP/HTTPS connector examples. You can't use the AJP connector (for example, with Apache mod_jk), as Synchrony, which is required for collaborative editing, can't accept AJP connections.

If you plan to use collaborative editing, there are a number of proxy and SSL considerations you'll need to take into account when deciding the best way to configure your proxy.

Step-by-step guides

In addition to the sample connectors, we also provide a number of step-by-step guides to help you enable HTTPS and configure your proxy correctly.

HTTPS:

- Running Confluence Over SSL or HTTPS(terminating HTTPS at Tomcat)
- Running Confluence behind NGINX with SSL(terminating HTTPS at your proxy)
- Securing your Atlassian applications with Apache using SSL(terminating HTTPS at your proxy)

Reverse proxy:

- Using Apache with mod_proxy(Confluence)
- Running Confluence behind NGINX with SSL(Confluence)
- Proxying Atlassian server applications with Apache HTTP Server (mod_proxy_http)(any Atlassian product)
- Proxying Atlassian server applications with Microsoft Internet Information Services (IIS)(any Atlassian product)

Outbound proxy:

- Configuring Web Proxy Support for Confluence (Confluence)
- How to Configure Outbound HTTP and HTTPS Proxy for your Atlassian application(any Atlassian product)



Although we provide guides for some third-party solutions, and mention Apache and Nginx in the serve r.xml file, you can choose your own proxy solution.

Atlassian Supportcan't provide assistance with configuring third-party tools like NGINX, Apache, or IIS. If you have questions, check your proxy server's documentation, ask the Atlassian Community, or get help from a Solution Partner.

Configuring Web Proxy Support for Confluence



The content on this page relates to platforms which are not supported. Consequently, Atlassian Support cannot quarantee providing any support for it. Please be aware that this material is provided for your information only and using it is done so at your own risk.

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.

- 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
```

If you're using Confluence 6.0 or later with Synchrony, you'll need to pass the following to ensure Confluence can connect directly to Synchrony. Replace localhost | 127.0.0.1 with your Synchrony IP if you have used the synchrony.host system property to change the IP Synchrony uses.

```
-Dhttp.nonProxyHosts=localhost | 127.0.0.1
-Dhttps.nonProxyHosts=localhost | 127.0.0.1
```

Note: You may need to escape the pipe character | in some command-line environments.

If thehttp.nonProxyHostsproperty 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 approriate 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.proxyUser=yourUserName
http.proxyUser=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 authenticatoravailable.

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 CONFSERVER-16941 CLOSED

Connecting to LDAP or Jira applications or Other Services via SSL



This page documents configuration of SSL, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with SSL. Please be aware that this material is provided for your information only, and that you use it at vour own risk.

Related pages:

- Configuring an SSL Connection to Active **Directory**
- Configuring Web Proxy Support for Confluence
- Running Confluence Over SSL or HTTPS

This page describes how to get Confluence connecting to external servers over SSL, via the various SSLwrapped 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, seeConfiguring an SSL Connection to Active Directory.
- You want to set up your Jira application as atrusted applicationin Confluence, when Jira is running
- You want to refer to an https://... URL in a Confluence macro.

If you want to run Confluence itself over SSL, seeRunning Confluence Over SSL or HTTPS.



There's a Confluence SSL plugin that facilitates this process.

Importing SSL Certificates

For more information on these commands, see the Keytool documentation.

1. Add the root certificate to your default Java keystore with the following command. This is the certificate that was used to authorize 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 -importcert -alias serverCert -file RootCert.crt -keystore %JAVA_HOME%/jre/lib/security
/cacerts (Windows)
keytool -importcert -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 -importcert -alias ldapCert -file JIRAorLDAPServer.crt -keystore %JAVA_HOME%/jre/lib
/security/cacerts (Windows)
keytool -importcert -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 CATALINA_OPTS to specify the path to the keystore, as specified in C onnecting to SSL services before restarting 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

Using Apache with mod proxy



Atlassian applications allow the use of reverse-proxies, however Atlassian Support does not provide assistance for configuring them. Consequently, Atlassiancan not guarantee providing any support for them.

If assistance with configuration is required, please raise a question on Atlassian Answers.

This page describes one possible way to use Apache HTTP Server 2.4 to proxy requests for Confluence running in a standard Tomcat container. You can find additional documentation that explains how to useNGINXfor the same purpose.

You might use this configuration when:

- You have an existing Apache website, and want to add Confluence (for example, http://w ww.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/confluenceandhttp://example:8080 /jiraand want to make them both available on the regular HTTP port (80) (for example, athtt p://www.example.com/confluenceandhttp:/ /www.example.com/jira). Each application can 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

http://example:8091 - the hostname and port Synchrony, the service that powers collaborative editing, defaults to

/confluence - the intended context path for Confluence (the part after hostname and port)

/synchrony - the context path for Synchrony, the process that powers collaborative editing

You'll need to replace these URLs with your own URLs.

1 Set the context path

Alf you want to access Confluence without a context path, such aswww.example.com, skip this step.

Set your Confluence application path (the part after hostname and port) in Tomcat. In this example the context path will be/confluence.

On this page:

- Base configuration
 - 1 Set the context path
 - 2 Set the URL for redirection
 - 3 Configure mod_proxy
 - 4 Restart Apache
 - 5 Disable HTTP Compression
 - 6 Change the Confluence Base URL
- Adding SSL
- More information

Edit <installation-directory>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">
```

In this example we've used/confluenceas the context path. Note that you can't use/resourcesas your context path, as this is used by Confluence, and will cause problems later on.

Restart Confluence, and check you can access it at http://example:8090/confluence.

2 Set the URL for redirection

Next, set the URL for redirection. In the same<installation-directory>conf/server.xmlfile, use the example connectors as a starting point.

Comment out the default connector (for unproxied access).

In XML a comment starts with<! -- and ends with-->, and is used to make sure only the relevant portions of the file are read by the application.

Add <!-- and--> around the **default** connector. It should now look like this.

Uncomment the connector listed under the HTTP - Proxying Confluence via Apache or Nginx over HTTPh eading.

To uncomment a section, remove the<!-- and--> surrounding the connector.

Here's an example showing the default connector commented out, and the HTTP connector uncommented. The headings remain commented out.

```
<!--
______
DEFAULT - Direct connector with no proxy, for unproxied HTTP access to Confluence.
_____
<!--
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
  maxThreads="48" minSpareThreads="10"
  enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
 protocol="org.apache.coyote.http11.Http11NioProtocol"/>
<!--
HTTP - Proxying Confluence via Apache or Nginx over HTTP
_____
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"</pre>
  maxThreads="48" minSpareThreads="10"
  enableLookups="false" acceptCount="10" debug="0"URIEncoding="UTF-8"
  protocol="org.apache.coyote.http11.Http11NioProtocol"
  scheme="http" proxyName="<subdomain>.<domain>.com" proxyPort="80"/>
```

Insert your **proxyName**and**proxyPort**as shown in the last line below:

```
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
  maxThreads="48" minSpareThreads="10"
  enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
  protocol="org.apache.coyote.http11.Http11NioProtocol"
  scheme="http" proxyName="www.example.com" proxyPort="80"/>
```

If you plan to enable HTTPS, use the connector underHTTPS - Proxying Confluence via Apache or Nginx over HTTPS.

3 Configure mod_proxy

Use one of the examples below to edit your Apache http.conffile to proxy requests to the application server.

You will need to enable the following required Apache modules if they are not already enabled:

- mod proxy
- mod_proxy_http
- proxy_wstunnel
- mod_rewrite

(proxy_wstunnel and mod_rewrite are new requirements in Confluence 6.0)

The format of the http.conffile, and location of the modules may differ on your operating system. We recommend Windows users specify the absolute path to the module files.

Example 1: Configuration with context path

Use this example if you set a context path in step 1, and will access Confluence with a context path like thisht tp://www.example.com/confluence.

In this example, users will connect to Synchrony, which is required for collaborative editing, directly via WebSockets.

▲The order of directives in the config is important.

Apache HTTP server 2.4 # Put this after the other LoadModule directives LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so LoadModule proxy_wstunnel_module /usr/lib/apache2/modules/mod_proxy_wstunnel.so LoadModule rewrite_module /usr/lib/apache2/modules/mod_rewrite.so # Put this in the main section of your configuration (or virtual host, if using Apache virtual hosts) ProxyRequests Off ProxyPreserveHost On <Proxy *> Require all granted </Proxy> ProxyPass /synchrony http://<domain>:8091/synchrony <Location /synchrony> Require all granted RewriteEngine on RewriteCond %{HTTP:UPGRADE} ^WebSocket\$ [NC] RewriteCond %{HTTP:CONNECTION} Upgrade\$ [NC] RewriteRule .* ws://<domain>:8091%{REQUEST_URI} [P] </Location> ProxyPass /confluence http://<domain>:8090/confluence ProxyPassReverse /confluence http://<domain>:8090/confluence <Location /confluence> Require all granted </Location>

Note: It's not possible to use Apache HTTP Server 2.2 with Confluence 6.0 or later. If you plan to use SSL, you will need version 2.4.10 or later.

Example 2: Configuration without context path

Use this example if you skipped step 1, and will access Confluence without a context path like this http://www.example.com.

As in the previous example, users will connect to Synchrony, which is required for collaborative editing, directly via WebSockets.

The order of directives in the config is important.

```
Apache HTTP server 2.4
# Put this after the other LoadModule directives
LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so
LoadModule proxy_wstunnel_module /usr/lib/apache2/modules/mod_proxy_wstunnel.so
LoadModule rewrite_module /usr/lib/apache2/modules/mod_rewrite.so
# Put this in the main section of your configuration (or virtual host, if using Apache virtual hosts)
  ProxyRequests Off
  ProxyPreserveHost On
 RewriteEngine On
  RewriteCond %{REQUEST_URI} !^/synchrony
  RewriteRule ^/(.*) http://<domain>:8090/$1 [P]
  <Proxy *>
     Require all granted
  </Proxy>
  ProxyPass /synchrony http://<domain>:8091/synchrony
  <Location /synchrony>
     Require all granted
      RewriteEngine on
     RewriteCond %{HTTP:UPGRADE} ^WebSocket$ [NC]
     RewriteCond %{HTTP:CONNECTION} Upgrade$ [NC]
     RewriteRule .* ws://<domain>:8091%{REQUEST_URI} [P]
  ProxyPass / http://<domain>:8090
  ProxyPassReverse / http://<domain>:8090
  <Location />
     Require all granted
  </Location>
```

Note: It's not possible to use Apache HTTP Server 2.2 with Confluence 6.0 or later. If you plan to use SSL, you will need version 2.4.10 or later.

4 Restart Apache

This is needed to pick up on the new configuration. To restart Apache, run the following command:

```
sudo apachectl graceful
```

5 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 ourCompressing an HTTP Response within Confluencedocs.

6 Change the Confluence Base URL

The last stage is to set the Base URL to the address you're using within the proxy, for example http://www.example.com/confluence.

Adding SSL

If you plan to enable HTTPS, seeSecuring your Atlassian applications with Apache using SSL, andmake sure you choose the HTTPS sample connector.

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.

Running Confluence behind NGINX with SSL

This page describes how to set up NGINXas arevers e proxyfor Confluence.

The configuration described on this page results in a scenario where:

- External clientconnections with NGINX are secured using SSL. Connections between NGINXand Confluence Server are unsecured.
- Confluence Server and NGINXrun on the same machine.

On this page

- Step 1: Set the context path
- Step 2: Configure the Tomcat connector
- Step 3: Configure NGINX
- Step 4: RestartConfluence and NGINX

We assume that you already have a running instance of NGINX. If not, refer to the NGINX documentation for instructions on downloading and installing NGINX.SSL certificates must be installed on the server machine. You'll an NGINX version that supports WebSockets (1.3 or later).

If your team plans to use the Confluence Server mobile app, you'll need a certificate issued by a trusted Certificate Authority. You can't use the app with a self-signed certificate, or one from an untrusted or private CA.



Atlassian Supportcan't provide assistance with configuring third-party tools like NGINX. If you have questions, check the NGINX documentation, ask the Atlassian Community, or get help from a Solutio n Partner

Step 1: Set the context path

▲If you want to access Confluence without a context path (www.example.com),or via a sub-domain (confluence.example.com)skip this step.

Set your Confluence application path (the part after hostname and port) in Tomcat. Edit <installationdirectory>/conf/server.xml, locate the "Context" definition:

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

and change it to:

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

In this example we've used /confluenceas the context path. Note that you can't use /resourcesas your context path, as this is used by Confluence, and will cause problems later on.

Restart Confluence, and check you can access it athttp://example:8090/confluence

Step 2: Configure the Tomcat connector

In the same<installation-directory>conf/server.xmlfile, use the example connectorsas a starting point.

Comment out the default connector (for unproxied access).

In XML a comment starts with<! -- and ends with-->, and is used to make sure only the relevant portions of the file are read by the application.

Add <!-- and--> around the **default** connector. It should now look like this.

Uncomment the connector listed under the HTTPS - Proxying Confluence via Apache or Nginx over HTTPSheading.

To uncomment a section, remove the <! -- and --> surrounding the connector.

Here's an example showing the default connector commented out, and the HTTPS connector uncommented. The headings remain commented out.

```
<!--
______
DEFAULT - Direct connector with no proxy, for unproxied HTTP access to Confluence.
______
<!--
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
 maxThreads="48" minSpareThreads="10"
  enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
 protocol="org.apache.coyote.http11.Http11NioProtocol"/>
<!--
______
HTTPS - Proxying Confluence via Apache or Nginx over HTTPS
-----
-->
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
  maxThreads="48" minSpareThreads="10"
  enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
  protocol="org.apache.coyote.http11.Http11NioProtocol"
  scheme="https" secure="true" proxyName="<subdomain>.com" proxyPort="443"/>
```

Insert your proxyName and proxyPort as shown in the last line below:

```
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
  maxThreads="48" minSpareThreads="10"
  enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
  protocol="org.apache.coyote.http11.Http11NioProtocol"
  scheme="https" secure="true" proxyName="www.example.com" proxyPort="443"/>
```

Make sure you've included correct values for protocol and proxyName.

Step 3: Configure NGINX

You will need to specify a listening server in NGINX, as in the example below. Add the following to your NGINX configuration.

Replace your server name and the location of your SSL certificate and key.

In this example, users will connect to Synchrony, which is required for collaborative editing, directly.

```
server {
   listen www.example.com:80;
   server name www.example.com;
   listen 443 default ssl;
    ssl_certificate /usr/local/etc/nginx/ssl/nginx.crt;
    ssl_certificate_key /usr/local/etc/nginx/ssl/nginx.key;
    ssl session timeout 5m;
    ssl_protocols TLSv1 TLSv1.1 TLSv1.2;
    ssl ciphers 'ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-
POLY1305:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:
ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-
AES128-SHA256:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA:ECDHE-RSA-
AES256-SHA384:ECDHE-RSA-AES128-SHA:ECDHE-ECDSA-AES256-SHA384:ECDHE-
ECDSA-AES256-SHA: ECDHE-RSA-AES256-SHA: ECDHE-ECDSA-DES-CBC3-SHA: ECDHE-
RSA-DES-CBC3-SHA: EDH-RSA-DES-CBC3-SHA: AES128-GCM-SHA256: AES256-GCM-
SHA384:AES128-SHA256:AES256-SHA256:AES128-SHA:AES256-SHA:DES-CBC3-
SHA: !DSS';
    ssl_prefer_server_ciphers
location /synchrony {
       proxy set header X-Forwarded-Host $host;
       proxy_set_header X-Forwarded-Server $host;
       proxy set header X-Forwarded-For $proxy add x forwarded for;
        proxy_pass http://localhost:8091/synchrony;
       proxy http version 1.1;
        proxy_set_header Upgrade $http_upgrade;
        proxy_set_header Connection "Upgrade";
    }
  location /confluence {
       client_max_body_size 100m;
        proxy_set_header X-Forwarded-Host $host;
        proxy_set_header X-Forwarded-Server $host;
       proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
       proxy_pass http://localhost:8090/confluence;
    }
}
```

See https://nginx.org/en/docs/http/ngx_http_proxy_module.htmlfor more information.

Note: do not includesslon; if you are configuring SSL and Confluence on the same server, as in this example.

If you're not sure what to include forssl_ciphers,https://mozilla.github.io/server-side-tls/ssl-configgenerator/is a useful resource.

If you experience 413Request Entity Too Large errors, make sure that the client_max_body_size in the /confluencelocation block matches Confluence's maximum attachment size. You may also need to increase the client_max_body_size in the /synchronylocation block if you experience errors when editing large pages.

If you plan to allow users to use the Confluence mobile app with your site, and you have configured a context path, as in the example above, you may also need to add the following line to your nginx configuration.

```
location /server-info.action {
    proxy_pass http://localhost:8090/confluence/server-info.action;
}
```

If you're accessing Confluence via a sub-domain, your config will look like this:

```
server {
   listen confluence.example.com:80;
    server_name confluence.example.com;
    listen 443 default ssl;
                      /usr/local/etc/nginx/ssl/nginx.crt;
    ssl_certificate
    ssl_certificate_key /usr/local/etc/nginx/ssl/nginx.key;
    ssl_session_timeout 5m;
    ssl_protocols TLSv1 TLSv1.1 TLSv1.2;
    ssl ciphers 'ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-
POLY1305:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:
ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-
AES128-SHA256:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA:ECDHE-RSA-
AES256-SHA384:ECDHE-RSA-AES128-SHA:ECDHE-ECDSA-AES256-SHA384:ECDHE-
ECDSA-AES256-SHA: ECDHE-RSA-AES256-SHA: ECDHE-ECDSA-DES-CBC3-SHA: ECDHE-
RSA-DES-CBC3-SHA: EDH-RSA-DES-CBC3-SHA: AES128-GCM-SHA256: AES256-GCM-
SHA384:AES128-SHA256:AES256-SHA256:AES128-SHA:AES256-SHA:DES-CBC3-
SHA: !DSS';
    ssl_prefer_server_ciphers
    location / {
       client max body size 100m;
       proxy_set_header X-Forwarded-Host $host;
       proxy_set_header X-Forwarded-Server $host;
       proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
       proxy_pass http://localhost:8090;
    location /synchrony {
       proxy_set_header X-Forwarded-Host $host;
       proxy_set_header X-Forwarded-Server $host;
       proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
       proxy_pass http://localhost:8091/synchrony;
       proxy_http_version 1.1;
       proxy_set_header Upgrade $http_upgrade;
       proxy_set_header Connection "Upgrade";
}
```

Step 4: RestartConfluence and NGINX

- 1. Restart Confluence and NGINX for all the changes to take affect.
- Update Confluence's base URL to include the context path you set earlier seeConfiguring the Server Base URL.

Running Confluence Over SSL or HTTPS



Atlassian applications can be accessed via HTTPS, however Atlassian Support does not provide assistance for configuring it. Consequently, Atlassian cannot guarantee providing any support for

- If assistance with conversions of certificates is required, please consult with the vendor who provided the certificate.
- If assistance with configuration is required, please raise a question on Atlassian Community.

This page provides a basic outline of how to configure Confluence to enable access via HTTPS (HTTP Secure), so that your Confluence logins and data are encrypted during transport to and from Confluence. This is a good way to safeguard your Confluence data and user logins from being intercepted and read by outsiders.

In this article we use 'SSL' as a general term to refer to the protocol used to encrypt traffic. In most cases the protocol will be TLS.

On this page:

- Step 1. Create or request an SSL certificate
- Step 2. Modify your Confluence server.xml
- 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

These instructions cover terminating SSL at Tomcat, the application server shipped with Confluence.

If you want to terminate SSL at your web server or proxy, seeApache with mod_proxyorRunning Confluence behind NGINX with SSL for examples of how to terminate SSL at an external web server.

You'll need the JDK for some of the steps in this guide. The JRE is not enough.



 Running Confluence without HTTPS enabled may leave your site exposed to vulnerabilities, such as man-in-the-middle or DNS rebinding attacks. We recommend you enable HTTPS on your site.

Step 1. Create or request an SSL certificate

You'll need a valid certificate before you can enable HTTPS. If you already have a certificate, skip to step 2.

You can create your own self-signed certificate, or acquire one from a trusted Certificate Authority.

If your team plans to use the Confluence Server mobile app, you'll need a certificate issued by a trusted Certificate Authority. You can't use the app with a self-signed certificate, or one from an untrusted or private CA.

Option 1: Create a self-signed certificate

Self-signed certificates are useful if you require encryption but don't 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. Users won't be able to log in to your site at all via the Confluence Server mobile app if you use a self-signed certificate.

Linux (and MacOS)

In this example, we'll use Java'skeytool utility, which is included with the JDK.If you're not comfortable using command line utilitiesKeyStore Exploreris a usefulalternative to the command line.

To generate a self-signed certificate using keytool:

1. From the command line, run the appropriate command for your operating system:

Windows "%JAVA_HOME%\bin\keytool" -genkeypair -keysize 2048 -alias tomcat -keyalg RSA -sigalg SHA256withRSA

```
2. When prompted, createa passwordfor the certificate (private key).
```

• Only use alphanumeric characters. Tomcat has a known issue with special characters.

\$JAVA_HOME/bin/keytool -genkeypair -keysize 2048 -alias tomcat -keyalg RSA -sigalg SHA256withRSA

- Make a note of the password, you'll need it in the next step.
- The default password is 'changeit'.
- 3. Follow the prompts to specify the certificate details. This info is used to construct the X.500 Distinguished Name (DN) of the entity.
 - First and last name: this is not your name, it is the Common Name (CN), for example 'confluence.example.com'. The CN must match the fully qualified hostname of the server running Confluence, or Tomcat won't be able to use the certificate for SSL.
 - Organizational unit: this is the team or department requesting the certificate, for example 'marketing'.
 - Organization: this is your company name, for example 'SeeSpaceEZ'.
 - City, State / province, country code: this is where you're located, for example Sydney, NSW, AU.
- 4. The output will look something like the example below. Hit 'y' to confirm the details.

```
CN=confluence.example.com, OU=Marketing, O=SeeSpaceEZ, L=Sydney, ST=NSW, C=AU
```

- 5. When asked for the password for 'tomcat', enter the password you created in step 2 (or hit return to use the same .
 - 'tomcat' is the alias we entered in the keytool command above, it refers to your.
 - Your keystore entry must have the same password as your private key. This is a Tomcat requirement.
- 6. You certificate is now ready. Go to step 2 below.

Option 2: Use a certificate issued by a Certificate Authority (recommended)

Production environments will need a certificate issued by a Certificate Authority (CA). These instructions 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 (follow the steps in option 1, above).
- 2. From the command line, run the following command to generate a certificate signing request.

```
keytool -certreq -keyalg RSA -alias tomcat -file certreq.csr -keystore <MY_KEYSTORE_FILENAME>
```

Replace<MY_KEYSTORE_FILENAME> with the path to and file name of the .keystorefile generated for your local certificate.

- 3. Submit the generated file called certreq.csr to your chosen certificate authority.

 Check your CA's documentation 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>
```

Some CAs require you to install an intermediate certificate before importing your certificate. You should follow your CA documentation to successfully install your certificate.

- If you receive an error, and you use Verisign or GoDaddy, you mayneed to export the certificate to PKCS12 format along with the private key.
 - 1. First, remove thecertificate 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 PKCS12to 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 your Confluence server.xml file

The next step is to configure Confluence to use HTTPS.

- Edit <install-directory>/conf/server.xml.
- 2. Uncomment the following lines:

```
<Connector port="8443" maxHttpHeaderSize="8192"
  maxThreads="150" minSpareThreads="25"
  protocol="org.apache.coyote.http11.Http11Nio2Protocol"
  enableLookups="false" disableUploadTimeout="true"
  acceptCount="100" scheme="https" secure="true"
  clientAuth="false" sslProtocol="TLSv1.2"
  sslEnabledProtocols="TLSv1.2" SSLEnabled="true"
  URIEncoding="UTF-8" keystorePass="<MY_CERTIFICATE_PASSWORD>"/>
```

- 3. Replace <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.
- Change the value of maxThreads to be at least 10 threads (or 25%) less than the size of your
 database connection pool. 48 is usually about right. SeeHTTP MaxThreads configuration for more
 information about this.
- 6. Save the server configuration file.

♠Don't remove or comment out thehttp connector, as the Synchrony proxy health check, still requires HTTP.If you don't want to include the httpconnector, you can use the synchrony.proxy.healthcheck. disabledsystem property to disable the health check.

You should also **not** disable the internal Synchrony proxy (by setting the synchrony.proxy.enabled syste m property to false) as this is known to cause problems when you're terminating SSL at Tomcat.

The default connector port for Confluence is 8090.

We recommend you use TLS 1.2, as the Confluence mobile app requires minimum TLS 1.2, and there are currently some known issues running Confluence with TLS 1.3.

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:\users\#CURRENT USER#\.keystore
- On OS X and UNIX-based systems: ~/.keystore



Don't store your keystore file in your Confluence installation directory as the contents of that directory are removed when you upgrade Confluence.

You may decide to move the certificate to a custom location. If your certificate is not in the default location, you'll need to update your server configuration file as outlined below, so that Tomcat can find the certificate.

- 1. Edit<confluence-install-directory>/conf/server.xml
- 2. Add the attribute keystoreFile="<MY CERTIFICATE LOCATION>" to the Connectorelement, so that the element looks like this:

```
<Connector port="8443" maxHttpHeaderSize="8192"</pre>
  maxThreads="150" minSpareThreads="25"
  protocol="org.apache.coyote.http11.Http11Nio2Protocol"
  enableLookups="false" disableUploadTimeout="true"
  acceptCount="100" scheme="https" secure="true"
  clientAuth="false" sslProtocol="TLSv1.2"
  sslEnabledProtocols="TLSv1.2" 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 configuration file.

Step 4. Change your confluence base URL to HTTPS

- 1. In your browser, go to Co Seneral Configuration.
- 2. Click Edit.
- 3. Change the Server Base URL to HTTPS. See the documentation on configuring the server base URL.
- 4. Restart Confluenceand 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:

- 4. Restart Confluence and access http://localhost:8090. You should be redirected to https://localhost: 8443/login.action.
- 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 11 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:

```
<Connector port="8443" address="your.confluence.url.com"
   maxHttpHeaderSize="8192" maxThreads="150" minSpareThreads="25"
   protocol="org.apache.coyote.http11.Http11Nio2Protocol"
   enableLookups="false" disableUploadTimeout="true"
   acceptCount="100" scheme="https" secure="true"
   clientAuth="false" sslProtocol="TLSv1.2"
   sslEnabledProtocols="TLSv1.2" SSLEnabled="true"
   URIEncoding="UTF-8" keystorePass="<MY_CERTIFICATE_PASSWORD>"
   keystoreFile="<MY_CERTIFICATE_LOCATION>"/>
```

- HTTPS must be configured for your whole site. It can't be enabled for individual pages or spaces.
- Before you upgrade Confluence, make a note of the changes you have made to your server.xmla nd web.xmlfiles. It is always best to re-apply these changes manually after upgrading, rather than copying over your existing files.
- TLS 1.2 recommended. The Confluence Server mobile app requires TLS 1.2. There are also known issues running Confluence with TLS 1.3, and connecting to other applications using TLS1.3. If you use Jira and Confluence together, we recommend configuring both applications to use TLS 1.2.

Troubleshooting

Check the Confluence knowledge base articles on troubleshooting SSL

- SSL Configuration HOW-TOin the Apache Tomcat 9.0 documentation
 keytool Key and Certificate Management Toolin the Java 11 documentation

Using Apache to limit access to the Confluence administration interface

As well as limiting access to the Confluence administration console to users who really need it, and using strong passwords, you can consider limiting access to certain machines on the network or internet. If you are using Apa che web server, this can be done with Apache's **Location** functionality.

To limit access to admin screens to specific IP addresses in Apache:

 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
```

In your Apache Virtual Host, add the following lines to restrict the administration actions to the Systems Administrator.

```
Define segmentregex (?:;[^/]*)?(?:/)?(?:(?:;[^/]*)?(?:/)?)*
<LocationMatch (?i)/confluence${segmentregex}/admin>
     Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}</pre>
/oauth${segmentregex}/consumers${segmentregex}/list>
        Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}</pre>
/oauth${segmentregex}/view-consumer-info>
        Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}</pre>
/oauth${segmentregex}/service-providers${segmentregex}/list>
        Include sysadmin_ips_only.conf
</LocationMatch>
< Location Match \ (?i)/confluence \\ \{segmentregex\}/plugins \\ \{segmentregex\}/servlet \\ \{segmen
/oauth${segmentregex}/service-providers${segmentregex}/add>
        Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}</pre>
/oauth${segmentregex}/consumers${segmentregex}/add>
        Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}</pre>
/oauth${segmentregex}/consumers${segmentregex}/add-manually>
       Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}</pre>
/oauth${segmentregex}/update-consumer-info>
        Include sysadmin ips only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/pages${segmentregex}/templates${segmentregex}</pre>
/listpagetemplates.action>
       Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/pages${segmentregex}/templates${segmentregex}</pre>
/createpagetemplate.action>
        Include sysadmin ips only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/spacepermissions.action>
       Include sysadmin_ips_only.conf
</LocationMatch>
< Location Match \ (?i)/confluence \\ \{ segmentregex \}/pages \\ \{ segmentregex \}/list permission pages. \\ action > list permission pages. \\ ac
        Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/removespace.action>
                       Include sysadmin_ips_only.conf
```

```
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/importmbox.action>
                                Include sysadmin_ips_only.conf
</LocationMatch>
< Location Match \ (?i)/confluence \\ \{segmentregex\}/spaces \\ \{segmentregex\}/view mail accounts.action > 1 \\ (?i)/confluence \\ \{segmentregex\}/spaces 
           Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/addmailaccount.action?>
           Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/importpages.action>
          Include sysadmin_ips_only.conf
</LocationMatch>
$$ \end{tabular} $$$ \e
/flyingpdf.action>
          Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/exportspacehtml.action>
           Include sysadmin ips only.conf
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/exportspacexml.action>
           Include sysadmin_ips_only.conf
</LocationMatch>
Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}/upm>
          Include sysadmin_ips_only.conf
</LocationMatch>
```

This configuration assumes you're running Confluence with the context path '/confluence'. If you are running with a different context path, or no context path, adjust the sample above accordingly.

Using Apache with mod_jk



⚠ It's not possible to use only mod_jk to proxy Confluence 6.0 or later. This is because Synchrony, which is required for collaborative editing, cannot accept AJP connections. The preferred configuration is Using Apache with mod_proxy.

If you are unable to switch to mod_proxy, see How to configure Apache mod_jk to proxy Confluence 6.x or later for a workaround.

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 Secure Administrator Sessions

Secure administrator sessions allows you to require administrators to re-enter their password before they can access administrative functions. This feature is sometimes known as "websudo" and is turned on by default.

Start a secure administrator session

When an administrator attempts to access an admin function (including some space admin functions like delete space), they will be prompted to re-enter their password. This starts the secure administrator session.

Administrators can click **Drop access** in the banner to manually end the session. This won't log them out of Confluence, it will just end the secure administrator session.



You have temporary access to administrative functions. Drop access if you no longer require it. For more information, refer to the documentation.

Change the secure administrator session timeout

The secure administrator session has a rolling timeout which defaults to 10 minutes. If there's no activity for a period of time, the administrator will be logged out of the session. They'll remain logged in to Confluence.

To change the timeout value:

- 1. Go to General Configuration > Security Configuration.
- 2. Select Edit.
- 3. Under Secure administrator sessions, enter the Minutes before automatic invalidation.
- 4. Save your changes.

Turn off secure administrator sessions

If you're using single sign-on, or have other security measures in place, you may want to disable secure administrator sessions. We don't recommend doing this unless you need to.

To turn off secure administrator sessions:

- 1. Go to O > General Configuration > Security Configuration.
- 2. Select Edit.
- 3. Under Secure administrator sessions, deselect the Enable checkbox.
- Save your changes.

Troubleshooting

Known issues with single sign-on and just-in-time user provisioning

You may need to disable secure administrator sessions if your users are not stored in Confluence's internal user directory. See

CONFSERVER-60263 - Ability to have the Websudo functionality working with SAML / SSO GATHERING INTEREST

for more information and some suggested workarounds.

Known issues for app developers

Secure administrator sessions can cause exceptions when developing against Confluence or deploying a plugin. SeeHow do I develop against Confluence with Secure Administrator Sessions?

Note that REST and XML-RPC APIs are not affected by secure administration sessions.

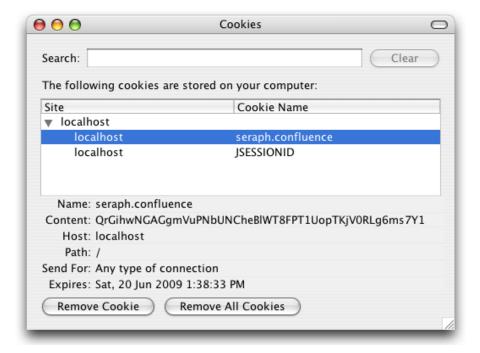
Confluence Cookies

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-partyConfluence 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. This cookie is ignored when Confluence is clustered.
- The 'remember me' cookie, seraph.confluence, is generated by Confluence when the user selects the Remember me check box on the login page. Remember me is enforced by default when Confluence is clustered.
- 1 You can read about cookies on the Wikipedia page about HTTP cookies.



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.

The default time to live of this cookie is two weeks.

When running Confluence in a cluster, Remember me is enabled by default to allow users to move seamlessly between nodes. See How to configure the 'Remember Me' feature in Confluence if you need to configure this.

Cookie key and contents

On this page:

- Authentication cookies
- The 'remember me' cookie
- Other Confluence cookies

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:

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 clearExpiredRememberMe TokensTrigger.

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
confluen ce.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.

confluen ce. 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.
confluen ce- 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. conglom erate. 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 autocompletion, which is typically set through a browser's settings.

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.

Prerequisites

- Requires Python 2.4 or higher to be installed
- Requires Apache Reverse Proxy 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.
[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

When opened in a viewport, the user will be redirected to: Proxy and HTTPS setup for Confluence.

Best Practices for Configuring Confluence Security

This page outlines a number of approaches you can use to make your Confluence site as secure as possible. 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.

On this page:	

Best practices

Not everything in this guide will be applicable to your environment, but the principles described can be adapted to most environments.

It's worth noting that none of these practices can provide 100% security. They are measures to reduce impact and to slow down an intruder in case your system does become compromised.

Subscribe to advisory alerts

Subscribe to advisory alerts and keep technical contact details up to date to make sure you receive security advisory alerts and other important technical updates.

Atlassian email and privacy preferences

Secure your installation and data directories

It's important to make sure your Confluence installation directory, home directory, and any storage locations you may define for attachments, space exports, or data pipeline exports are secure.

We strongly recommend you:

- run Confluence with a dedicated non-root user account.
- limit the user accounts who can access any Confluence directories.

To find out how to do this, see Creating a Dedicated User Account on the Operating System to Run Confluence.

You should also monitor your binaries. If an attacker compromises an account on your system, they 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. Consider how you might regularly verify that no malicious changes have been made.

Secure your database

Make sure the Confluence database user (and all datasource database users) only have the amount of database privilege they really need.

Limit database access to just the Confluence host (using iptables or built in database security tools). Refer to your database documentation to find out how to do this.

Limit access to administrator functions

As a general rule, you shouldkeep the number of Confluence administrators as low as possible, and review these user accounts every so often to make sure the access level is still appropriate.

- Avoid shared administrator or user accounts, and easily guessed usernames like 'admin'.
- Provide administrators with two separate accounts, to allow them to separate day-to-day work such as creating pages, from administration tasks.
- Limit the number of people in the confluence-administrators group. Members of this 'super group' can access all admin functions and access all content, including restricted pages. Consider limiting the

members of this group and instead create a new group with system administrator global permissions. Learn about the confluence-administrators super group

- Use secure administrator sessions to require admins to re-enter their password to access admin functions, and set a short timeout for the administrator session.
 Learn how to turn on secure administrator sessions
- Use Apache to lock down the administration interface to specific IP addresses. This can be used as a template for your reverse proxy of choice.
 Using Apache to limit access to the Confluence administration interface.

Limit incoming and outgoing connections

There are a number of ways you can limit incoming and outgoing connections, including using firewalls and proxy servers.

 Use the allowlist tolimit incoming and outgoing connections to avoid Server-Side Request Forgery (SSRF) attacks. Confluence relies on the allowlist for things like macros, that may be displaying content from external sites.
 Learn how to turn on the allowlist

Manage user accounts

Good user management practices can help prevent user accounts being compromised.

- Consider integrating Confluence with an identity provider for single sign-on and two-factor authentication.
 - Learn about the various SSO options available
- Use personal access tokens for integrations. This provides your users a more secure way to authenticate API requests than basic authentication (username and password) Learn how to manage personal access tokens
- Disable basic authentication. If you've configured single sign-on, you can disable basic authentication
 for login and REST requests. Basic authentication is less secure than single sign-on and personal
 access tokens. Learn how to disable basic authentication
- Disable user accounts when people leave your organisation. If required, you can also delete user accounts which will replace their name with an anonymized ID.
 Delete and disable users
- Restrict the number of users with powerful roles or group memberships. If only one department should have access to particularly sensitive data, then restrict access to the data to only those users. Do not let convenience over-rule security. Do not give all staff access to sensitive data when there is no need.

Limit login attempts and monitor access and activity

There are several things you can do to reduce the risk of brute force or denial of service attacks.

- Fail2Ban can help you reduce the risk of brute force attacks.
 Using Fail2Ban to limit login attempts
- Use rate limiting to block REST API requests from anonymous users if you don't have a reason to allow them, or limit the number of requests to reduce the risk of DoS attacks Learn how to use rate limiting to block requests
- Review your audit log settings, to make sure you're logging important administrator and end user actions.
 - Learn which events you can write to the audit log
- Access logs can help you identify unusual activity. Logs are written to the install directory, and you
 may want to monitor these logs using your preferred monitoring tool.

 Learn about access logging

Perform regular security audits

Regular security audits can help you identify potential threats, and also provide an opportunity to review your security policies and procedures.

- Know who can help if a security breach occurs. What is the process if a potential threat is identified?
- Perform 'what if' planning exercises. Consider questions like 'What is the worst thing that could happen if a privileged user's password were stolen while on vacation? What can we do to minimize damage?'.
- Document your security measures, and regularly monitor that all measures are still in place, and are adequate. For example after upgrading or migrating, someone may forget to apply the rule to the new system or version.
- Perform a security check-up when preparing for a major upgrade. It's a good time to check your current configuration against our current recommendations.

Hiding the People Directory

The People Directory provides a list of all users in your Confluence system.

If you need to disable the People Directoryset the following system properties on your application server command line:

• To disable the People Directory foranonymous users:

-Dconfluence.disable.peopledirectory.anonymous=true

• To disable the People Directory entirely:

-Dconfluence.disable.peopledirectory.all=true

This workaround will prevent the People directory from appearing on the dashboard, but if you navigate to the profile of a user, and then click on the "People" in the breadcrumb link (Dashboard >> People >> FullName >> Profile) or you go to the URL directly <CONFLUENCE_INSTALL>/browsepeople.action, you will be able to access the people directory.

To workaround this, set up your Apache webserver in front of Confluence and redirect requests to this URL.

Configuring Captcha for Spam Prevention

If your Confluence site is open to the public (you allow anonymous users to add comments, create pages etc) you may find that automated spam is being added, in the form of comments or new pages.

You can configure Confluence to deter automated spam by asking users to prove that they are human before they are allowed to:

- · Sign up for an account.
- Add a comment.
- Create a page.
- · Edit a page.
- Send a request to the Confluence administrators.

Related pages:

Configuring Confluence Security

Captcha is a test that can distinguish a human being from an automated agent such as a web spider or robot. When Captcha is switched on, users will see a distorted picture of a word, and must enter it in a text field before they can proceed.

Screenshot: Example of a Captcha test



By default, Captcha is disabled. When enabled, the default is that only anonymous users will have to perform the Captcha test when creating comments or editing pages. You can also choose to enforce Captcha for all users or members of particular groups.

You need System Administrator permissions to configure Captcha for spam prevention in Confluence.

To enable Captcha for spam prevention in Confluence:

- 1. Choose the cog icon , then choose General Configuration
- 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.
 - SelectSigned in usersif you want only anonymous users to see Captchas.
 - If you want everyone to see Captchas except members of specific groups, selectMembers 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. ChooseSave

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.

3 Shortcut links (e.g. CONF-2622@JIRA) and internal links to other pages within Confluence are not tagged.

To hide external links from search engines:

- 1. Choose the cog icon , then choose General Configuration
- 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.

(i) Background to the nofollow attribute

As part of the effort to combat the spamming of wikis and blogs (Confluence being both), Google came up with some markup which instructs search engines not to follow links. By removing the main benefit of wiki-spamming it's hoped that the practice will stop being cost-effective and eventually die out.

Configuring Captcha for Failed Logins

If you have confluence administrator permissions, you can configure Confluence to impose a maximum number of repeated login attempts. After a given number of failed login attempts (the default is three) Confluence will display a Captcha form asking the user to enter a given word when attempting to log in again. This will prevent brute force attacks on the Confluence login screen.

Similarly, after three failed login attempts via the XML-RPC or SOAP API, an error message will be returned instructing the user to log in via the web interface. Captcha will automatically be activated when they attempt this login.

On this page:

- Enabling, Disabling and Configuring Captcha for Failed Logins
- Notes

'Captcha' is a test that can distinguish a human being from an automated agent such as a web spider or robot. When Captcha is activated, users will need to recognize 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



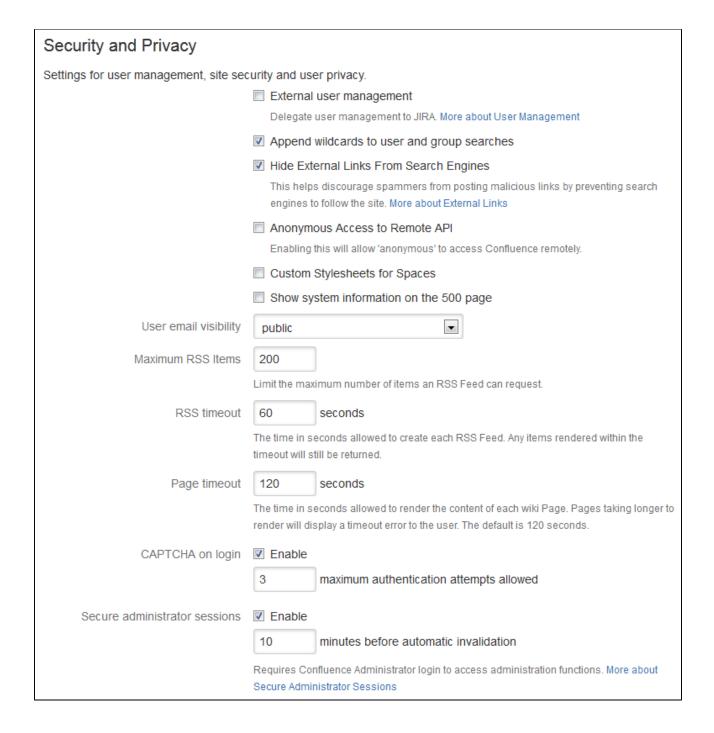
Enabling, Disabling and Configuring Captcha for Failed Logins

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
- 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



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 ad ministrative actions, change of email address and Captcha for failed logins. See Recognized 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
- 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

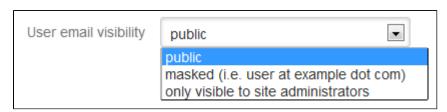
Confluence provides three options for email address privacy which can be configured by a Confluence administrator from the **Administration Console**:

- 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
- 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
- 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.

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.

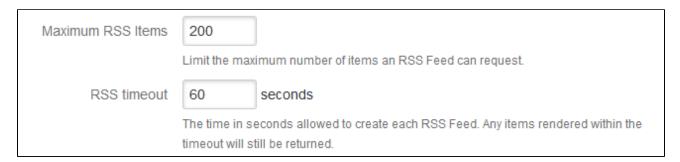
Both of these are set in the 'Edit Security Configuration' screen.

To configure RSS feeds:

- Choose the cog icon then choose Genera I Configuration
- 2. ChooseSecurity Configuration.
- 3. ChooseEdit.
- 4. Enter a value for **Maximum RSS Items**. The default value is 200.
- 5. Enter a value for RSS timeout.
- 6. ChooseSave.

Screenshot: Configuring RSS feeds

On this page: • Notes Related pages: • The RSS Feed Builder



Notes

- When using the RSS Feed Builder, a user could potentially enter such a large value for the number of feed items returned that Confluence would eventually run out of memory.
- When using the Feed Builder, if a users a value greater than this setting (or less than 0) they will get a
 validation error.
- If any pre-existing feeds are set to request more than the configured maximum, they will be supplied
 with only the configured maximum number of items. This is done silently there is no logging and no
 message is returned to the RSS reader.
- If Confluence times out when responding to an RSS feed request, any items already rendered are returned.

Preventing and Cleaning Up Spam

If your Confluence site is public-facing you may be affected by spammers.

Stopping Spammers

To prevent spammers:

- 1. Enable Captcha. See Configuring Captcha for Spam Prevention.
- 2. Run Confluence behind an Apache webserver and create rules to block the spammer's IP address.

Blocking Spam at Apache or System Level

If a spam bot is attacking your Confluence site, they are probably coming from one IP address or a small range of IP addresses. To find the attacker's IP address, follow the Apache access logs in real time and filter for a page that they are attacking.

For example, if the spammers are creating users, you can look for signup.action:

```
$ tail -f confluence.atlassian.com.log | grep signup.action
1.2.3.4 - - [13/Jan/2010:00:14:51 -0600] "GET /signup.action HTTP/1.1" 200 9956 "-" "Mozilla/4.0
(compatible; MSIE 6.0; Windows NT 5.1; SV1)" 37750
```

Compare the actual spam users being created with the log entries to make sure you do not block legitimate users. By default, Apache logs the client's IP address in the first field of the log line.

Once you have the offender's IP address or IP range, you can add it to your firewall's blacklist. For example, using the popular Shorewall firewall for Linux you can simply do this:

```
# echo "1.2.3.4" >> /etc/shorewall/blacklist
# /etc/init.d/shorewall reload
```

To block an IP address at the Apache level, add this line to your Apache vhost config:

```
Deny from 1.2.3.4
```

You can restart Apache with a "graceful" command which will apply the changes without dropping any current sessions.

If this still does not stop the spam, then consider turning off public signup.

Deleting Spam

Profile Spam

By 'profile spam', we mean spammers who create accounts on Confluence and post links to their profile page.

If you have had many such spam profiles created, the easiest way to delete them is via SQL.

To delete a spam profile:

- Shut down Confluence and back up your database.
 Note: This step is essential before you run any SQL commands on your database.
- Find the last real profile:

```
SELECT bodycontentid, body FROM bodycontent WHERE contentid IN (SELECT contentid FROM content WHERE contenttype='USERINFO')
ORDER BY bodycontentid DESC;
```

- 3. Look through the bodies of the profile pages until you find where the spammer starts. You may have to identify an number of ranges.
- 4. Find the killset:

```
CREATE TEMP TABLE killset AS SELECT bc.bodycontentid,c.contentid,c.username FROM bodycontent bc JOIN content c ON bc.contentid=c.contentid WHERE bodycontentid >= BOTTOM_OF_SPAM_RANGE AND bodycontentID <= TOP_OF_SPAM_RANGE AND c.contenttype='USERINFO';

DELETE FROM bodycontent WHERE bodycontentid IN (SELECT bodycontentid FROM killset);

DELETE FROM links WHERE contentid IN (SELECT contentid FROM killset);

DELETE FROM content WHERE prevver IN (SELECT contentid FROM killset);

DELETE FROM content WHERE pageid IN (SELECT contentid FROM killset);

DELETE FROM content WHERE contentid IN (SELECT contentid FROM killset);

DELETE FROM os_user_group WHERE user_id IN (SELECT id FROM killset k JOIN os_user o ON o.username=k.username);

DELETE FROM os_user WHERE username IN (SELECT username FROM killset);
```

If you're using Confluence 5.6 or earlier use the SQL commands below:

```
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.

Configuring a Confluence Environment

This section describes the external setup of your Confluence installation. It includes information on configuring the web server, application server, directories and files everything to do with the environment that Confluence runs in. For guidelines on modifying settings inside the application, see Configuring Confluence instead.

Confluence is a J2EE web application. On the client side, users access Confluence primarily via a web browser.

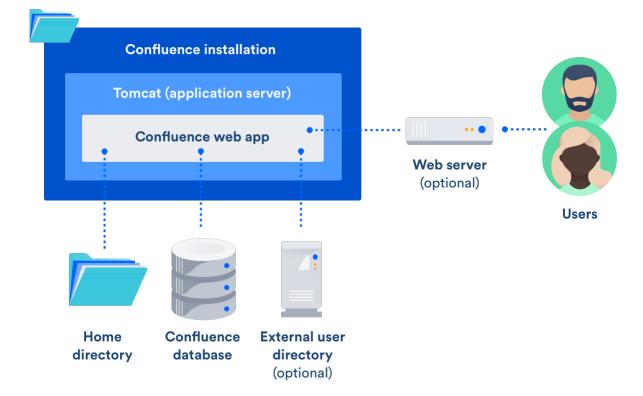
This section contains the following guidelines:

- Confluence Home and other important directories
- Application Server Configuration
- Starting Confluence Automatically on System Startup

Diagram: A Confluence installation

Related pages:

- Getting Started as Confluence Administrator
- Supported Platforms



Confluence Home and other important directories

Confluence installation directory

The 'Confluence Installation directory' is the directory where Confluence was installed. This directory is also sometimes called the 'Confluence Install directory'.

Important files in the installation directory:

- bin/setenv.batorbin/setenv.sh
 This file is used to edit CATALINA_OPTS memory and garbage collection settings and define system properties.
- confluence/WEB-INF/classes /confluence-init.properties
 This file contains the location of the Confluence Home directory.

On this page:

- Confluence installation directory
- Confluence home directory
- Changing the location of the home directory
- Database
- Temp directory (installation directory)

Confluence home directory

The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. Another term for 'Home directory' would be 'data directory'. We also refer to this as the 'local home directory' in Data Center.

Finding the home directory

The location of the Confluence home directory is defined when youinstall Confluence. This location is stored in the confluence-init.properties file, which is located in the confluence/WEB-INF/classes directory of your Confluence Installation directory.

When Confluence is running you can find the location of the home directory in General Configuration > S ystemInformation > Confluence Information - Confluence Home.

If you're using Confluence Data Center (a clustered instance), 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. The location of your shared home directory can be found in your<local-home>/confluence.cfg.xmlfilein theconfluence.cluster.homeproperty.

Contents of the home directory

The Confluence home directory contains some of the configuration data used by Confluence. This section outlines the purpose of the files and directories in the Confluence home directory.

File or directory	Purpose
conflue nce. 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

attachm	This directory contains every version of each attachment stored in Confluence.	
ents/	You can specify an alternative directory for attachment storage by setting theattachments. dirproperty inconfluence.cfg.xml.	
	In Data Center this directory is usually found in the Shared Home directory.	
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 thedaily.backup.dirprope rty inconfluence.cfg.xml.	
	In Data Center this directory is usually found in the Shared Home directory.	
bundled - plugins/	Confluence includes a set of bundled plugins. 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, as it will be replaced the next time Confluence starts up.	
databas e/	This is where Confluence stores its database when configured to run with the Embedded H2 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.	
	The H2 database is provided for evaluating Confluence and is not supported as a production database.	
index/	The Confluence index is heavily used by the application for content searching and recently updated lists and is critical for a running Confluence instance. If data in this directory is lost or corrupted, it can be restored by running a full reindex from within Confluence. This process can take a long time depending on how much data is stored Confluence's database.	
	An alternative directory may be specified for the index by setting the lucene.index.dirproperty inconfluence.cfg.xml.	
journal/	Entries are added to the journal when changes occur (such as a comment, like, new page). Journal entries are then processed and the entries added to the index (about every 5 seconds). In a cluster, the journal keeps the indexes on each node in sync.	
logs/	Confluence's application logs are stored in this directory.	
plugin- cache/	All Confluence plugins are stored in the database. To allow for quicker access to classes contained within the plugin JARs, Confluence will cache these plugins in theplugin-cached irectory. 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.	
temp/	Thetempdirectory is used for runtime functions such as exporting, importing, file upload and indexing. Files in this directory are temporary and can be safely removed when Confluence is offline. A daily job within Confluence deletes files that are no longer needed.	
	You can specify a different temp directory location, if necessary. Edit <confluence-home> /confluence.cfg.xml and set the new location in the webwork.multipart.saveDir property. Your new location can't be in the installation directory, as this will cause some functions, such as download, to fail. We recommend you keep the temp directory in the local home directory.</confluence-home>	
thumbna ils/	Stores temporary files for image thumbnails. This directory is essentially a thumbnail cache, and files deleted from this directory will be regenerated the next time the image is accessed.	
	In Data Center this directory is usually found in the Shared Home directory.	

sharedhome/ This sub-directory is created in your home directory when you install Confluence Server. If you choose move to Data Center, you'll move the contents of this directory to a seperate shared home directory, that is accessible to all nodes.

Cache files for some features, including Office document and PDF previews are also stored in this directory (in both Server and Data Center).

Changing the location of the home directory

When Confluence first starts up, it reads the confluence-init.properties file to determine where to look for the Home directory.

To change the location of the home directory edittheconfluence.homeproperty in theconfluence-init. propertiesfile as follows:

Windows

In Windows, the pathC:\confluence\datawould be written as: confluence.home=C:/confluence/data

Note that all backslashes (\) are written as forward slashes (/)

Linux / Solaris

On any Linux-based system, the property is defined using the normal directory syntax: confluence.home=/var/confluence/

Symbolic links

There can be no symbolic links within the Confluence home directory. You must define an absolute path. If disk space is an issue, place the entireconfluence.homedirectory 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 theconfluence.homedirectory 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.xmlinside 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.dirif you have not configured your backups to be placed elsewhere already
- hibernate.connection.urlif you are using the embedded HSQL database.

Database

All other data, including page content, is kept in the database. If you installed Confluence as a trial, or chose to use theembedded HSQL databaseduring setup, the database will store its files underdatabase/in theConfluence Home Directory. Otherwise, the database management system you are connecting to is responsible for where and how your remaining data is stored.

Temp directory (installation directory)

The temp directory is configured in the Java runtime and some Confluence components write temporary files or lockfiles into this directory.

The tempdirectoryis located in theinstallationdirectory 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.

⚠Note: this is not the same as the temp directory in Confluence Home where exports, for example, are saved. See the table above to find out how to change the location of the <confluence-home > / temp directory.

There's a known issue with setting a temp directory in Confluence. See

CONFSERVER-59613 - java.io.tmpdir has no effect on changing the installation temp directory GATHERING IMPACT

Application Server Configuration

The following pages contain information about configuring your application server for Confluence:

Managing Application Server Memory Settings

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 optimizing for performance should use this document as a starting point only.



Por 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 maximized 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

Starting Confluence Automatically on System Startup

You can configure Confluence to start automatically on system startup, allowing it to recover automatically after a reboot.

- Start Confluence Automatically on Linux
- Start Confluence Automatically on Windows as a Service

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.In this example we're using/usr/local/confluence.

```
sudo mkdir /usr/local/confluence
sudo chown confluence: /usr/local/confluence
```

3. Log in as the confluence user to install Confluence:

```
sudo su - confluence
cd /usr/local/confluence/
tar zxvf /tmp/confluence-5.6.4.tar.gz
ln -s confluence-5.6.4/ current
```

- 4. Edit <<CONFLUENCE_INSTALL_DIRECTORY>>/confluence/WEB-INF/classes/confluence-init. properties file, and set confluence.home=/usr/local/confluence/<Confluence_Data_Home> (ensure you have removed the comment '#')
- 5. Then back as root, create the file /etc/init.d/confluence (code shown below), which will be responsible for starting up Confluence after a reboot (or when manually invoked).
 A 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 Confluence install directory
CATALINA HOME=/usr/local/confluence/current
# Location of Java JDK
export JAVA_HOME=/usr/lib/jvm/java-7-oracle
case "$1" in
  # Start command
  start)
    echo "Starting $APP"
   /bin/su -m $USER -c "$CATALINA_HOME/bin/start-confluence.sh &> /dev/null"
  # Stop command
  stop)
    echo "Stopping $APP"
    /bin/su -m $USER -c "$CATALINA_HOME/bin/stop-confluence.sh &> /dev/null"
   echo "$APP stopped successfully"
   ;;
   # Restart command
   restart)
       $0 stop
       sleep 5
       $0 start
    echo "Usage: /etc/init.d/$APP {start|restart|stop}"
   exit 1
    ;;
esac
exit 0
```

6. Make this file executable:

```
sudo chmod +x /etc/init.d/confluence
```

- 7. Set this file to run at the appropriate runlevel. For example, use sudo chkconfig --add confluence on Redhat-based systems, sudo update-rc.d confluence defaults or rcconf on Debianbased systems.
- 8. You should now be able to start Confluence with the init script. A successful startup output typically looks like this:

```
$ sudo /etc/init.d/confluence start
Starting Confluence:
If you encounter issues starting up Confluence, please see the Installation guide at
http://confluence.atlassian.com/display/DOC/Confluence+Installation+Guide
Using CATALINA_BASE: /usr/local/confluence/current
Using CATALINA_HOME: /usr/local/confluence/current
Using CATALINA_TMPDIR: /usr/local/confluence/current/temp
Using JRE_HOME: /usr/lib/jvm/java-1.7.0-oracle
done.
```

You should then see this running at http://<server>:8090/

1 The port for this will be whatever is defined in your Confluence server.xml file.

Adding Confluence as a service for Ubuntu Jaunty (or later)

To continue configuring Confluence to start automatically as a service on Ubuntu Jaunty (or later):

1. After logging in as the confluence user to install Confluence, create start and stop scripts in /usr /local/confluence:

Example startscript:

```
#!/bin/bash
export JAVA_HOME=/usr/lib/jvm/java-7-oracle-1.7.0.71/
export JDK_HOME=/usr/lib/jvm/java-7-oracle-1.7.0.71/
cd /usr/local/confluence/current/bin
./startup.sh
```

Example stopscript:

```
#!/bin/bash
export JAVA_HOME=/usr/lib/jvm/java-7-oracle-1.7.0.71/
export JDK_HOME=/usr/lib/jvm/java-7-oracle-1.6.0.71/
cd /usr/local/confluence/current/bin
./shutdown.sh
```

- 2. Make both of these scripts executable. For example, by issuing the command: sudo chmod a+x /usr/local/confluence/start /usr/local/confluence/stop.
- 3. Karmic and later: Create two text files in /etc/init/ called confluence-up.conf and confluence-down.conf:

confluence-up:

```
start on runlevel [2345]
script

date >> /tmp/confluence-startup.out
    exec sudo -u confluence /usr/local/confluence/start >> /tmp/confluence-startup.out 2>&1
end script
```

confluence-down:

```
start on runlevel [16]
expect fork
respawn
exec sudo -u confluence /usr/local/confluence/stop >> /tmp/confluence-shutdown.out 2>&1
```

... and make them readable to all users:

```
sudo chmod a+r /etc/init/confluence-up.conf /etc/init/confluence-down.conf
```

 Jaunty, Intrepid: Create two text files in /etc/event.d/ called confluence-up and confluencedown:

```
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:

 $\verb|sudo| chmod| a+r /etc/event.d/confluence-up /etc/event.d/confluence-down|$

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 as a service: using the Confluence installeror manually as described below.

On this page:

- Reasons for starting Confluence as a
- Changing the user running the service
- Manually installing Confluence 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\conflu ence **not** $z : \confluence$).

For more detail, see Creating a Dedicated User Account on the Operating System to Run Confluence.

Manually installing Confluence as a service

In Windows:

- 1. Open a command prompt and change directory to the <CONFLUENCE-INSTALL>/bin directory. You'll need to run the command prompt using 'Run as administrator' so that you can complete some
- 2. Confirm that the JAVA_HOME variable is set to the JDK base directory with the command:

echo %JAVA HOME%

If you installed the Java Runtime Environment (JRE) or used the Confluence installer, replace JAVA_H OME withJRE_HOME. See Setting the JAVA_HOME Variable in Windowsfor more info.

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. Use the following command to install the service with default settings:

```
service.bat install Confluence
```

The service will be called **Atlassian Confluence** and will be configured to start automatically by default, but will not automatically start up until the next server reboot.

4. If you have a large Confluence installation, you can increase the maximum memory Confluence can use (the default is 1024MB). For example, you can set the maximum memory to 2048MB using:

```
tomcat9w //US//Confluence --JvmMx 2048
```

5. If you don't have any JVM parameters that you pass toConfluence, you can skip this step. If you do, add them to the service using:

```
tomcat9w //US//Confluence ++JvmOptions="-Djust.an.example=True"
```

Alternatively you can use the following commandto launch the service properties dialog then navigate to the Java tab to add more JVM parameters.

```
tomcat9w //ES//Confluence
```

For further configuration options, please refer to the Tomcat Windows Service How-To guide.

- 6. Go to Control Panel > Administrative Tools > Services > Atlassian Confluence and right-clickProperties to verify the settings are correct. Start the Confluence service with the command:
- 7. Finally, start the Confluence service. From now on this will happen automatically after the a server reboot.

```
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 StartPathparameter to your new folder. Refer to the Tomcat documentation for help.

Troubleshooting Confluence while running as a Windows service

- Check the Knowledge Base articles:
 - Getting 'The image file tomcat6.exe is valid, but is for a machine type other than the current machine'
 - Confluence Does Not Start Due to Windows Firewall
 - Unable to start Confluence Windows service after allocating JVM memory
 - Unable to Configure Confluence to Run as a Service on Tomcat 5
 - Unable to Install Service on Windows Vista
- If none of the above solves your problem, please refer to the complete list of known issues in our Knowledge Base.
- When investigating memory issues or bugs, it may be useful to view information from Confluence's garbage collection. To turn on the verbose garbage collection seeHow to Enable Garbage Collection (GC) Logging.
- 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.



Mote

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:

- 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.
- A support zip, containing your application logs and configuration files.

Performance Tuning

This document describes tuning your application for improved performance. It is not a guide to troubleshooting Confluence outages. Check Trouble shooting Confluence hanging or crashing for help if Confluence is crashing.

Like any server application, Confluence may require some tuning as it is put under heavier use. We do our best to make sure Confluence performs well under a wide variety of circumstances, but there's no single configuration that is best for everyone's environment and usage patterns.

If you are having problems with the performance of Confluence and need our help resolving them, you should read Requesting Performance Support.

Performance Data Collector

ThePerformance Data Collectoris a server-side, standalone application that exposes a number of REST APIs for collecting performance data. It can be used to collect data, such as thread dumps, disk speed and CPU usage information, to troubleshoot performance problems.

SeeHow to use the Performance Data Collectorfor more information.

On this page:

- Performance Data Collector
- Use the latest version of your tools
- Avoid swapping due to not enough RAM
- Being aware of other systems using the same infrastructure
- · Choice of database
- Database connection pool
- Database in general
- Database statistics and query analyzers
- Cache tuning in Confluence and Apache
- Antivirus software
- Enabling HTTP compression
- Performance testing
- Access logs
- Built-in profiler
- Application server memory settings
- Web server configuration
- Troubleshooting possible memory leaks
- Enable faster permissions

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.

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 aCluster Panic due to Performance Problems. This is because swapping causes the JVM to pause duringGar bage Collection, which in turn can break the inter-node communication required to keep the clustered nodes in sync.

Being aware of other systems using the same infrastructure

It may sound tempting: Just have one powerful server hosting your database and/or application server, and run **all** your crucial programs on that server. **If** the system is set up perfectly, then you might be fine. Chances are however that you are missing something, and then one application's bug might start affecting other applications. So if Confluence is slow every day around noon, then maybe this is because another application is using the shared database to generate complicated reports at that time? Either make sure applications can't harm each other despite sharing the same infrastructure, or get these systems untangled, for example by moving them to separate instances that can be controlled better.

Choice of database

The **embedded H2 database**is provided for evaluating Confluence, not for production Confluence sites. After the evaluation finishes, you must switch to a <u>supported external database</u>. 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 andhibernate.hikari.maximumPoolSize property (if present)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

Database Latency

O

Confluence Usage

The latency is calculated by sending a trivial request to the database, querying a table which is known to have only one column and one row. ("select * from CLUSTERSAFETY"). Obviously this query should be blazing fast, and return within 1 or 2 milliseconds. If the value displayed is between 3 and 5 milliseconds, you might already have an issue. If the value is above 10ms, then you **definitely** need to investigate and improve something! A few milliseconds may not sound so bad, but consider that Confluence sends quite a few database queries per page request, and those queries are a lot more complex too! High latency might stem from all sorts of problems (slow network, slow database, connection-pool contention, etc), so it's up to you to investigate. Don't stop improving until latency is below 2ms on average.

Obviously, latency is just the very first thing to look at. You may get zero latency and still have massive database problems, e.g. if your tables are poorly indexed. **So don't let a low latency fool you either.**

Database statistics and query analyzers

Modern databases have query optimizers based on collecting statistics on the current data. Using the SQL EXPLAIN statement will provide you information on how well the query optimizer 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.

Antivirus software

Antivirus software greatly decreases the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental, and may even cause errors with Confluence. You should configure your antivirus software to ignore the Confluence home directory, its index directory and any database-related directories.

Enabling HTTP compression

If bandwidth is responsible for bottlenecking in your Confluence installation, you should consider enabling HTTP compression. This may also be useful when running an external facing instance to reduce your bandwidth costs.

Take note of the known issues with HTTP compression in versions of Confluence prior to 2.8, which may result in high memory consumption.

Performance testing

You should try out all configuration changes on a demo system. Ideally, you should run and customize loadtests that simulate user behavior.

Access logs

You can find out which pages are slow and which users are accessing them by enabling Confluence's built-in access logging.

Built-in profiler

You can identify the cause of page delays using Confluence's built-in profiler according to Troubleshooting Slow Performance Using Page Request Profiling.

Application server memory settings

See How to fix out of memory errors by increasing available memory.

Web server configuration

For high-load environments, performance can be improved by using a web server such as Apache in front of the application server. There is a configuration guide to Running Confluence behind Apache.

When configuring your new web server, make sure you configure sufficient threads/processes to handle the load. This applies to both the web server and the application server connector, which are typically configured separately. If possible, you should enable connection pooling in your web server connections to the application server.

Troubleshooting possible memory leaks

Some external plugins, usually ones that have been written a long time ago and that are not actively maintained anymore, have been reported to consume memory and never return it. Ultimately this can lead to a crash, but first this manifests as reduced performance. The Troubleshooting Confluence hanging or crashing guide is a good place to start. Some of the known causes listed there could result in performance issues short of a crash or hang.

Enable faster permissions

Confluence needs to regularly check the current user's permissions in order to determine what to display. The faster permissions service changes the way permissions information is stored in the database to optimize these permissions checks. Although this comes with some overhead, it can provide a significant performance improvement in sites with a lot of content and complex permissions. If you only have a small amount of content or very simple permissions (just a few groups, or few nested page restrictions), this service is unlikely to make your Confluence site significantly faster.

Learn more about the faster permissions service

Cache Performance Tuning

Confluence performance can be significantly affected by the performance of its caches.

Before you change the size of your caches, it's important to take a baseline so you can measure how effective each individual change is, and decide whether they are needed.

On this page we'll take you through some example statistics and discuss how you might be able to improve Confluence performance by resizing these caches.

If you just want to check your cache statistics, or make a change to your cache config, seeCache Statistics.

Cache tuning example

On this page:

- Cache tuning example
- Important caches

As an example of how to tune Confluence's caches, let's have a look at the following table:

Caches	% Utilization	% Effectiveness	Objects/Size	Hit/Miss/Expiry
Attachments	87%	29%	874/1000	78226/189715/187530
Content Attachments	29%	9%	292/1000	4289/41012/20569
Content Bodies	98%	81%	987/1000	28717/6671/5522
Content Label Mappings	29%	20%	294/1000	4693/18185/9150
Database Queries	96%	54%	968/1000	105949/86889/83334
Object Properties	27%	18%	279/1000	5746/25386/8102
Page Comments	26%	11%	261/1000	2304/17178/8606
Users	98%	5%	982/1000	6561/115330/114279

The maximum size of the caches above is 1000 (meaning that it can contain up to 1000 objects). You can tell when a cache size needs to be increased because the cache has both:

- a high usage percentage (above 75%)
- a low effectiveness percentage.

Check the 'effectiveness' versus the 'percent used'. A cache with a low percent used need not have its size lowered; it does not use more memory until the cache is filled.

Based on this, the sizes of the "Attachments", "Database Queries", and "Users" caches should be increased to improve their effectiveness.

As the stored information gets older or unused it will expire and be evicted from the cache. Cache expiry can be based on time or on frequency of use.

There's 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.

Important caches



The following suggestions are general guidelines. In cases of large databases, 20-30% of the size of the table may be unnecessarily large. Check the effectiveness and percent used categories in the cache for more specific assessments.

- Content Objects cache (com.atlassian.confluence.core.ContentEntityObject) should be set to at least 20-30% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the queryselect count(*) from CONTENT where prevver is null.
- Content Body Mappings cache (com.atlassian.confluence.core.ContentEntityObject. bodyContents) should be set to at least 20% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the queryselect count (*) from CONTENT where prevver is null.
- Embedded Crowd Internal User cache(com.atlassian.crowd.model.user.InternalUser) should be set to the number of users you have in the internal directory. You can discover this number by using the following SQL:

```
SELECT
   COUNT(*)
FROM
   cwd_user u
JOIN
   cwd directory d
   u.directory_id = d.id
AND d.directory_name = 'Confluence Internal Directory';
```

• Embedded Crowd Users cachecom.atlassian.confluence.user.crowd. CachedCrowdUserDao.USER CACHE should be set to the number of rows in the cwd_user table.

```
SELECT
COUNT(*)
FROM
cwd user u;
```

• Space permissions by ID cache (com.atlassian.confluence.security.SpacePermission

should be set to the number of space permissions in your deployment (a good rule of thumb is 20 times the number of spaces). You can find the number of space permissions using the queryselect count(*) from SPACEPERMISSIONS.

Cache Statistics

Caches help reduce the load on your database, and can make some operations faster. Track the size and hit ratio of each of Confluence's internal caches, and adjust the cache size for better performance.

On this page:

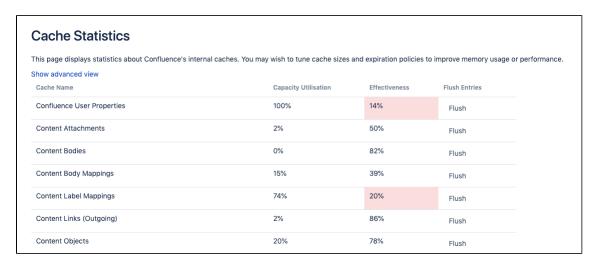
- View cache statistics
 - View cache statistics in a cluster
 - What the statistics mean
 - Cache types
- Change the size of a cache

View cache statistics

To view cache statistics:

- 1. Go to General Configuration>Cache Management.
- 2. SelectShow Advanced View to see the full details.

Screenshot: Cache statistics screen showing the utilisation and effectiveness of a selection of caches.



View cache statistics in a cluster

If you're running Confluence in a cluster, this screen shows the statistics for the node you're currently on.

To view cache statistics for another node in the cluster:

- 1. Go to General Configuration>Clustering.
- 2. Select ••• > Cache statistics next to the node you want to view.

You will only be able to view the statistics. To flush a cache or adjust the size, you'll need to access the Cache Management screen on each node directly.

What the statistics mean

Here's some information on how eachnumber is generated.

Capacity	=(Objects)/(Size)
Utilization	For example Percent Used = 4023 / 5000 = 80%
Effectiveness	=(Hits)/(Hits + Misses)
:	For example Effectiveness = 374550 / (374550 + 140460) = 73%

Current / Max Entries	The number of entries in the cache / the number of total possible entries allowed (this is the size of the cache).
Current Heap Size	Heap memory (in MB) allocated to this cache (if applicable)
Hit / Miss / Evicted	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.
Flush	Flushes the cache.

Cache types

When running in a cluster, Confluence has three types of caches:

- **local** cache data is replicated on each node.
- distributed -cache data isevenly partitioned across all Confluence nodes in the cluster (known as replicate-via-copy).
- **hybrid** cache data is replicated on each node, and invalidated remotely by other nodes when things change (known as replicate-via-invalidation).

The cache type is indicated with a lozenge beside the cache name in the advanced view.

Screenshot: Cache statistics advanced view showing the full details of each cache, including the cache type.

Cache Statistics							
his page displays statistics about Confluence's internal caches. You may wish to tune cache sizes and expiration policies to improve memory usage or performance.							
Cache Name	Capacity Utilisation	Effectiveness	Current / Max Entries	Current Heap Size (MB)	Hit / Miss / Evicted	Adjust Max Entries	Flush Entries
Confluence access annotations LOCAL	1%	89%	18 / 1000	0	304 / 36 / 0	Adjust Size	Flush
Confluence and Jira Content Connector DISTRIBUTED	0%	Unknown	7 / 1000	0	0/7/0	Adjust Size	Flush
Confluence User Properties DISTRIBUTED	100%	19%	10000 / 10000	0	7059 / 29981 / 19979	Adjust Size	Flush
Content Attachments HYBRID	3%	41%	34 / 1000	<1	24 / 34 / 0	Adjust Size	Flush
Content Bodies HYBRID	1%	85%	12 / 1000	<1	960 / 167 / 0	Adjust Size	Flush

Change the size of a cache

Tuning the size of a cache can speed up Confluence (if the caches are too small), or reduce memory (if the caches are too large). Larger caches will require more memory at runtime, so make sure you review the memory allocation of the Confluence Java process and the physical memory available on your server.

You need **System Administrator** global permission to change the size of a cache.

To change the size of a cache:

- 1. Go to General Configuration>Cache Management.
- 2. SelectShow Advanced View.
- 3. Select Adjust Size next to the cache you want to change.
- 4. Enter the maximum number of items to be stored in the cache and select **Submit**.

The changes will take effect immediately. You don't need to restart Confluence.

Any changes to cache sizes are recorded in:

- <home-directory>/shared-home/config/cache-settings-overrides.properties if you run Confluence on a single server.
- <shared-home>/config/cache-settings-overrides.propertiesif you run Confluence in a cluster.

To reset the values back to the default, you can delete the cache-settings-overrides.properties file and restart Confluence.

SeePerformance Tuning for a more general overview of tuning in 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.

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

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 customize Confluence's cache to meet your needs, seecache tuning.

To increase the amount of memory available to Confluence, seeHow to fix out of memory errors by increasing available memory.

Mail error queue

Confluence keeps a copy of all emails that it failed to send within an internal error queue. In the event of intermittent failures such as network connectivity issues, the emails in this queue can be manually resent when the problem is fixed. Under certain circumstances, the mail queue can fill up with large objects. The queue is regularly flushed, but if you get a *lot* of mail errors, you might get a spike in memory usage.

Attachments

The indexing of large attachments requires that the attachment be loaded into memory. In the case of large attachments, this can cause a temporary strain on the systems resources, and may result in indexing failing 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:

```
\label{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 minimize 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 minimize the impact that garbage collection has on the running application.

Requesting Performance Support

Basic performance troubleshooting steps

Begin with the following procedures:

- 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 optimize 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:

On this page:

- Basic performance troubleshooting steps
- Requesting basic performance support
- Advanced performance troubleshooting

Related pages:

- Memory Usage and Requirements
- Confluence for Enterprise
- Asupport zip, containing log files and configuration, ideally with a series ofthread dumps separated by 10 seconds.
- 2. A description with as much detail as possible regarding:
 - a. What changes have been made to the system?
 - b. When did performance problems begin?
 - c. When in the day do performance issues occur?
 - d. What pages or operations experience performance issues?
 - e. Is there a pattern?

Continue with as much of the advanced performance troubleshooting information as you can.

Advanced performance troubleshooting

Please gather **all** of the information listed below and include it in your support request, even if you think you have a good idea what's causing the problem. That way we don't have to ask for it later.

System information

Confluence server

- Take a screenshot of Confluence's Administration System Information (or save the page as HTML)
- Take a screenshot of Confluence's Administration Cache Statistics (or save the page as HTML)
- Find out the exact hardware Confluence is running on
 - How many CPUs? What make and model? What MHz?
 - How much memory is installed on the machine?
 - How much memory is assigned to Confluence's JVM? (i.e. what are the -Xmx and -Xms settings for the JVM?)
 - What other applications are being hosted on the same box?

Confluence content

- How many users are registered in Confluence?
- On average, to how many groups does each user belong?
- How many spaces (global and personal) are there in your Confluence server?
- How many of those spaces would be viewable by the average user?
- Approximately how many pages? (Connect to your database and perform 'select count(*)
 from content where prevver is null and contenttype = 'PAGE')
- How much data is being stored in Bandana (where plugins usually store data)? (Connect to your database and perform 'select count(*), sum(length(bandanavalue)) from bandana')

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?

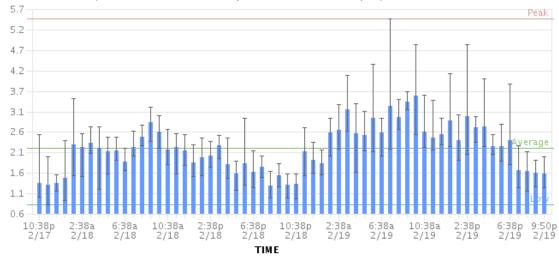
Monitoring data

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.

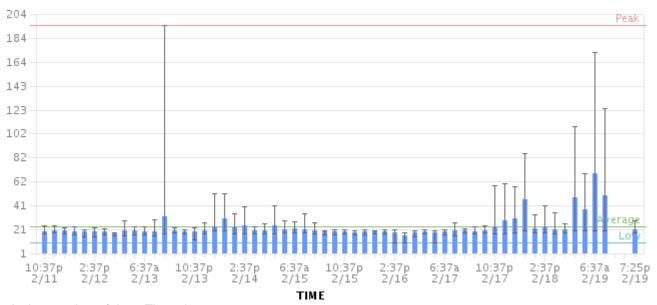
Here are sample graphs from the confluence.atlassian.com system, showing

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.

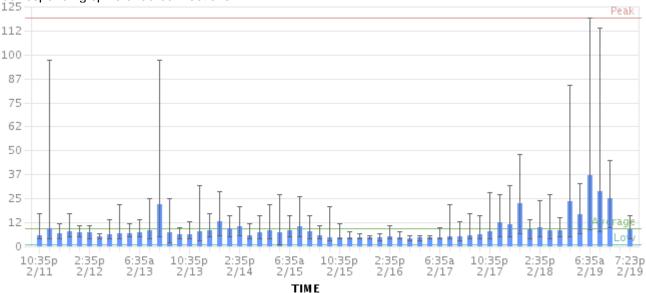


Resin threads and database connections



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

- Access logging is enabled by default. You can find Confluence access logs at <confluence install>/logs/conf_access_log.<date>.log. These logs are configured in the server.xml. Refer to the Tomcat Access Log Valve documentation for more information on the attributes that can be logged.
- You can also chose to enable access logging in your application logs via the log4j configuration. SeeH
 ow to Enable User Access Logging, including how to redirect the logs to a separate file

You can run these access log files through a log file analyzer such as AWStats, or manually look for pages which are slow to load.

Profiling and logs

- Enable Confluence's built-in profiling for long enough to demonstrate the performance problem using T roubleshooting 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.

Next steps

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.

Compressing an HTTP Response within Confluence

Confluence supports HTTP GZip transfer encoding. This means that Confluence will compress the data it sends to the user, which can speed up Confluence over slow or congested Internet links, and reduce the amount of bandwidth consumed by a Confluence server.

Turn on Confluence's GZip encoding if:

- Users are accessing Confluence over the Internet, or a WAN connection with limited bandwidth.
- You wish to reduce the amount of data transfer between the Confluence server and client.

If you are accessing Confluence over a Local Area Network or over a particularly fast WAN, you may wish to leave GZip encoding disabled. If the network is fast enough that transferring data from Confluence to the user isn't a limiting factor, the additional CPU load caused by compressing each HTTP response may slow Confluence down.

Enabling HTTP Compression

- 1. Choose the cog icon O, then choose General Configuration
- 2. Select 'General Configuration' in the left-hand panel.
- 3. Enable 'Compress HTTP Responses'.

It is possible to configure which types of content are compressed within Confluence. By default, the following mime types will be compressed:

- text/htmltext
- javascript
- text/css
- text/plain
- application/x-javascript
- application/javascript

If you wish to change the types of content to be compressed, add a replacement urlrewrite-gzip-default.xml file within the WEB-INF/classes/com/atlassian/gzipfilter/ directory in your Confluence Installation Directory. A sample file is provided as an attachment. It is unlikely that you will need to alter this file.

Garbage Collector Performance Issues

The information on this page relates to memory management with Oracle's Hotspot JVM. These recommendations are based on our support team's successful experiences with customers with large Confluence instances.

Which garbage collector?

Confluence uses the garbage first garbage collector (G1GC) by default. This is the garbage collector we recommend.

SeeGarbage First Garbage Collector Tuningin the Oracle documentation for useful information on tuning this garbage collector.

We have also observed that G1GC performs better with a larger heap (2gb). See the information below about how to increase your heap size gradually.



Don't use the Concurrent Mark Sweep (CMS) Collector with Confluence, unless advised by Atlassian Support. It requires extensive manual tuning and testing, and is likely to result in degraded performance.

Use the right size heap

Keep your heap as small as possible, without the instance experiencing OutOfMemory errors. If you experience OutOfMemory errors and need to increase this, we recommend you do it in 512mb or 1gb allotments, and monitor the instance. If you continue to receive OutOfMemory errors, increase the heap by another 512mb or 1gb, and continue this process until you are operating stably with no OutOfMemory errors. Do not increase the heap further than required, as this will result in longer garbage collections.

Remove any old tuning parameters

On every full GC, the JVM will resize the allocations of Eden, Survivor etc based on the throughput it is actually seeing. It will tune itself based on the real world data of the objects that are being created and collected. Most of the time simply allowing JVM to tune itself will give you better performance.

If you have added JVM parameters in the past and are experiencing difficulties with GC now, we'd recommend you remove all GC related parameters, unless you added them to solve a specific problem, and they did in fact solve that problem. You should also consider re-benchmarking now to ensure that they are still solving that problem, and are not causing you any other issues.

Check your VM resources

If you run Confluence on a VM, check that is it 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, see our Garbage Collection (GC) Tuning Guide. This document was put together a few years ago, but has some useful information on choosing performance goals (throughput/footprint/latency), and how to tune for those goals.

Viewing your GC logs

How to Enable Garbage Collection (GC) Logging, and use a tool likeChewiebug's GCViewerto view the resulting logs.

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.

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

You needSystem Administratorglobal permissions to enable or disable profiling.

To enable or disable page profiling:

- 1. Go to Co Seneral Configuration Logging and Profiling.
- 2. Choose Enable Profiling or Disable Profiling.
- There is a known issue with the logging and profiling screen. See
 - CONFSERVER 58861 Turning on profiling has no effect on generated logging CLOSED

As a workaround you can enable profiling manually:

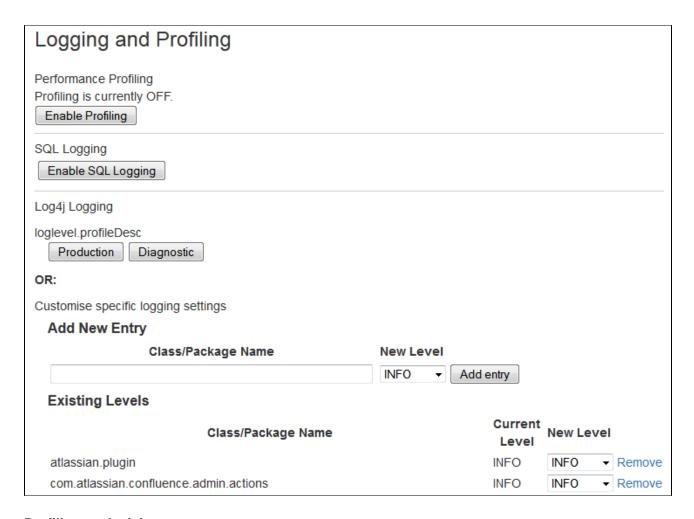
- 1. Go to Some Seneral Configuration > Logging and Profiling
- 2. Under Add a new entry, paste the following into the Class/package name field

com.atlassian.util.profiling.Timers

- 3. Set the logging level to **DEBUG**.
- 4. Choose Add Entry.

To disable profiling, find the com.atlassian.util.profiling.Timers entry in the list and choose Remove.

Screenshot: Changing Log Levels and Profiling



Profiling an Activity

- 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()
[28lms] - XW Interceptor: After defaultStack: /pages/viewpage.action (ViewPageAction.execute())
[28lms] - 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 section similar to the one below. Set the parameter value for **autostart** to **true**:

```
<filter>
    <filter-name>profiling</filter-name>
    <filter-class>com.atlassian.confluence.util.profiling.ConfluenceProfilingFilter</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.

Confluence Diagnostics

When investigating a performance problem or outage, it's useful to know as much as possible about what was happening in your site in the lead-up to the problem. This is when diagnostics information can help.

While often not individually actionable, diagnostic alerts can help you build up a detailed picture of your site's behaviour, and identify symptoms that may be contributing to the problem.



This feature is still experimental in Confluence 6.11. We plan to fine-tune the thresholds and provide a UI for this diagnostic information in an upcoming Confluence release. Stay tuned!

About diagnostic alerts

The purpose of the diagnostics tool is to continuously check for symptoms or behaviours that we know may contribute to problems in your site. An alert is triggered when a set threshold is exceeded.

For example, if the free disk space for your local home (or shared home) directory falls below 8192MB, an alert is triggered. This is useful because if you run out of space, your users may not be able to upload new files, export spaces, or perform other tasks that rely on writing files to that directory.

It's important to note that the thresholds are just the point at which the alert is triggered. It's not the same as a timeout, or other hard limit. For example a long running task may trigger an alert after 5 minutes, and still complete successfully after 8 minutes.

When an alert is triggered a message is written to the atlassian-confluence.logfile (your application log), and further details provided in the atlassian-diagnostics.log file. It's also included in support zips.

Some behaviours trigger a single alert, for others, multiple alerts are possible. Diagnostic information is stored in the database, and retained for 30 days. Old alerts are cleaned up automatically.

Types of alerts

There a several types of alerts.

Alert and KB	Level	Default threshold	Configurable	
Low free disk space	Critical	8192 megabytes	Yes	
Low free memory	Warn	256 megabytes	Yes	
Node left or joined the cluster	Info	N/A	No	
Long running task exceeded time limit	Warn	300 seconds	Yes	
Garbage collection exceeded time limit	Warn	10% (over the last 20 seconds)	Yes	

Availability

Some diagnostic alerts are disabled by default, because they may have a performance impact on your site, or are not designed to run continuously.

Our support team may ask you to enable one of the following alerts when troubleshooting a specific problem. They'll provide you with information on how to do this.

Alert and KB	Level	Default threshold	Configurable
HTTP request exceeded time limit	Warn	60 seconds	Yes
Macro rendering exceeded time limit	Warn	30 seconds	Yes

Thread memory allocation rate exceeded limit	Warn	5% over the last 20 seconds)	Yes
Sandbox crashed or was terminated during document conversion	Info	N/A	No

Alert levels

There are three levels of diagnostic alerts:

- Info information that might be useful when troubleshooting a problem, for example a node joined the cluster
- Warning a problem that may impact performance or availability in future, for example low memory
- **Critical** a serious problem that is likely to impact system stability or availability, for example low disk space.

Most alerts don't require any immediate action.

Change alert thresholds

Some alert thresholds are configurable. If you find you are seeing too many instances of an alert, you can change the threshold, so it's not triggered so easily.

Head to Recognized System Properties for a list of system properties for each alert. This info can also be found on the knowledge base article for each alert.

Change diagnostics behaviour

You can also change the way the diagnostics framework itself behaves. For example, you might change how often checks are performed, or how long diagnostics information is retained.

Head toRecognized System Properties for the full list of system properties.

Faster permissions service

Confluence's default permission checking method is very efficient when checking permissions for a single page, and guarantees strong permission consistency. However, this method can be quite slow and memory intensive when Confluence needs to check many thousands of pages. For example, to render the Task report macro, we need to find all the pages with a task assigned to the user, then check they have permission to see the spaces and pages the task appears on. The faster permissions service allows Confluence to check permissions on a large number of pages more quickly.

On this page:

- Enable the faster permissions service
- Disable the faster permissions service
- Troubleshooting

How it works

The faster permissions service replicates permissions information in a database structure that can be queried more efficiently. It uses database triggers to register changes to pages, spaces and their permissions, and records these in a change log. These change logs are processed regularly, and keep the fast permissions database tables up to date.

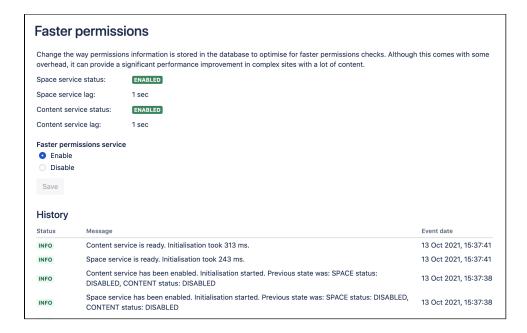
When faster permissions is enabled:

- the faster permissions method is used when listing pages in the task report, children display, and page index macros, search requests when the cache is cold, and on the dashboard the first time someone logs in.
- the default permissions checking method is always used when viewing a page, and in all other situations.

When space permissions or page restrictions are changed, it can take a few seconds for the faster permissions service to know about the change. Here's how we handle this delay:

- When a page is created and page restrictions are applied before being published for the first time, the page won't appear to anyone in macros or the dashboard until the fast permissions service has processed the change log.
- When a previously unrestricted page is restricted, it will continue to appear in macros and the
 dashboard for a few seconds, until the fast permissions service has processed the change log, then it
 will only appear to the relevant users.
- When a previously unrestricted page is restricted and the page title changed, it will continue to appear with the original title in macros and the dashboard for a few seconds, until the fast permissions service has processed the change log, then it will only appear (with the new title) to the relevant users.

The time it takes for a change to be reflected in faster permissions tables is shown in the faster permissions screen as **Space service lag** and **Content service lag**. If the lag exceeds 60 seconds, Confluence will automatically fall back to the slower default permission checking method until it decreases again. This is to ensure Confluence doesn't rely on out of date permissions information.



Screenshot of the faster permissions administration screen.

Enable the faster permissions service

To enable the faster permissions service:

- 1. Go to < confluence base url > /plugins / servlet / fasterpermissions
- 2. SelectEnable.

The status of each service will change to initializing as it starts up and begins to populate the new database tables. This may take some time depending on the size of your site, and is done in small batches to avoid any performance impact. Confluence is still completely operational during this time, and will start using each service as soon as the status changes to enabled.

Disable the faster permissions service

To disable the faster permissions service:

- 1. Go to < confluence base url > /plugins / servlet / fasterpermissions
- SelectDisable.
- 3. Wait for the services to shut down.

Confluence will use thedefault permission checking method from this point on. The additional database tables that were added when the service was first initialised will not be removed. If you later re-enable the service, the service will re-populate the tables.

Troubleshooting

Database requirements for faster permissions service

- If you use SQL Server, your database user needs full create, read and write permissions for the
 database tables. Confluence must be able to create its own schema, and have the ability to create
 /drop triggers and functions. Refer to your SQL Server documentation for how to do this.
- If you use Oracle, you must use Oracle 12c Release 2, there are known problems with Release 1 (which is no longer supported by Confluence).

Data Collection Policy

Why does Confluence collect usage data?

We're proud that Confluence is one of the most versatile collaboration tools on the planet, and we will continue to deliver innovative new features as quickly as we can. In order to prioritize the features we deliver, we need to understand how our customers use Confluence, what's important, what's not, and what doesn't work well. The collection of usage data allows us to measure the user experience across many thousands of users and deliver features that matter.

What data is collected?

The type of data we collect is covered inour Privacy Policy. Please read it - we've tried to avoid legal jargon and made it as straightforward as possible.

To view a sample of data that might be collected from your specific installation, go to General Configuration > Analytics.

Data is always collected in Confluence Cloud.

How is data collected from Confluence?

Older versions of Confluence (prior to Confluence 5.6 or Confluence Questions 1.0.618) didn't collect usage data. Analytics are collected using the Atlassian Analytics system app. The app collects analytics events in a log file which is located in <confluence-home>/analytics-logs. The logs are periodically uploaded using an encrypted session and then deleted. If Confluence is unable to connect to the Internet, no logs are ever uploaded.

Enabling/disabling data collection in Confluence

You can turn off analytics collection at any time. Go to General Configuration > Analytics.

Administering Collaborative Editing

Collaborative editing takes teamwork to the next level. This page covers everything you need to know about administering collaborative editing.

Head to Collaborative editing to find out how your team canwork together in real time on software requirements, meeting notes, retros, and any other Confluence page you can think of.

About Synchrony

Collaborative editing is powered by Synchrony which synchronizes data in real time. Confluence manages Synchrony, so administrators should rarely need to interact with it directly.

Synchrony runs on port 8091 by default, and an internal Synchrony proxy means that you shouldn't need to open this additional port.

How you connect to Synchrony will depend on your environment, and your Confluence license. See Pos sible Confluence and Synchrony Configurations.

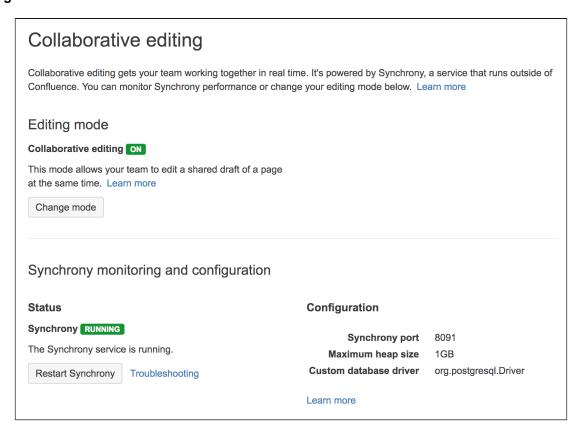
On this page:

- About Synchrony
- Change the editing mode
 - What happens to existing drafts when the mode changes?
- Maximum editor limit
- Auditing considerations
 - No version history in unpublished drafts
 - Visibility of edits made by anonymous users
- Proxy and SSL considerations
 - SSL
 - Proxies
 - WebSockets
- Change your Synchrony configuration
- Start and stop Synchrony
- Monitor Synchrony
- Accessing Synchrony logs
- Managing Synchrony data

Related pages:

• Troubleshooting Collaborative Editing

To see your collaborative editing and Synchrony setup, head to General Configuration>Collaborative editing.



The editing mode determines the editing experience for all users in your site. This is how you turn collaborative editing on or off.

To change the editing mode:

- 1. Go to Co General Configuration>Collaborative editing.
- 2. Choose Change mode.
- 3. Select either On or Off and choose Change.

Changing the editing mode is not trivial, so it is good to understand the implications of each mode.

Mode	Implications
On	This mode allows your team to edit a shared draft of a page at the same time, and see each others' changes in real time.
	This is the recommended editing mode.
Off	This mode means that your team can only edit their own personal draft of a page. Confluence will attempt to merge any conflicts on save. Consider turning collaborative editing on for the full experience
	This mode is useful if you are unable to run Synchrony successfully in your environment, or if you have decided that collaborative editing is not for you (for example if you have auditing requirements that would prohibit using collaborative editing just yet).
	It's a good idea to prompt your users to publish any shared drafts before you turn collaborative editing off, as they will not be able to resume editing existing shared drafts or unpublished changes.

What happens to existing drafts when the mode changes?

Users can always access any existing personal drafts and shared drafts from the **Drafts** page in their profile. Whether they can resume editing the draft depends on the editing mode.

When collaborative editing is **ON**,users will be able to discard or resume editing any personal or shared drafts. A personal draft will be converted to a shared draft when a user resumes editing.

When collaborative editing is **OFF**,users will be able to discard or resume editing any personal drafts. They can't resume editing existing shared drafts, but can view and copy the contents of those drafts.

Shared drafts will only appear in a user's Drafts page if, when collaborative editing was on, they:

- created a draft, and never published it
- edited a published page, and did not publish their changes.

Maximum editor limit

A maximum of 12 people can edit a page at the same time. This means that people can't enter the editor if there are already 12 other people editing the page, and will need to wait until someone leaves.

Administrators can increase or decrease this limit using a system property. If you experience performance issues when many people are editing, you might want to decrease this limit.

Auditing considerations

We know that auditing is a major consideration for some customers. We don't yet have very granular auditing capabilities with collaborative editing. All page changes are currently attributed to the person that publishes the page, rather than the people who made each specific change.

If this is going to be a problem in your site, we recommend turning collaborative editing off in your site for now.

No version history in unpublished drafts

We're saving all the time in collaborative editing, but we don't save versions of unpublished changes. When restoring an earlier page version, you can only roll back to an existing published version. Any unpublished changes will be lost when you restore a previous version.

Visibility of edits made by anonymous users

There are some additional things to be aware of if you have granted the **Add**page permission (and **Can use** global permission) to anonymous users.

You won't be alerted, when closing the editor or publishing a page, if anonymous users made the only unpublished changes on the page. This means a logged in user may inadvertently publish changes they were not aware had been made to the page.

The changes themselves are visible in the page, but the usual warning dialog will not appear if the only people to have made changes were not logged in.

If there are unpublished changes from both logged in users and anonymous users, the warning dialog will appear, but only the logged in users will be listed in the dialog. Changes made by all users (including anonymous) will be included if you view the changes from that dialog.

Proxy and SSL considerations

How you connect to Synchrony will depend on your environment. We know that most Confluence sites run behind a reverse proxy, often with SSL. Here's some information to help you identify the right configuration for your environment, and any changes you might need to make to your environment to use collaborative editing in your site.

SSL

Synchrony runs in a separate JVM, and does not support direct HTTPS connections. If you are not using a reverse proxy, SSL should be terminated at Tomcat. If you are using a reverse proxy or load balancer, SSL should be terminated at your reverse proxy or load balancer.

SeePossible Confluence and Synchrony Configurations for detailed diagrams and examples.

Proxies

If you run Confluence behind a reverse proxy, you should take a look at the Possible Confluence and Synchrony configurations for guidance on how your Confluence and Synchrony setup may impact your proxy.

SeePossible Confluence and Synchrony Configurations for detailed diagrams and examples, plus links to example proxy configuration files.

WebSockets

For best results, your load balancer and proxies should allow WebSocket connections. If your users cannot get a WebSocket connection, Confluence will fall back to an XML HTTP Request (XHR), allowing them to edit pages successfully.

XHR fallback is enabled by default, but can be disabled using asystem property(passed to Confluence)if necessary. You shouldn't need to change this.

Change your Synchrony configuration

You can't change your Synchrony configuration through the Confluence UI. In most cases you shouldn't need to make changes to the default configuration.

If you need to change the port Synchrony runs on or the maximum memory available, for example, you can do this using a system property, or in your start-synchrony script (if you're running your own Synchrony cluster).

SeeConfiguring Synchronyfor more information.

Start and stop Synchrony

If Synchrony is**managed by Confluence**(recommended), Confluence will automatically start Synchrony for you when it starts up. You can also restart Synchrony from the collaborative editing admin screen in Confluence.

If you're running Synchrony standalone in a cluster, you'll use the start-synchrony.shorstart-synchrony.bat. scripts on each Synchrony node. A process ID (PID) file will be created in your synchrony directory.

Stop Synchrony the same way, usingstop-synchrony.shorstop-synchrony.bat. This will destroy the PID file that the start script created in your Synchrony directory. If you've customized the location for storing the PID file in the start-synchronyscript, you'll need to also update this in the stop-synchronyscript.

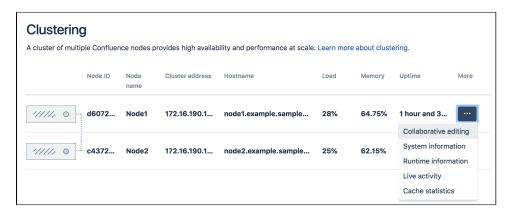
If you're unable to start Synchrony, check that there isn't an existing PID file in your Synchrony directory.

Monitor Synchrony

To check if Synchrony is running, go to General Configuration>Collaborative editing.

If you're running Confluence in a cluster, you can check the status of Synchrony on each node from the clustering screen.

Go to General Configuration Clustering, then on each node choose Collaborative editing. You can access all nodes in this way, you don't need to hit a specific node in your browser.



From here you can see the Synchrony status, mode, and URL Confluence is using to connect to it. Here's what it looks like when Synchrony is managed by Confluence.



All Confluence nodes must use the same Synchrony mode. For example, you can't have one node using managed Synchrony, and another node connecting to a standalone Synchrony cluster.

Accessing Synchrony logs

If Synchrony ismanaged by Confluence(recommended), Synchrony logs will be stored in your<local-home>/logs directory, with the Confluence application logs.

If you're running **Synchrony standalone in a cluster**, your Synchrony logs will be stored in the Synchrony directory on each Synchrony node (wherever you run the start and stop scripts from).

To learn how to change the logging level, seeConfiguring Synchrony.

Managing Synchrony data

Each page and blog post has its own Synchrony change log, which contains a graph of all edits to that page or blog post. In busy Confluence sites the database tables that store the Synchrony change logs can grow very quickly. Because the change logs store all changes as they happen, they may retain personally identifiable information, even after the page they relate to has been deleted.

We provide two scheduled jobs for removing Synchrony data:

- Synchrony data eviction (soft)
- Synchrony data eviction (hard).

The soft eviction job runs regularly in the background. The hard eviction job is available for when you need to remove Synchrony data more aggressively, and is disabled by default.

SeeHow to remove Synchrony datafor more information on how these jobs work.

Possible Confluence and Synchrony Configurations

Synchrony is the engine that powers collaborative editing in Confluence.

There are a few different options for running Synchrony, and it is worth taking some time to determine which option will best meet the needs of your organisation.

On this page:

- Possible configurations for Confluence Server
- Possible configurations for Confluence Data Center

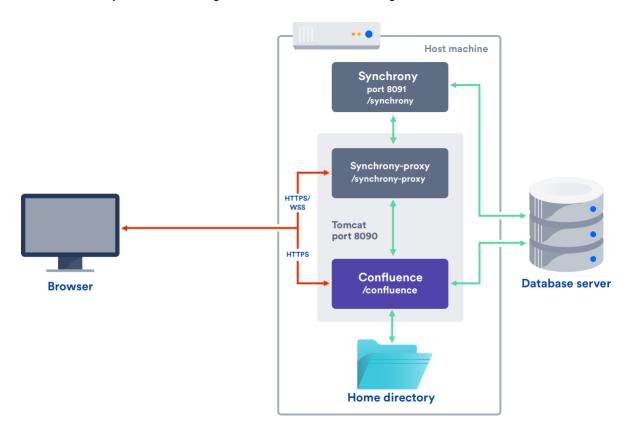
Possible configurations for Confluence Server

If you have a Confluence Server license, Synchrony ismanaged by Confluence.

The diagrams below show examples of a common implementation where Confluence is running under the /confluence context path (e.g. www.mysite.com/confluence). The concepts are the same if you use Confluence without a context path (e.g.www.myconfluence.com).

No reverse proxy

If you don't run Confluence behind a reverse proxy, you'll connect to Synchrony via Confluence's internal Synchrony proxy. SSL, if used, is terminated at Tomcat. This is the default configuration, and you shouldn't need to make any additional changes to use collaborative editing.

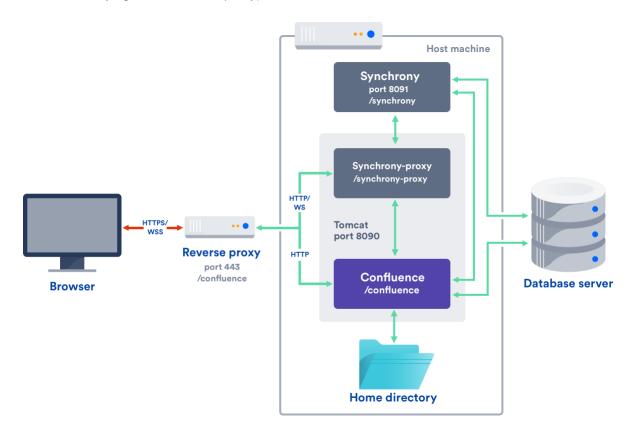


With a reverse proxy

If you run Confluence behind a reverse proxy, you will connect to Synchrony via Confluence's internal Synchrony proxy. This is the default configuration with a reverse proxy, and a good choice if you do not want to open port 8091. SSL should be terminated at your reverse proxy.

You do not need to make any additional changes to your reverse proxy configuration for Synchrony, but for best results your reverse proxy must support WebSocket connections (you may need to manually enable this in your proxy).

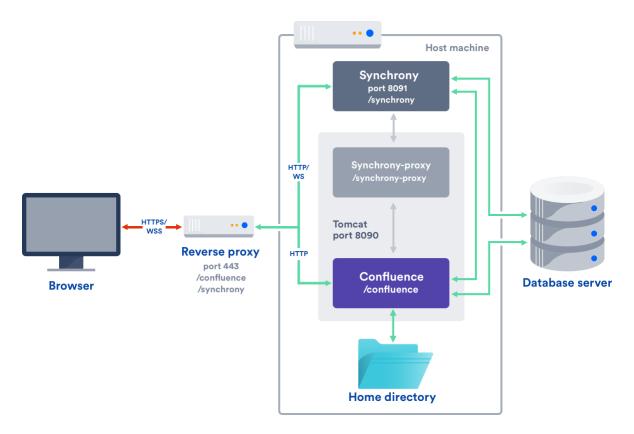
To tell Confluence that you want to use the internal proxy, set the synchrony.proxy.enabledsystem propertytotrue. (This is optional, but will prevent Confluence from trying to reach Synchrony via /synchrony first, before retrying via the internal proxy).



If Synchrony can't be reached via /synchrony-proxy we'll automatically try /confluence/synchrony-proxy (where /confluence is your Confluence context path).

Direct to Synchrony with a reverse proxy

If you run Confluence behind a reverse proxy, and experience latency or other issues connecting to Synchrony via Confluence's internal Synchrony proxy, you can choose to connect direct to Synchrony. This is the optimal setup, but does require some changes to your environment. You will need to open port 8091 and add /synchrony to your reverse proxy configuration.SSL will still be terminated at your reverse proxy, as Synchrony does not accept direct HTTPS connections.



If Synchrony can't be reached via /synchrony we'll automatically try the internal Synchrony proxy via /confluence/synchrony-proxy (where /confluence is your Confluence context path).

See the following guides for example reverse proxy configurations. The order of directives is important, so check our examples.

- Using Apache with mod_proxy
- Running Confluence behind NGINX with SSL
- Proxying Atlassian server applications with Apache HTTP Server (mod proxy http)
- Proxying Atlassian server applications with Microsoft Internet Information Services (IIS)
- How to configure Amazon Web Service Elastic Load Balancer with Confluence

Possible configurations for Confluence Data Center

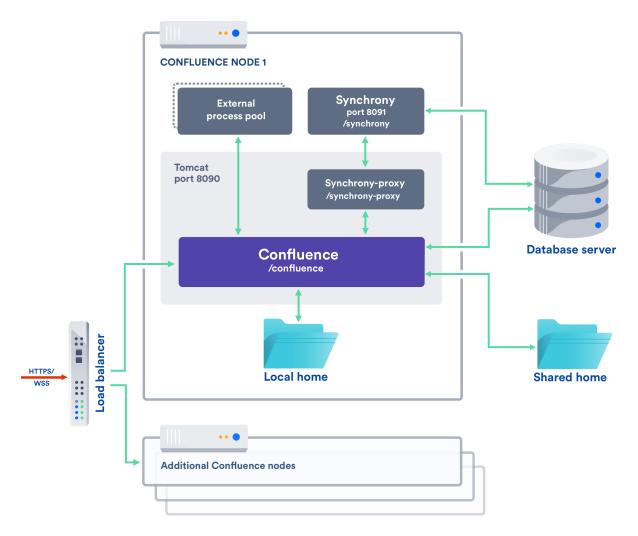
If you have a Confluence Data Center license, two methods are available for running Synchrony:

- managed by Confluence(recommended)
 Confluence will automatically launch a Synchrony process on the same node, and manage it for you.
 No manual setup is required.
- Standalone Synchrony cluster (managed by you)
 You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need.
 Significant setup is required. During a rolling upgrade, you'll need to upgrade the Synchrony separately from the Confluence cluster.

If you want simple setup and maintenance, we recommend allowing Confluence to manage Synchrony for you. If you want full control, or if making sure the editor is highly available is essential, then managing Synchrony in its own cluster may be the right solution for your organisation.

Managed by Confluence

Here's a simplified view of the architecture when Synchrony is**managed by Confluence**. This is the recommended approach, as no manual set up, or ongoing upgrades are required - it works right out of the box.

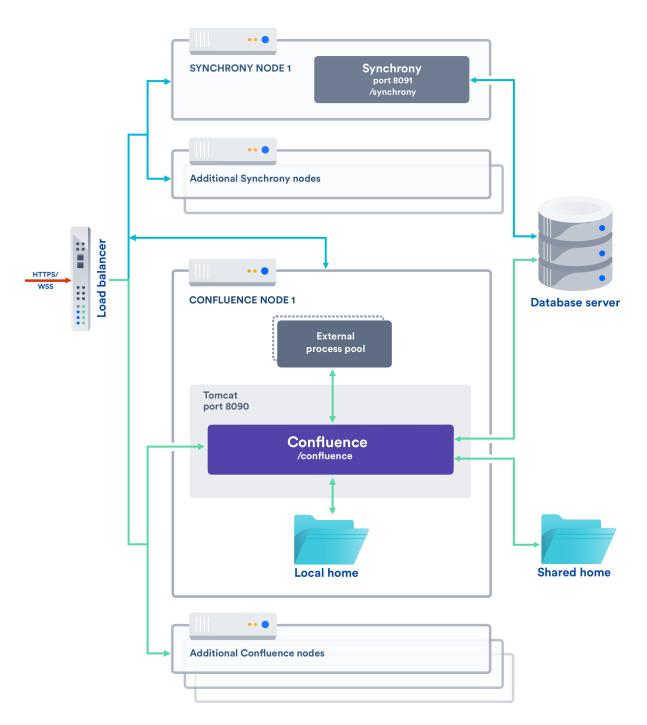


If you choose this approach, the guidance provided for Confluence Server above, also applies to your load balancer and any proxies.

This option is only available in Confluence 6.12 and later.

Standalone Synchrony cluster

If you choose to**manage Synchrony yourself**, the architecture looks more like this. Again the diagram has been simplified, and doesn't show communication between nodes.



If you choose this approach you will need to:

- Set up and manage multiple Synchrony cluster nodes.
- Always start Confluence with the synchrony.service.urlsystem property (this tells Confluence where to find your Synchrony cluster, instead of launching a Synchrony process on the current node).
- Open the Synchrony port (8091), as the Synchrony proxy is never used when you manage Synchrony yourself.
- Terminate SSL at your load balancer. Synchrony cannot accept HTTPS connections.
- Upgrade Synchrony manually, each time you upgrade Confluence.

In most cases, two Synchrony nodes will be adequate for multiple Confluence nodes.

For a step-by-step guide to setting up your Synchrony cluster, seeSet up a Synchrony cluster for Confluence Data Center.



i If you enabled collaborative editing prior to Confluence Data Center 6.12, standalone Synchrony will be your default setup.

If you would prefer a less complex setup, seeMigrate from a standalone Synchrony cluster to managed Synchronyto find out how to allow Confluence to manage Synchrony for you.

Configuring Synchrony

Synchrony is the engine that powers collaborative editing in Confluence.

There's no UI for configuring Synchrony. Configuration changes, such as changing the Synchrony port or memory settings, are made via system properties. How you pass these properties depends on whether Synchrony is managed by Confluence, or deployed as a seperate cluster.

In most cases, Synchrony is managed by Confluence.

If you have a Data Center license, you may choose to deploy and manage **S ynchrony standalone in a cluster**, instead of allowing Confluence to manage Synchrony for you. SeePossible Confluence and Synchrony Configurations for more information.

On this page:

- Passing recognized system properties to Synchrony
- Common configuration changes
- Change the logging level for managed Synchrony
- Change the logging level for Synchrony standalone
- Troubleshooting

Passing recognized system properties to Synchrony

If Synchrony is **managed by Confluence** (the most common setup), you make changes to Synchrony by passing system properties to Confluence. SeeConfiguring System Properties to find out the best way to do this for your operating system.

You can find a full list of system properties at Recognized System Properties.

If you're running **Synchrony standalone in a cluster**, you pass properties directly to Synchrony via the start-synchrony scripts.

Note that the properties are not always the same as those used when Synchrony is managed by Confluence. A full list of required and optional properties can be found atSet up a Synchrony cluster for Confluence Data Center.

Passing JVM arguments to Synchrony

Sometimes you may want to pass additional arguments, that are not already provided by a system property, directly to Synchrony's JVM.

If Synchrony is**managed by Confluence**, you will need to create a file called synchrony-args. propertiesin your home directory (or shared home if you have a Data Center license) and include the arguments you want to pass to Synchrony, one per line, as follows:

property1=value1
property2=value2

Remember, you can't use this method for passing any value that is already handled by a system property, such port, Xmx or Xss etc.

If you're running **Synchrony standalone in a cluster**, you pass arguments to Synchrony's JVM directly, by adding them to your start-synchrony script, in the **Optional Overrides** section.

Common configuration changes

The two most common changes people make to Synchrony is to change the port that Synchrony runs on, if port 8091 is already in use, and to change the maximum heap memory allocated to Synchrony.

Change the port Synchrony runs on

Synchrony runs on port 8091 by default. If this port is already in use by another application on your server you can use the thesynchrony.port system property to change it to an available port.

If you're running Confluence 6.0.3 or earlier you'll need to use reza. port instead of synchrony.port.

To change the maximum heap for Synchrony

Synchrony has a maximum heap size of 2 GB by default.

If you experience out of memory errors related to Synchrony, you can change the heap size allocated to Synchrony using the synchrony.memory.maxsystem property.

If you're running Confluence in a cluster, you may want to increase the maximum heap size to 4gb on each node.

Change the logging level for managed Synchrony

The logging level for managed Synchrony is set to INFOby default. If you find this too verbose, you can decrease the logging level to WARNOT ERROR.

To change the managed Synchrony logging level:

Create a file called synchrony-log4j.properties with the following content:

```
log4j.rootLogger=WARN, stdout
log4j.appender.stdout=org.apache.log4j.ConsoleAppender
log4j.appender.stdout.Target=System.out
log4j.appender.stdout.layout=org.apache.log4j.PatternLayout
log4j.appender.stdout.layout.ConversionPattern=%d %p [%t] [%c{4}] %m%n
```

In this example we'll set the logging level to WARN. Replace this with ERRORif you only want to log errors.

- 2. Save the file. You can place the file anywhere, but we recommend your home directory (or shared home) alongside the synchrony-args.propertiesfile.
- 3. Edit your<home-directory>/synchrony-args.propertiesfile. If you're running Confluence in a cluster, this will be in your shared home directory.
- 4. Add the following line to tell Synchrony where to find your log configuration.

```
log4j.configuration=file://<path-to-file>/synchrony-log4j.properties
```

Replace <path-to-file>with your file path. In Linux this will be something like=file:///var/confluence/local-home/synchrony-log4j.properties, for example.

5. In Confluence, go to General Configuration Collaborative editing and select Restart Synchronyto pick up the changes.

Exclude the Confluence DEBUG prefix

Because Synchrony is managed by Confluence, the Synchrony logs include a prefix with information from Confluence itself. You can omit this prefix to make the logs easier to read.

To omit the Confluence DEBUG prefix from the Synchrony logs:

- 1. Edit the <install-directory > /confluence / WEB-INF / classes / log4j. properties file.
- 2. Change thelog4j.appender.synchronylog.layout.ConversionPatternline to remove%d % p [%t] [%c{4}]as follows:

```
log4j.appender.synchronylog.layout.ConversionPattern=%m%n
```

3. Save the file, then restart Confluence to pick up the changes.

If you're running Confluence in a cluster, you'll need to repeat this process on each Confluence node.

Change the logging level for Synchrony standalone

If you choose to deploy and manage **Synchrony standalone in a cluster**, you can configure the logging levelin your start-synchrony script.

To change the Synchrony standalone logging level:

1. Create a file called synchrony-log4j.properties with the following content:

```
log4j.rootLogger=WARN, stdout
log4j.appender.stdout=org.apache.log4j.ConsoleAppender
log4j.appender.stdout.Target=System.out
log4j.appender.stdout.layout=org.apache.log4j.PatternLayout
log4j.appender.stdout.layout.ConversionPattern=%d %p [%t] [%c{4}] %m%n
log4j.category.com.hazelcast=INFO
log4j.category.hazelcast=INFO
```

In this example we want to set the logging level to WARN. Replace this with ERRORIf you only want to log errors. We keep the Hazelcast logging level at INFOso you can still see the Synchrony nodes communicating with each other.

- 2. Save the file. You can place the file anywhere, but we recommend your Synchrony directory.
- 3. Edit your < synchrony-directory > / start-synchrony.sh or start-synchrony.batfile.
- 4. Add the following line in the Optional Overrides section to tell Synchrony where to find your log config:

```
log4j.configuration=file://<path-to-file>/synchrony-log4j.properties
```

5. Restart Synchrony to pick up the changes.

Repeat this process on each Synchrony node.

Troubleshooting

• If you have a Data Center license, and Synchrony is managed by Confluence, we recommend storing the synchrony-args.propertiesfile in the shared home directory, so that all Synchrony processes are started with the same JVM arguments. If you do locate the synchrony-args. propertiesfile in the local home, the arguments will only be passed to the Synchrony process on that node.

Set up a Synchrony cluster for Confluence Data Center

If you have a Confluence Data Center license, two methods are available for running Synchrony:

- managed by Confluence(recommended)
 Confluence will automatically launch a
 Synchrony process on the same node, and manage it for you. No manual setup is required.
- Standalone Synchrony cluster (managed by you)

You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need. Significant setup is required. During a rolling upgrade, you'll need to upgrade the Synchrony separately from the Confluence cluster.

If you want simple setup and maintenance, we recommend allowing Confluence to manage Synchrony for you. If you want full control, or if making sure the editor is highly available is essential, then managing Synchrony in its own cluster may be the right solution for your organisation.

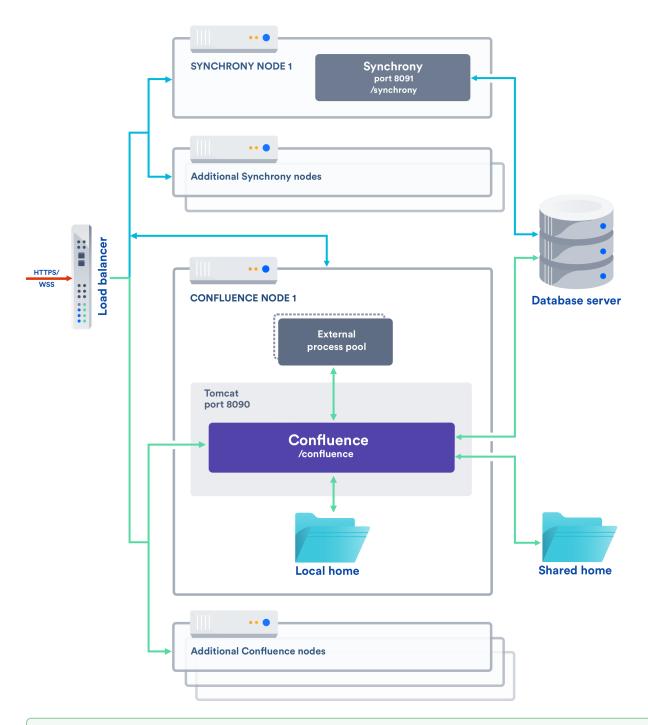
On this page:

- Architecture overview
- Set up a Synchrony standalone cluster
 - 1 Provision your Synchrony nodes
 - 2 Create the Synchrony home directory
 - 3 Edit the start and stop scripts
 - 4 Add additional Synchrony nodes and configure your load balancer
 - 5 Start Confluence one node at a time
 - 6 Enable collaborative editing
- Required properties for Synchrony standalone
- Optional properties for Synchrony standalone
- Run Synchrony standalone in an IPv6 environment
- Run Synchrony standalone as a service
- Provide credentials to Synchrony standalone using environment variables

On this page we'll guide you through the process of setting up a standalone Synchrony cluster, hosted on your own infrastructure. The ability to run your own Synchrony cluster is only available with a Data Center license.

Architecture overview

Here's a simplified view of the architecture when you manage Synchrony yourself, in a seperate cluster. Note that this diagram doesn't show communication between nodes.



If you enabled collaborative editing prior to Confluence Data Center 6.12, standalone Synchrony will be your default setup.

If you would prefer a less complex setup, seeMigrate from a standalone Synchrony cluster to managed Synchronyto find out how to allow Confluence to manage Synchrony for you.

Set up a Synchrony standalone cluster

This page will guide you through setting up a Synchrony standalone cluster on your own infrastructure.

If you're using AWS or Azure, using one of our templates may be a more efficient way to set up Confluence with a standalone Synchrony cluster.

1 Provision your Synchrony nodes

For the purposes of this guide, we assume you have already provisioned the hardware or virtual instances for your Synchrony nodes. We recommend starting with 2 Synchrony nodes.

You should allow 2GB memory for Synchrony, and enough disk space for the Synchrony application and logs.

2 Create the Synchrony home directory

To create the Synchrony directory on your first Synchrony node:

- 1. Grab the <install-directory > /bin/synchronydirectory from one of your Confluence nodes and move it to your new Synchrony node. We'll call this your <synchrony-home > directory.
- 2. Copysynchrony-standalone.jarfrom your Confluence local home directory to your<synchrony -home>directory.
- 3. Copy your database driver from your Confluence<install-directory>/confluence/web-inf /libto your<synchrony-home>directory or other appropriate location on your Synchrony node.

3 Edit the start and stop scripts

We provide scripts to start and stop Synchrony on each node. These need to be edited to add information about your environment:

- 1. Edit the < synchrony-home > / start-synchrony.shorstart-synchrony.batfile
- 2. Enter details for all of the required parameters listed under**Configure parameters**. SeeRequired properties below, for a description of each.
- 3. Enter detail for any optional properties you may want to specify. SeeOptional properties below for a description of each.
- 4. Save the file.
- 5. Start Synchrony by running the start-synchrony script.
- 6. Visithttp://<SERVER_IP>:<SYNCHRONY_PORT>/synchrony/heartbeat to check Synchrony is running.

4 Add additional Synchrony nodes and configure your load balancer

To create your second Synchrony node:

- 1. Copy your<synchrony-home>directory to the second Synchrony node.
- 2. Start Synchrony on that node using the start-synchrony script. As each node joins you'll see something like this in your console.

3. Configure your load balancer for Synchrony traffic.

For best results, your load balancer should allow WebSocket connections. SSL connections must be terminated at your load balancer, as Synchrony can't accept HTTPS requests.

You can choose to use the same load balancer for both Confluence and Synchrony, or two seperate load balancers. When we refer to the Synchrony load balancer, we mean whichever load balancer is handling Synchrony traffic.

4. Make sure the Synchrony port (8091) is open. Ports used by Atlassian Applicationshas a good summary of all ports Synchrony uses in Data Center. This is the only one that needs to be open.

5 Start Confluence one node at a time

Now that Synchrony is running in a cluster, it's time to get Confluence involved. It is essential that you stop Confluence on **all nodes** before continuing.

- 1. Stop Confluence on all nodes.
- Start Confluence on one node with the followingsystem property.
 This property is used to tell Confluence where to find Synchrony, and prevents Confluence fromautomatically launching a Synchrony process on your Confluence node.

-Dsynchrony.service.url=http://<synchrony-load-balancer-url>/synchrony/v1

For example http://42.42.42.42/synchrony/vlor http://synchrony.example.com/synchrony/vl

3. Check that Confluence can connect to Synchrony. Head to ♥ > General Configuration > Clustering then choose • • • > Collaborative editing beside the Confluence node you just started.

The Synchrony mode should be **Standalone Synchrony cluster**.



If the mode is 'Managed by Confluence', your Confluence node is not connected to your Synchrony cluster. Make sure you're passing the Synchrony service URL system property correctly.

4. Repeat this process, starting each Confluence node, one at a time, with the synchrony.service. url.

SeeHow to check the status of Synchrony for Confluence Data Centerfor more info on how to check Synchrony is running.

6 Enable collaborative editing

If you're installing Confluence for the first time, collaborative editing is enabled by default. If you've upgraded from an earlier Confluence version, or have disabled it in the past, collaborative editing may still be disabled.

To enable collaborative editing:

- 1. Head to General Configuration>Collaborative editing.
- 2. Choose Change mode.
- 3. Select On and choose Change.

You can now try editing a page. You'll need to access Confluence via your load balancer. You can't create or edit pages when accessing a node directly.

Any users who had the editor open before you made this change will need to refresh in order to continue editing, as the Synchrony URL they're connected to will have changed.

Required properties for Synchrony standalone

These properties **only apply** when you're running Synchrony standalone in its own cluster. If Synchrony is managed by Confluence (Server or Data Center) these properties don't apply.

The following properties must be provided in the start-synchrony script.

Property name	Description
SERVER_IP	Public IP address or hostname of this Synchrony node. It could also be a private IP address - it should be configured to the address where Synchrony is reachable by the other nodes.

DATABASE _URL	This is the URL for your Confluence database. For example jdbc:postgresql://yourserver:5432/confluence. You can find this URL in <local-home>/confluence.cfg.xml.</local-home>
DATABASE _USER	This is the username of your Confluence database user.
DATABASE _PASSWORD	(Optional)This is the password for your Confluence database user. If your password contains special characters, Synchrony may silently fail to connect to the database.
	Rather than hardcoding your password, we recommend setting your password with the environment variable SYNCHRONY_DATABASE_PASSWORD. Any dots (".") in variable names (identifiers) will need to be replaced with underscores ("_").
CLUSTER_ JOIN_PRO PERTIES	This determines how Synchrony should discover nodes. You'll be prompted to uncomment a set of parameters for either: • TCP/IP • Multicast • AWS Follow the prompts in the script for the values you need to enter for each of these.
	Follow the prompts in the script for the values you need to enter for each of these.
DATABASE _DRIVER_ PATH	This is the path to your database driver file. If you're running Synchrony on its own node, you'll need to copy your database driver to an appropriate location then provide the path to this location.
SYNCHRON Y_JAR_PA TH	This is the path to the synchrony-standalone. jar file you copied to this node.
SYNCHRON Y_URL	This is the URL that the browser uses to contact Synchrony. Generally this will be the full URL of the load balancer Synchrony will run behind plus the Synchrony context path, for example http://yoursite.com:8091/synchrony.
	Note that it $\textit{does not}$ end with $/v1$, unlike the <code>synchrony.service.url</code> system property passed to Confluence. If this URL doesn't match the URL coming from a users' browser, Synchrony will fail.
OPTIONAL _OVERRID ES	You can choose to specify additional system properties. See the table below for recognised Synchrony system properties.

Optional properties for Synchrony standalone

These properties only apply ifyou're running Synchrony standalone in a cluster.

When you start Synchrony, we pass default values for the properties listed below. You can choose to override these values by specifying any of these properties when you start Synchrony.

Property name	Default	Description
cluster. listen. port	5701	This is Synchrony's Hazelcast port. Specify this property if you do not want to use port 5701 or if it is not available. As with the Confluence Hazelcast port (5801) you should ensure that only permitted cluster nodes are allowed to connect to Synchrony's Hazelcast port, through the use of a firewall and or network segregation.

synchron	25500	This is the Aleph binding port. Synchrony uses Aleph to communicate between nodes. Specify this property if you don't want to use the default.
y. cluster. base. port		between nodes. Specify this property if you don't want to use the default.
cluster. join. multicas t.group	224.2.2.3	If the cluster join type is multicast, you can specify an IP address for the multicast group if you don't want to use the default.
cluster. join. multicas t.port	54327	If the cluster join type is multicast, you can specify a multicast port if you don't want to use the default.
cluster. join. multicas t.ttl	32	If the cluster join type is multicast, this is the time to live threshold. The default, 32, means the scope is restricted to the same site, organization or department. Specify this property if you want to use a different threshold.
cluster. join. aws. access. key		If the cluster join type is AWS, this is your AWS access key.
cluster. join. aws. secret. key		If the cluster join type is AWS, you can authenticate by IAM role or Secret key. This is your AWS secret key.
cluster. join. aws.iam		If the cluster join type is AWS, you can authenticate by IAM role or Secret key. This is your AWS IAM role.
cluster. join. aws. region	us-east-1	If the cluster join type is AWS, this is the AWS region your Synchrony nodes will be running in.
cluster. join. aws. security .group		If the cluster join type is AWS, and you want to narrow the members of your cluster to only resources in a particular security group, specify the name of your AWS security group.
cluster. join. aws.tag. key		If the cluster join type is AWS, and you want to narrow the members of your cluster to only resources with particular tags, specify the AWS tag key.
cluster. join. aws.tag. value		If the cluster join type is AWS, and you want to narrow the members of your cluster to only resources with particular tags, specify the AWS tag key value.
cluster. join. aws. host. header		If the cluster join type is AWS, t his is the AWS endpoint for Synchrony to use (the address where the EC2 API can be found, for example 'ec2.amazonaws.com').

cluster. join. aws. timeout	5	If the cluster join type is AWS, this is the joining timeout (in seconds).
cluster. interfac es	Defaults to the same value as SE RVER_IP	This is the network interface Synchrony will use to communicate between nodes. Specify this property if you don't want to use the default, which uses the value of the required property Defaults to the same value as <code>SERVER_IP</code> (also known as <code>synchrony.bind</code>).
synchron y. cluster. bind	Defaults to the same value as SE RVER_IP	This is the Aleph binding address. This should be set to the same value ascluster.interfaces. Specify this property if you did not use the default value for cluster.interfaces.
synchron y.port	8091	This is the HTTP port that Synchrony runs on. If port 8091 is not available, specify this property to choose a different port.
synchron y. context. path	Defaults to the context path of SYN CHRONY_URL	This is the context path for Synchrony. There should be no need to change this.
hazelcas t. prefer. ipv4. stack	True	If you're running Confluence in an IPv6 environment, you will need to set this property to False.

Run Synchrony standalone in an IPv6 environment

If you're running a **Synchrony standalone in a cluster**in an IPv6 environment, you will need to start Synchrony with the following JVM argument:

```
-Dhazelcast.prefer.ipv4.stack=false
```

If you're using the start-synchrony scripts, simply uncomment this line in the script.

Run Synchrony standalone as a service

If you're running **Synchrony standalone in a cluster**, and you'd prefer to run Synchrony as a service on each node, seeRun Synchrony-standalone as a service on Linux.

It's not possible to run Synchrony standalone as a service on Windows. Consider switching to managed Synchrony instead.

Provide credentials to Synchrony standalone using environment variables

If you're running **Synchrony standalone in a cluster**, and you prefer to store sensitive information in your environment, rather than directly in the Synchrony startup scripts you can create asynchronyenvfile, and use it to provide your database credentials. This is only available in Linux environments.

SeeProvide credentials to Synchrony standalone using environment variables (Linux)

Migrate from a standalone Synchrony cluster to managed Synchrony

If you have a Confluence Data Center license, and enabled collaborative editing prior to Confluence 6.12, you will likely be running standalone Synchrony, either in it's own cluster, or manually on each Confluence node.

If you'd prefer a simpler setup, with less ongoing maintenance, you can choose to let Confluence manage Synchrony for you. Confluence will automatically start up a Synchrony process when Confluence is started.

Some Confluence downtime is required for this process.

To switch from managing your own Synchrony cluster to letting Confluence manage Synchrony:

- 1. Configure your load balancer to direct traffic away from all Confluence and Synchrony nodes.
- 2. Stop Confluence and Synchrony on all nodes.
- 3. Remove the thesynchrony.service.urlsystem property. This property tells Confluence where to find your external Synchrony cluster.

The way you remove this system property depends on how you run Confluence. Note that this system property is passed to **Confluence**, not Synchrony itself.

If you start **Confluence manually on Windows**, edit the <install directory>/bin/setenv. batfile and remove the following line:

```
set CATALINA_OPTS=-Dsynchrony.service.url=http://example-synchrony.com/synchrony/v1 %
CATALINA_OPTS%
```

If you start **Confluence manually on Linux**, edit the <install directory>/bin/setenv.shfile and remove the following line:

```
CATALINA_OPTS="-Dsynchrony.service.url=http://example-synchrony.com/synchrony/v1 ${CATALINA_OPTS}"
```

If you're running as a **Confluence as a Windows Service**, you'll need to edit the service and remove the following from the Java options:

```
- \texttt{Dsynchrony.service.url=http://example-synchrony.com/synchrony/v1} \\
```

SeeConfiguring System Properties for a step-by-step guide to passing system properties to Windows services via the command line or Windows Registry.

4. Set the synchrony.memory.maxsystem property to increase the maximum heap memory available to Synchrony to 2gb (or the amount of memory previously allocated to the Synchrony standalone service).

The way you set this system property depends on how you run Confluence. Note that this system property is passed to **Confluence**, not Synchrony itself.

If you start **Confluence manually on Windows**, edit the <install directory>/bin/setenv. batfile and remove the following line:

```
set CATALINA_OPTS=-Dsynchrony.memory.max=2g %CATALINA_OPTS%
```

If you start **Confluence manually on Linux**, edit the <install directory>/bin/setenv.shfile and remove the following line:

```
CATALINA_OPTS="-Dsynchrony.memory.max=2g ${CATALINA_OPTS}"
```

If you're running as a **Confluence as a Windows Service**, you'll need to edit the service and remove the following from the Java options:

-Dsynchrony.memory.max=2g

SeeConfiguring System Properties for a step-by-step guide to passing system properties to Windows services via the command line, Windows Registry, or in AWS.

- Make sure all required ports are open, especiallyespecially5701 and 25500which are used by the Synchrony cluster. SeeConfluence Server and Data Center portsfor a full list.
- 6. Start Confluence on one node.
- 7. In Confluence, edit a page and check that you can successfully make changes.
- 8. Repeat this process on each Confluence node, starting each nodeone at a time.

Once all nodes are back up and running, and you've confirmed that collaborative editing is working as expected, you can decommission your external Synchrony cluster, including removing any startup scripts or services you may have configured.

Any users who had the editor open before you made this change will need to refresh in order to continue editing, as the Synchrony URL they're connected to will have changed.

You may also need to make some changes to your load balancer configuration. SeePossible Confluence and Synchrony Configurations for more information.

Troubleshooting Collaborative Editing

Collaborative editing is powered by Synchrony which synchronizes data in real time. Under normal circumstances it should not need to be managed manually by an administrator.

This page will help you troubleshoot problems with Synchrony in your instance.

First steps

Check Synchrony is running

To check if Synchrony is running, go to Configuration Collaborative editing.

Note: if you're running Confluence Data Center, this page will only be able to tell you if the current Confluence node is connected to your Synchrony cluster. You may want to use a third party monitoring tool to help you monitor your Synchrony cluster. SeeHow to check the status of Synchrony for Confluence Data Centerfor more info.

Check you can edit a page

If you see an error when you edit a page, but Synchrony is running, something is preventing your browser from connecting to Synchrony.

The most common issue is a misconfigured reverse proxy. See our proxy troubleshooting tips later in this page or head toAdministering Collaborative Editingto find out more about possible proxy and SSL configurations.

Check the logs

You can find the Confluence application logs at <home-directory>/logs/atlassian-confluence. logand Synchrony specific logs at <home-directory>/logs/atlassian-synchrony.log.

Restart Synchrony

If Synchrony is managed by Confluence, go to General Configuration Collaborative editing and choose Restart Synchrony.

If you run your own standalone Synchrony cluster, manually restart Synchrony on each node.

Check port 8091 is available

Synchrony runs on port 8091 by default. If this port is already in use by another application on your server you can use the thesynchrony.port system property to change it to an available port.

(If you're using Confluence 6.0.3 or earlier you'll need to use reza.port instead of synchrony.port.)

SeeConfiguring System Properties to find out how to change this.

For Confluence Data Center the way you run Synchrony is a little different. SeeConfiguring Synchrony for more information.

Reverse proxy issues

On this page:

- First steps
 - Check Synchrony is running
 - Check you can edit a page
 - Check the logs
 - Restart Synchrony
- Check port 8091 is available
- Reverse proxy issues
- Forward proxy issues
- Websocket issues
- SSL issues
- Memory issues
- Multiple Synchrony processes
- Mixed Synchrony modes in cluster
- Incompatible browser extensions
- Firewall or anti-virus interference
- Too many people in the editor

Related pages:

Administering Collaborative Editing

If you have configured your reverse proxy, but can't edit pages, here's some things to check in your configuration:

- Go to installation-directory>/conf/server.xml and check the Connectordirective. Make sure that you have correct values for cprotocol>and cproxyName>. See the examples in the guides below for more information.
- Thehttpconnector always needs to be present in the <installation-directory>/conf /server.xmlfile, even if you're configuring SSL or using the AJP connector. The Synchrony health check uses HTTP and will fail if this connector is not present. Alternatively, if you do not want to include the http connector, you can use the synchrony.proxy.healthcheck.disabled system property to disable the health check.
- If you're using Apache, make sure you're using Apache 2.4 (with WebSockets support) and all required modules have been enabled (mod-proxy,mod_rewrite,proxy_wstunnel).
- If you're using Apache and want to connect directly to Synchrony, in your proxy config file, make sure you've included /synchronyand that the order of the Confluence and Synchrony directives and location blocks is correct. See the examples in the guides below for more information.

SeeAdministering Collaborative Editingto find out more about possible proxy and SSL configurations then check out the following guides for more information on how to include Synchrony in your reverse proxy config, if you want to connect direct to Synchrony:

- Using Apache with mod_proxy
- Running Confluence behind NGINX with SSL
- Proxying Atlassian server applications with Apache HTTP Server (mod_proxy_http)
- Proxying Atlassian server applications with Microsoft Internet Information Services (IIS)
- How to configure Amazon Web Service Elastic Load Balancer with Confluence

Forward proxy issues

If you're using a forward or outbound proxy, you will need to add the IP that Synchrony listens on to your config to ensure it is bypassed. SeeConfiguring Web Proxy Support for Confluencefor more info.

By default, the IP is 127.0.0.1, or it will be the value of the synchrony. hostsystem property, if you've customized the hostname or IP that Confluence uses to connect to Synchrony.

Websocket issues

Collaborative editingworks best with a WebSocket connection. If one can't be established due to a timeout, or a proxy server or firewall that doesn't allow WebSocket connections, the editor will attempt to connect via an XML HTTP Request (XHR).

You can usehttp://websocket.org/echo.htmlto perform a quickHTML5 WebSocket test against an echo server.

It's possible to verify the websocket connection using the developer tools in your browser. In the following example, we'll use Chrome to check the websocket connection:

- 1. Open the Chrome**Developer Tools(Shift+CTRL+** J)
- 2. Select the Network tab
- 3. Select the WS filter to show only WebSocket traffic
- 4. Edit a Confluence page
- 5. Select the websocket connection:ws://<confluence-url>/synchrony-proxy/v1/bayeux-sync1
- 6. Select the Frames/Messages tab to see the traffic between browser and the WebSocket.

SSL issues

Synchrony cannot accept direct HTTPS connections, so you will need to terminate SSL at your reverse proxy or load balancer, or at Tomcat if you are not using a reverse proxy.

Memory issues

If you experience out of memory errors related to Synchrony, you can change the heap size allocated to Synchrony using the synchrony memory maxsystem property.

If you're Confluence 6.0.3 or earlier you'll need to use reza.memory.max instead of synchrony.memory. max.

SeeConfiguring System Properties to find out how to change this.

For Confluence Data Center the way you run Synchrony is a little different. SeeConfiguring Synchrony for more information.

Multiple Synchrony processes

If you see an error immediately in the editor, but Confluence reports that Synchrony is running, check to make sure that you only have one Synchrony process running.

If you do have multiple Synchrony processes running, stop Confluence, kill the additional Synchrony processes and then restart Confluence.

You can avoid this problem by always using stop-confluence.sh/stop-confluence.batto stop Confluence, rather than simply closing the Tomcat window.

Mixed Synchrony modes in cluster

If you're running Confluence in a cluster, all of your Confluence nodes must connect to Synchrony in the same way.

If users are able to use collaborative editing on one Confluence node, but not on another Confluence node, go to Configuration Clustering, then on each node choose Collaborative editing.

Make sure all of your Confluence nodes are reporting the same Synchrony mode - either Managed by Confluence, or Standalone Synchrony cluster.



You can access all nodes in this way, you don't need to hit a specific node in your browser.

Incompatible browser extensions

Some third party browser extensions that interact with the editor, such as Grammarly,may not function correctly with collaborative editing. SeeConfluence Collaborative Editing blocks Grammarly Extensionto find out how to disable Grammarly for just your Confluence site.

Firewall or anti-virus interference

We've had a few reports of firewalls or anti-virus software blocking some requests to the server, resulting in unexpected behavior in the editor. You may need to add Confluence to your whitelist / trusted URLs if you experience issues. SeeWeird Page or Editor Behaviors with Kaspersky Internet Securityfor more information.

Too many people in the editor

A maximum of 12 people can edit a page at the same time. This means that people can't enter the editor if there are already 12 other people editing the page, and will need to wait until someone leaves.

Administrators can increase or decrease this limit using a theconfluence.collab.edit.user.limitsys tem property.

Using read-only mode for site maintenance

This feature is available with a Confluence Data Center license.

If you need to perform maintenance while Confluence is still running, or if you're preparing to migrate to a new site, you can put your site into read-only mode to limit what users can do. Your users will be able to view pages, but not create or change them.

Turn on read-only mode

You need System Administrator global permissions to do this.

To enable read-only mode:

- 1. Go to General Configuration Maintenance
- 2. In the Read-only mode section, choose Edit.
- SelectRead-only mode.
- Update the wording of thebanner message, if you'd like to provide a customised message.
- 5. ChooseSave.

The banner message will display above the header on all pages in your site. It's not possible to disable this banner while read-only mode is enabled, but you can customise the message, for example to let your users know when you expect the maintenance to be complete.

It's also possible to turn on the banner before you enable read-only mode. This can be helpful if you want to warn users that you'll be doing some maintenance later that day.

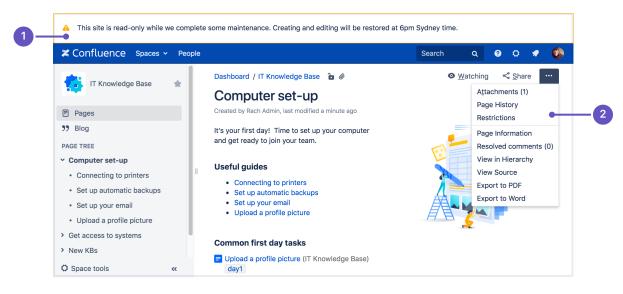
Impact of read-only mode on your site and database

Read-only mode limits the actions that an end user can perform. Some operations may still write to your database, but for the most part people will be unable to make any changes.

While read-only mode is on, youwon't be able to:

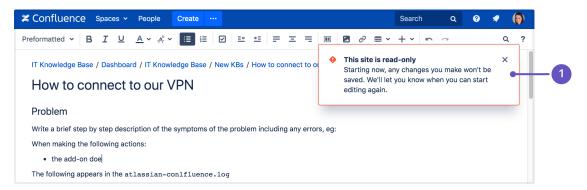
- Create, edit, rename, move, delete or otherwise interact with pages.
- · Create, delete or rename spaces.
- Access most space tools, including reorder pages, make changes to the look and feel, or add integrations.

Here's how a page looks when read-only mode is enabled:



- 1. **Customizable banner** the banner appears on all pages in your site. Admins can customize the message to let you know when the site will be available again.
- 2. **Options are limited** we hide buttons and menu items that are not available, including create, edit, move, and delete.

If you happen to be in the editor at the point read-only mode is enabled, you'll be able to keep typing, but any further changes won't be saved.



Read-only warning - although you can keep typing in the editor (including comment fields), changes you
make after read-only mode is enabled won't be saved. It's best to stop editing at this point.

While read-only mode is on, people with system administrator global permissions will be able to perform some administrative functions, such as:

- Install, uninstall, enable, disable system and user installed apps
- Manage users, groups, and permissions
- Change the site appearance
- Export and import spaces
- Change logging levels, and other configuration.

Not all admin features will be available, and just like end-users, admins won't be able to create, edit, or delete any content.

People with Confluenceadministratorglobal permissions will also be able to perform some administrative functions, but they won't be able to make changes to space permissions while read-only mode is enabled.



It's important to note that read-only mode**does not prevent data being written to the database**, but will significantly limit the changes that can be made.

If you're doing database maintenance, and need to make sure that absolutely nothing is written to the database during that time, it may be best to stop Confluence, rather than using read-only mode.

User-installed app compatibility

Not all apps (also known as plugins or add-ons) are compatible with read-only mode, and may continue to allow users to create or update content while read-only mode is enabled.

To check if your apps are compatible:

- Go to O Seneral Configuration Maintenance
- Check whether any of your user-installed apps are listed as incompatible.

If an app is incompatible, you may want to disable it while you perform maintenance, to avoid users being able to create content via the app.

If you've developed your own custom apps, seeHow to make your app compatible with read-only modeto find out how to test your app and mark it as compatible.

Ways you might use read-only mode

If you're excited by the possibilities of read-only mode, but not sure when you might use it, here are some examples.

Upgrading Confluence

The way you upgrade Confluence hasn't changed, but read-only mode can help you minimize the impact on your organization.

If some downtime is acceptable, the simplest option is to enable read-only mode while you perform the preupgrade steps, such as checking Marketplace app compatibility and backing up your file system and database (if your database supports online backups). This helps you keep the overall downtime to a bare minimum, as users can view pages right up to the point you need to stop Confluence.

If you need to provide uninterrupted access, the approach you take may depend on whether Confluence is running on virtualized or physical hardware.

- If virtualized, you might want to take a 'move forwards' approach. You could enable read-only on your
 production site, clone your database, install, and home directories, then upgrade the clone. Once the
 upgrade is complete and you've validated that everything is working fine, you candirect traffic to the
 upgraded site, and tear down the old site.
- If you're running Confluence on physical hardware it might be more appropriate to create a temporary read-only site. You couldclone your production database, install, and home directories to create a temporary read-only site (similar to the process involved in creating a staging site), and direct traffic to that sitewhile you upgrade your production Confluence site in place.

You should also always test the upgrade on a staging or test instancefirst. As when creating astaging site, it's essential to make sure Confluence is always pointing to the correct database and home directory.

Upgrading your infrastructure

Need to move Confluence to another server, or provision more space for your shared home directory? The approaches outlined above for upgrading Confluence can also be useful when upgrading parts of your infrastructure.

Note that some data may still be written to the database while read-only mode is enabled, so if you're doing database maintenance of any sort, directing your users to a secondary site (with a copy of your database) that has read-only enabled, may be a good approach. You can't, for example, upgrade your production database while Confluence is still running, even if read-only mode is enabled.

Again, always make sure Confluence is pointing to the correct database!

Consolidating multiple confluence sites

It's quite common for multiple Confluence sites to pop up in big organisations. If you're consolidating or merging sites, read-only mode can help limit changes to content while you work through the process of exporting spaces and importing them into your new site.

Administering the Atlassian Companion App



This page applies to Confluence Server and Data Center

The Companion app will be removed from Confluence Cloud at the end of March 2022. Read announcement

We have no plans to end support for the Companion app in Confluence Data Center and Server.

The Atlassian Companion app enables users to edit Confluence files in their preferred desktop application, then save the file back to Confluence automatically.

The download and re-upload of files is managed by the Atlassian Companion app, which needs to be installed on each user's machine (not in the Confluence installation directory) to enable file editing.

On this page:

- Download and install the Atlassian Companion app
- Compatibility with virtual desktop environments
- Recover edited files
- How to delete the Cache folder
- Disable file editing
- Alternatives to the Atlassian Companion app

Download and install the Atlassian Companion app

To edit files, users need to install the Atlassian Companion app and have it running in the background. The first time a user clicks the Edit button in file preview, we prompt them to download and install the app. SeeEdit Filesfor details.

If your users aren't able to install applications themselves, you may want to distribute the app to them or deploy using the Microsoft Installer.



Download the latest Companion version

Download the Atlassian Companion app for Mac or Windows.

Single sign-on considerations

If you've configured single sign-on (SSO) in such a way that your reverse proxy redirects the requests to your SSO gateway, and only successfully authenticated requests ever reach Confluence, your users won't be able to edit files using the Atlassian Companion app. This is because the Atlassian Companion app uses JWT tokens to authenticate requests, and only Confluence can authenticate these requests, not your SSO authenticator.

To make sure requests from the Atlassian Companion app can be authenticated, you should configure your reverse proxy to always allow requests from the following URLs:

- <base-url>/rest/token-auth/api/*
- <base-url>/download/token-auth/attachments/*
- <base-url>/plugins/servlet/imgFilter*
- <base-url>/rest/analytics/1.0/publish/bulk (only necessary if you haveopted in to data collection)

If an unauthenticated user tries to access these URLs directly, they would be redirected to the Confluence login screen. The wouldn't be able to access any content or download files while unauthenticated.

There's a known issue where the token-authpath is not included in the download URL that Confluence provides Companion. See

CONFSERVER-63189 - Companion App does not perform download request using token-auth endpoint CLOSED

Content security policy (CSP) considerations

If you have arestrictive content security policy, your browser will refuse to launch companion, and you'll see a content security policy error in the browser console. This error occurs becauseConfluence 7.3 and later uses a hidden iframe to attempt to launch Companion's custom protocol (atlassian-companion). To resolve this problem you will need to add atlassian-companion: to the default-srcor frame-src list. For example:

frame-src atlassian-companion:;

The content security policy is most commonly configured in your reverse proxy.

Install the Companion app via Microsoft Installer (MSI)

We also provide a Microsoft Installer package (.msi file) to deploy the Atlassian Companion app for Windows across multiple users or machines. By default, the Companion app installs to the Program Files directory, but you can customize this.

Download the Atlassian Companion MSI (69 MB)

If the link above downloads an .exe file instead of the MSI, copy the URL below into your browser to download the file.

 $\label{lem:https://update-nucleus.atlassian.com/Atlassian-Companion/291cb34fe2296e5fb82b83a04704c9b4/latest/win32/atlassian \end{align*} 20 \end{align*} $$ (ia 32/Atlassian \end{align*} 20 \end{align*} $$ (ia 32/Atlassian \end{align*} 20 \end{align*} $$ (ia 32/Atlassian \end{align*} $$ (ia 32/A$

Use the Microsoft Installerto install the Companion app for all users on a given computer:

msiexec /i Atlassian Companion.msi COMPANION_TRUSTED_DOMAINS=https://confluence.atlassian.com; https://support.atlassian.com; /qb ALLUSERS=1"

If you deploy using the Microsoft Installer, the Companion app wont automatically get the latest updates, including security and bug fixes, so some maintenance is required.

We may update the Companion app before or after we release a new version of Confluence. Check the Atlas sian Companion app release notes to make sure you're on the latest version.

Standard install switches

The Microsoft Installer supports standard install switches. For example, install with the TARGETDIR and APPL ICATIONROOTDIRECTORY parameters to change the installation directory:

msiexec /i "Atlassian Companion.msi" TARGETDIR="C:\Users\Emma\AppData\Local\Companion"
APPLICATIONROOTDIRECTORY="C:\Users\Emma\AppData\Local\Companion" /qb

Set trusted domains

In Companion 1.2.0 and later, set your Confluence URL as a trusted domain so users dont have to select 'Trust this domain' when they edit a file for the first time.

System administrators have two options for setting trusted domains/sites before rolling out the Companion app to all users. Either set an environment variablecalled COMPANION_TRUSTED_DOMAINSON each user's computer, or pass the parameter COMPANION_TRUSTED_DOMAINSto the Microsoft Installer (MSI). Set multiple trusted domains by using semicolons (;) as separators.

To set trusted domains when installing using the MSI:

Compatibility with virtual desktop environments

From Confluence 7.3 onwards, Atlassian Companion app should work in most session-based virtual desktops.

Recover edited files

When a user edits a file, that file is also downloaded and saved to the Atlassian Companion folder on their computer. Files are cleared every time the Companion app restarts.

Follow our guide toaccessing Confluence files edited with the Atlassian Companion app.

How to delete the Cache folder

If youd like to free up disk space, its safe to manually delete the cache folder. Deleting individual files in the cache folder may cause errors, so you should delete the entire Cache folder. If the cache folder is locked while Companion App is running, quit Companion, delete the cache folder, then open Companion.

For Windows, Companion 1.0.0, go to C:\Users\admin\AppData\Roaming\Atlassian Companion\Cache.

For Windows, Companion 1.1.0, go to C:\Users\admin\AppData\Local\Atlassian Companion\Cache.

For Mac, Companion 1.0.0 and 1.1.0, go toHome/Library/Application Support/Atlassian Companion /Cache.

Disable file editing

From Confluence 7.3 onwards, it is not possible to disable file editing completely. However, you can choose to revert to the previous Edit in Office functionality, which disables the Companion app integration.

Alternatives to the Atlassian Companion app

In some versions of Confluence, you can revert to the previousEdit in Officefunctionality. This is a workaround for customers who are unable to use the Companion app in their environment.

To enable the legacy Edit in Office functionality:

- 1. Go to General Configuration>Office Connector.
- 2. ChooseEnable Edit in Office for all users and save your changes.

This will disable Companion app functionality for all users in the site.

Notifications from Atlassian

The Atlassian Notifications system app provides targeted in-app notifications in Confluence.

These notifications are mostly directed at administrators and provide information about things like new Confluence versions, upcoming license renewals, and important security announcements.

To disable these notifications:

- Go to > Manage apps.
 Search for Atlassian Notifications.
- 3. Expand the listing choose **Disable**.

Administer analytics

This feature is available with a Confluence Data Center license.

If you have Confluence Administrator or System Administrator global permissions you can configure Confluence analytics to suit the needs of your organisation.

To find out what data is collected, seeAnalytics.

On this page:

- Disable analytics
- Permissions
- Data retention
- Rate limiting
- Increased privacy mode
- Known issues
- There are some known issues with analytics that you should be aware of.
- Analytics for Confluence marketplace app

This page is not about the analytics data which is sent to Atlassian to help us improve the product. See Data Collection Policy to find out what is collected, and how you can opt in or out of this data collection.

Disable analytics

If you don't want Confluence to collect user activity and page view data, you can disable the system app that provides the Analytics feature.

To disable analytics:

- Go to > Manage apps.
- 2. Search for Analytics for Confluence.
- 3. Expand the listing then select Disable.

Disabling the system app won't remove any existing analytics data from your database, but will stop Confluence collecting data.

Permissions

By default, all logged in users can view the Analytics option in the header, and access analytics data for spaces they have permissions to see. Anonymous users can't view analytics.

Limit who can view analytics reports

If you don't want everyone to be able to view analytics, you can chose to limit access to specific groups.

To limit analytics tospecific groups:

- 1. Go to General Configuration Analytics Permissions.
- 2. Search for a group.
- 3. Select Add.
- 4. Repeat this process for each group.

Only people who are a member of these groups will see the Analytics option in the header, or in a space.

To revert to the default, and allow all users to view analytics:

1. Go to O S General Configuration Analytics Permissions.

2. Remove all groups from the list.

You don't need to add every group individually to give everyone access.

Limit who can view analytics reports for specific spaces

You can further limit who can view analytics reports for specific spaces. You need Space Administrator permission to do this.

To change who can see analytics reports in a space:

- 1. Go to the space and choose Space tools>Permissions from the bottom of the sidebar.
- 2. Select the Analytics Permissions tab.
- 3. SelectViewing Analytics Restricted from the drop down.
- 4. Enter the users and/or groups you want to allow to view analytics reports.
- 5. Save your changes.

Good to know:

- If a Confluence administrator has denied a group permission to view analytics, adding the group at the space level will not grant this permission.
- The space will still appear in the site analytics report (if the user has permission to see the space, and
 use analytics globally), but they will be prevented from viewing the space analytics report if they don't
 have space permissions.

Data retention

If you have a very big, busy site, the amount of data collected can have an impact on your database size and general site performance. To avoid any problems, you can choose how long to retain events for, and change the maximum number of events that can be stored at any time.

If you have a very large, active site, we recommend keeping the data retention settings quite low, and using rate limiting.

Set a data retention period

By default we store analytics events in the database for 12 months. You can reduce or increase this limit if you have different requirements. A long retention period can affect the size and performance of your database, so if you have a big, busy site you may need to choose a shorter retention period.

To change the data retention period:

- 1. Go to > Manage apps.
- 2. Search for Analytics for Confluence.
- 3. Expand the listing then select**Configure**.
- 4. Under Data Retention Period, select Edit.
- 5. Choose either 6, 12, 18, or 24 months.
- 6. Save your changes.

Your change will take effect after 24 hours. This is to prevent data being immediately deleted if you accidentally choose the wrong retention period.

Confluence regularly deletes any events older than the chosen retention period (every 60 seconds). Note that the event retention limit also applies, so events exceeding the maximum number of events will be deleted, regardless of whether they fall within the data retention period.

Set an event retention limit

By default, we store a maximum of 20,000,000 analytics events in the database. This is to ensure analytics queries are reasonably performant in all databases. You can choose to further reduce this limit.

To change the maximum number of events to retain:

- 1. Go to > Manage apps.
- 2. Search for Analytics for Confluence.
- 3. Expand the listing then selectConfigure.
- 4. UnderEvent retention, selectEdit.
- 5. Enter the maximum number of events.
- 6. Save your changes.

Confluence regularly deletes any events exceeding limit (every 60 seconds), starting with the oldest events.

We don't recommend increasing this limit beyond 20 million events, as this will result in very large database tables, and eventually cause performance issues.

Rate limiting

Too many people accessing analytics data at the same time can slow down your site. To help avoid performance bottlenecks you can limit the number of reports that can be generated concurrently, and set a timeout for reports that are taking too long to load. This helps protect the overall performance of your site.

To change the maximum number of concurrent reports:

- 1. Go to > Manage apps.
- 2. Search for Analytics for Confluence.
- 3. Expand the listing then selectConfigure.
- 4. UnderRate Limiting, selectEdit.
- 5. Enter the maximum number of of concurrent reports.
- 6. Save your changes.

If someone tries to view an analytics report, and the limit has been reached, they'll see a message to try again in a few minutes.

To change the timeout:

- Go to > Manage apps.
- 2. Search for Analytics for Confluence.
- 3. Expand the listing then selectConfigure.
- 4. UnderData Retention Period, selectEdit.
- 5. Enter the new timeout value in seconds.
- 6. Save your changes.

If someone tries to view a report that can't be loaded within the time limit, they'll see a timeout error. When this happens, the best option is to change the report filters, such as reducing the date range.

Increased privacy mode

Privacy requirements can differ greatly between organisations and between different jurisdictions. In some circumstances it will not be appropriate for people to see the names of users who have viewed content in your site, or for you to be collecting this data in an identifiable format.

Turn on Increased privacy mode to minimise the amount of personally identifiable information (PII) that is collected and displayed to users.

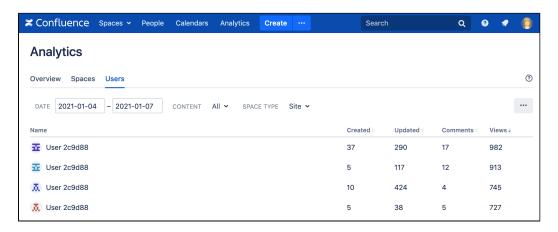
To turn on increased privacy mode:

- 1. Go to O > Manage apps.
- 2. Search for Analytics for Confluence.
- 3. Expand the listing then selectConfigure.
- 4. UnderIncreased Privacy Mode, selectEdit.
- 5. Select the **Enabled** toggle.
- 6. Save your changes.

From this point onwards, space or page activity is no longer linked to individuals, but instead attributed to an anonymised user. In analytics reports, people will be represented as "User 12345" with an anonymised avatar. This means you still get an accurate picture of the engagement with your content, but without revealing user information.

It is very important to note that changing this setting does not affect any previously collected data. This means:

- You may need to manually remove any previously collected data after turning on increased privacy mode.
- Unique user counts will be temporarily inaccurate, as we do not attempt to connect a named user with their anonymised alias, so unless you have manually removed the named user data, both will exist for a time.



Screenshot showing the site analytics users tab, with increase privacy mode turned on.

Known issues

There are some known issues with analytics that you should be aware of.

Site analytics reports may take a long time to load

There are a few situations where the site analytics reports (accessed from Analytics in the header) take longer than expected to load, including:

- if your spaces have complex permissions, or most spaces in your site are restricted
- if you're running Confluence on MySQL.

Some reports show data for deleted content

Some analytics reports continue to show aggregate data for content or user accounts that have subsequently been deleted.

The following reports continue to show views data for pages that have been deleted:

- Site analytics report spaces tab
- Site analytics report users tab
- Space analytics report users tab

The following reports continue to show views data for user accounts that have been deleted:

- Site analytics report spaces tab
- Space analytics report content tab

Comments column always shows data for pages and blogs

In the Users tab of any analytics report, the comments column always lists the total number of comments on pages and blog posts, regardless of whether the content filter is set to just pages, or just blog posts.

Analytics for Confluence marketplace app

The Good Software Analytics for Confluence was previously available on Marketplace. We ended new sales and renewals for this app in 2019, and subsequently made the app's features part of Confluence Data Center (version 7.11 and later).

If you have a Confluence Server license, you can continue to use the app until it reaches end of life, but features in your legacy version will differ from features available in Confluence Data Center.

If you have a Confluence Data Center license, the Good Software Analytics for Confluence marketplace app will be automatically disabled when you upgrade to Confluence Data Center 7.11 or later. Previously collected analytics data will be preserved.

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 Documentation.

Enterprise releases

A Long Term Support release is a feature release that gets backported critical security updates and critical bug fixes during its entire two-year support window. If you can only upgrade once a year, consider upgrading to a Long Term Support release.Learn more

Long Term Support releases were originally referred to as Enterprise Releases.

- System Requirements
 - Server Hardware Requirements Guide
 - Running Confluence in a Virtualized Environment
- Confluence Installation Guide
 - Installing Confluence
 - Installing Confluence Data Center
 - Installing Java for Confluence
 - 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
 - Migration from Wiki Markup to XHTML-Based Storage Format
 - Migration of Templates from Wiki Markup to XHTML-Based Storage Format
 - Upgrading Confluence Manually
 - Create a staging environment for upgrading Confluence
 - Upgrade Confluence without downtime
- Supported Platforms
 - End of Support Announcements for Confluence
 - Bundled Tomcat and Java versions
 - Supported Platforms FAQ

Downloads



Download the Confluence documentation in PDF format.

Other resources

Confluence Release Notes

Confluence administrator's guide

Confluence Knowledge Base

Atlassian Answers

System Requirements

Confluence can run on a wide range of operating systems and databases, on physical or virtualized servers.

SeeSupported Platformsfor the full list of platforms that we support in this version of Confluence orSupported Platforms FAQfor details on our support handling procedures.

Software requirements

Operating systems

Atlassian supports the operating systems listed on the Supported Platformspage.

If you would like to run Confluence on virtualized hardware, please read ourRunning Confluence in a Virtualized Environmentdocument first.

On this page:

- Software requirements
 - Operating systems
 - Application server
 - Databases
 - Java
 - Antivirus considerations
- Hardware requirements
- Hosted solutions Confluence Cloud

Application server

We only support running Confluence on the version of Apache Tomcat that is bundled with the Confluence distribution.

Databases

You'll need an external database to run Confluence. See the Supported Platforms page for a list of all the databases we support.

When evaluating Confluence, you can use the embedded H2 database included in the Confluence installation, but you will need to migrate to a supported external database once you're ready to roll Confluence out to your team.

Java

The Java Runtime Environment (JRE) is packed up and ready to go when you install Confluence using the Windows or Linux installer. You don't need to install Java yourself.

If you choose to install Confluence from an archive file, you'll need asupportedJRE or JDK, and your JAVA_HOME variable set correctly. SeeInstalling Java for Confluencefor more information.

Antivirus considerations

Antivirus software on the operating system running Confluence can greatly decrease the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental and may even cause errors in Confluence. This 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.

You should configure your antivirus software to ignore the following directories:

- Confluence home directory
- Confluence's index directory
- All database-related directories

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 more information.

You may also like to check out our tips on reducing out of memory errors, in particular the section on Perman ent Generation Size.

Hosted solutions Confluence Cloud

If you do not have the resources to set up and maintain a Confluence installation locally, consider trying Confluence Cloud. Atlassian can run and maintain your installation of Confluence, handling all the testing, monitoring and upgrading processes for you.

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.

(i) Enterprise-scale Confluence sites

These recommendations are designed for small or medium sized Confluence sites. For guidance on large or extra large sites, read our enterprise-scale infrastructure recommendations.

We've also created load profiles to help you determine the size of vour site.

Minimum hardware requirements

The values below refer to the minimum available hardware required to run Confluence only; for example, the minimum heap size to allocate to Confluence is 1 GB and 1 GB for Synchrony (which is required for collaborative editing). You'll need additional physical hardware, of at least the minimum amount required by your Operating System and any other applications that run on the server.

1 On small instances, server load is primarily driven by peak visitors, so minimum system requirements are difficult to judge. We provide these figures as a guide to the absolute minimum required to run Confluence, and vour configuration will likely require better hardware.

Here is our minimum hardware recommendation:

- CPU:Quad core 2GHz+ CPU
- RAM:6GB
- Minimum database space: 10GB

Note: Please be aware that while some of our customers run Confluence on SPARC-based hardware, we only officially support 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-intense startup process.

Example hardware specifications

These are example hardware specifications for non-clustered Confluence instances. It is not recorded whether the amount of RAM refers to either the 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 (MB)	Notes	
----------	--------	-------	------	--------------	-------------	-------	--

On this page:

- Minimum hardware requirements
- Example hardware specifications
- Server load and scalability
- Maximum reported usages
- Hard disk requirements
 - Private and public comparison
- **Professional** assistance
- Example site

Related pages:

- Confluence **Installation Guide**
- Managing **Application Server Memory Settings**
- Running Confluence in a Virtualized **Environment**

150	30	1,000	1	2.6	1,024	
350	100	15,000	2	2.8	1,536	
5,000	500		4	3	2,048	
10,000	350	16,000	2	3.8	2,048	
10,000	60	3,500	2	3.6	4,096	
21,000	950		2	3.6	4,096	
85,000	100	12,500	4	2.6	4,096	3 machines total: application server, database server, Apache HTTPD + LDAP tunnel server.

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 512MBof memory
- Alwaysuse an external database, and check out theperformance tuning guides.

Maximum reported usages

These values are largest customer instances reported to Atlassian or used for performance testing. Clustering, 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 page content is stored in the database, while attachments are stored in the file system. 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:

Case Accounts To Revisions Viewer Ratio	ents Comment
---	--------------

Online Docu mentat ion	140	11,500	1,000	9%	8,800	65,000	7,300	11,500
Privat e Intran et	130	180	140	78%	8,000	84,000	3,800	500
Comp any- Wide Collab oration	100	85,000	1,000+	1%+	12,500	120,000	15,000	

Professional assistance

For large instances, it may be worthwhile contacting anAtlassian Solution Partnerfor expertise on hardware sizing, testing and performance tuning.

Example site

Here's a breakdown of the disk usage and memory requirements of a large documentation site 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.

Running Confluence in a Virtualized Environment

This page provides pointers for things to look at when running Confluence on virtualized hardware.

Summary

Running Confluence in a virtual machine (VM) requires specialized skills to set up and manage the virtualized 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 running Confluence and Confluence Data Center in a virtualized environment, but we cannot offer support for problems which are related to the environment itself.

On this page:

- Summary
- Recommendations
- Further help

Related pages:

- Server Hardware Requirements Guide
- Confluence Installation Guide
- Confluence Data Center documentation
- AWS Quick Start (Data Center only)

Recommendations

The following recommendations come from our experience in running and testing Confluence in virtualized 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 virtualization 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 optimize 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 virtualization infrastructure has the flexibility and capacity to deal with
 peak load, not just idle load.
- Synchronize time correctly. Some customers have had problems with time synchronization
 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 virtualized environment.

Further help

For further assistance in setting up a virtualized environment for running Confluence, you may want to consult an Atlassian Solution Partner. Several experts have experience with installation and performance tuning, and can help you with your Confluence configuration.

Confluence Installation Guide

Before you start

Before installing Confluence, please check that you meet the minimum system requirementsandSupported Platforms.

If you're planning to run Confluence in a virtualized environment seeRunning Confluence in a Virtualized Environment.

Choose your installation method

There are a number of ways to install Confluence. Choose the method that is best for your environment.

Install method	Is this right for you?
Install a Confluence trialWindows, Linux or OS X	This is the fastest way to get a Confluence site up and running. If you want to see what Confluence can do, use this option or try Confluence Cloud free.
Install Confluence using an installer • Windows • Linux	This option uses an installer, and is the most straightforward way to get your production site up and running on a Windows or Linux server.
Install Confluence from a zip or archive file • Windows • Linux	This option requires you to manually install files and configure some system properties. It gives you the most control over the install process. Use this option if there isn't an installer for your operating system.
Run Confluence Server in a Docker container • Docker	This option gets Confluence Server up and running using a preconfigured Docker image. Head to https://docs.docker.com/ to find out more about Docker. Atlassian supports running Confluence in a Docker container, but we cannot offer support for problems which are related to the environment itself.
Install Confluence Data Center in a cluster • Windows and Linux • Kubernetes • AWS Quick Start • Azure	You can deploy a Confluence Data Center cluster on your own infrastructure or a public cloud platform like AWS or Azure. Read the Confluence Data Center Technical Overview for more details on clustering

Note:we do not support installing Confluence as a production system on OS X. An OS X download is available for the purposes of evaluating Confluence only. There are no limitations to using Confluence on a mac with any one of thesupported browsers.

The EAR/WAR distribution is no longer available, you'll need to install Confluence from a zip or archive file if you previously deployed Confluence into an existing application server.

Installing Confluence

There are a number of ways to install Confluence. Choose the method that is best for your environment.

Install method	Is this right for you?
Windows, Linux or OS X	This is the fastest way to get a Confluence site up and running. If you want to see what Confluence can do, use this option or try Confluence Cloud free.
Install Confluence using an installer • Windows • Linux	This option uses an installer, and is the most straightforward way to get your production site up and running on a Windows or Linux server.
Install Confluence from a zip or archive file • Windows • Linux	This option requires you to manually install files and configure some system properties. It gives you the most control over the install process. Use this option if there isn't an installer for your operating system.
Run Confluence Server in a Docker container • Docker	This option gets Confluence Server up and running using a pre- configured Docker image. Head to https://docs.docker.com/ to find out more about Docker. Atlassian supports running Confluence in a Docker container, but we cannot offer support for problems which are related to the environment itself.
Install Confluence Data Center in a cluster • Windows and Linux • Kubernetes • AWS Quick Start • Azure	You can deploy a Confluence Data Center cluster on your own infrastructure or a public cloud platform like AWS or Azure. Read the Confluence Data Center Technical Overview for more details on clustering

Installing a Confluence trial

Want to get up and running with Confluence quickly? This page will quide you through the steps to install and set up an evaluation Confluence Data Center site.



On this page:

Before you begin

- 1. Download the installer
- 2. Install Confluence
- 3. Set up Confluence





We no longer offer server evaluation licenses, but you can still trial our Data Center products. Learn more about these changes

If you're ready to set up a production Confluence site or you want more control, check out ourfull installation guides.

Before you begin

Our installers come with all the bits and pieces you need to run the application, but there's a few thingsyou'll need to get up and running:

 A computer or laptop with a supported operating system - you'll be installing Confluence so you'll need admin rights.

You can install Confluence on a Windows or Linux operating system.

Apple Mac isn't supported for production sites, but if you're comfortable setting up applications on your Mac from scratch, you can download the tar.gz file and follow the instructions for Installing Confluence on Linux from Archive File as the process is similar.

- A supported web browser you'll need this to access Confluence, we support the latest versions of Chrome and Mozilla Firefox, Internet Explorer 11, and Microsoft Edge.
- A valid email address you'll need this to generate your evaluation license and create an account.
- An external database. To run Confluence you'll need an external database. Check the Supported Platforms page for the version you're installing for the list of databases we currently support. If you don't already have a database, PostgreSQL is free and easy to set up.

Good to know:

- Set up your database before you begin. Step-by-step guides are available for PostgreSQL, Ora cle, MySQL, and SQL Server.
- If you're using Oracle or MySQL you'll need to download the driver for your database.
- To use a datasource seeConfiguring a datasource connection as there are some steps you need to perform before running the setup wizard.

Ready to get going? Let's start with grabbing the installer.

1. Download the installer

Head towww.atlassian.com/software/confluence/downloadand download the installer for your operating system.

2. Install Confluence

The installer allows you to choose Express or Custom installations.

The Custom installation allows you to pick some specific options for Confluence, but for this guide we'll use the Express installation.

- Run the installer we recommend running with a Windows administrator account.
 If prompted, make sure you allow the installer to make changes to your computer. This will allow you to install Confluence as a service.
- Choose Express Install, then click Next.
- 3. Once installation is complete, it will ask you if you want to open Confluence in your browser. Make sure this option is selected then click **Done**.
- 4. Confluence will open in your default browser, and you're ready to start the set up wizard.
- Change to the directory where you downloaded Confluence then execute this command to make it executable:

```
$ chmod a+x atlassian-confluence-X.X.X-x64.bin
```

Where x.x.x is is the Confluence version you downloaded.

2. Run the installer - we recommend usingsudoto run the installer as this will create a dedicated account to run Confluence and allow you to run Confluence as a service.

```
$ sudo ./atlassian-confluence-X.X.X-x64.bin
```

- 3. When prompted, choose Express Install (option 1).
- 4. Once installation is complete head tohttp://localhost:8090/in your browser to begin the setup process.

3. Set up Confluence

The set up wizard is the last step in getting Confluence up and running. You'll need your email address to generate your evaluation license.

- 1. Select**Trial**, then select**Next**.
- 2. Select**Get an evaluation license** and follow the prompts to generate your trial Data Center license.
- 3. Choose whether you want to try a standalone (single node) or clustered installation. 'Standalone' is the fastest way to get started. If you choose 'Clustered', you'll need to configure your cluster before continuing.
- Enter the details for your database, or set one up. See the 'Before you begin' section of this page for details and connection options.
 - It will take a few minutes to get everything connected and operational.
- 5. SelectManage users with Confluence, and selectNext.
- 6. Enter and confirm the details you want to use for your administrator account, and select**Done**.

That's it! You're ready to team up with some colleagues and start using Confluence.

Installing Confluence on Windows

In this guide we'll run you through installing Confluence in a production environment, with an external database, using the Windows installer.

This is the most straightforward way to get your production site up and running on a Windows server.



On this page:

Before you begin Install Confluence

- 1.Download Confluence
- 2. Run the installer

Set up Confluence

- 3. Choose installation type
- 4. Enter your license
- 5. Connect to your database
- 6. Populate your new site with content
- 7. Choose where to manage users
- 8. Create your administrator account
- 9. Start using Confluence

Troubleshooting

Other ways to install Confluence:

- Evaluation -get your free trial up and running in no time.
- Zip install Confluence manually from a zip file.
- Linux install Confluence on a Linux operating system

Before you begin

Before you install Confluence, there's a few questions you need to answer.

Are you using a supported operating system?

Check the Supported Platforms page for the version of Confluence you are installing. This will give you info on supported operating systems, databases and browsers.

Good to know:

- We don't support installing Confluence on OSX.
- The Confluence installer includes Java (JRE) and Tomcat, so you don't need to install these separately.

Do you want to run Confluence as a Windows Service? Running Confluence as a service in Windows means that Confluence will automatically start up when Windows is started.

If you choose to run Confluence as a service:

- You must run the installer as administrator to be able 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.
- We strongly recommend creating a dedicated user account (e.g. with username 'confluence') for running Confluence. See Creating a Dedicated User Account on the Operating System to Run Confluence to find out what directories this user will need to be able to read and write to.

If you choose not to run Confluence as a service:

- You will start and stop Confluence using the Windows Start menu, or by running a file in your Confluence installation directory.
- Confluence will be run as the Windows user account that was used to install Confluence, or you can choose to run as a dedicated user.
- Confluence will need to be restarted manually if your server is restarted.

Are ports 8090 and 8091 available?

Confluence runs on port 8090 by default. If this port is already in use, the installer will prompt you to choose a different port.

Synchrony, which is required for collaborative editing, runs on port 8091 by default. If this port is already in use, you will need to change the port that Synchrony runs on after your Confluence installation is complete. See Administering Collaborative Editing to find out how to change the port Synchrony runs on. You won't be able to edit pages until Synchrony has an available port.

SeePorts used by Atlassian Applicationsfor a summary of all the ports used.

Is your database set up and ready to use?

To run Confluence you'll need an external database. Check the Supported Platforms page for the version you're installing for the list of databases we currently support. If you don't already have a database, PostgreSQL is free and easy to set up.

Good to know:

- Set up your database before you begin. Step-by-step guides are available for Postgre SQL, Oracle, MySQL, and SQL Server.
- If you're using Oracle or MySQL you'll need to download the driver for your database.
- To use a datasource seeConfiguring a datasource connection as there are some steps you need to perform before running the setup wizard.

Do you have a Confluence license?

You'll need a valid license to use Confluence.

Good to know:

- If you have not yet purchased a Confluence license you'll be able to create an evaluation license during setup.
- If you already have a license key you'll be prompted to log in to my.atlassian.com to retrieve it, or you can enter the key manually during setup.
- If you're migrating from Confluence Cloud, you'll need a new license.
- Weve ended sales for new server licenses and will end support for server on February 2, 2024. Were continuing our investment in Data Center. Learn more



Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Install Confluence

1.Download Confluence

Download the installer for your operating system -https://www.atlassian.com/software/confluence/download

2. Run the installer

- 1. Run the installer. We recommend using a dedicated Windows administrator account.
- 2. Follow the prompts to install Confluence. You'll be asked for the following info:
 - a. **Destination directory**this is whereConfluencewill be installed.
 - b. Home directorythis is where Confluence data like logs, search indexes and files will be stored.
 - c. TCP portsthese arethe HTTP connector port and control port Confluence will run on. Stick with the default unless you're running another application on the same port.
 - d. Install as servicethis option is only available if you ran the installer as administrator.
- 3. Confluence will start up in your browser once installation is complete.

Set up Confluence

3. Choose installation type

- 1. Choose**Production installation**.
- 2. Choose any appsyou'd also like to install.

4. Enter your license

Follow the prompts to log in tomy.atlassian.comto retrieve your license, or enter a license key.

5. Connect to your database

- 1. If you've not already done so, it's time to create your database. See the 'Before you begin' section of this page for details and connection options.
- 2. For MySQL and Oracle, follow the prompts to download and install the required driver.
- 3. Enter your database details. Usetest connection to check your database is set up correctly.

If you want to specify particular parameters, you can choose to connect By connection string. You'll be prompted to enter:

- Database URL the JDBC URL for your database. If you're not sure, check the documentation for your database.
- Username and Password A valid username and password that Confluence can use to access your database.

6. Populate your new site with content

Choose whether you'd like Confluence to populate your site with content:

This option will create a space that you and your users can use to get to know Confluence. You can delete this space at any time.

Use this option if you have a full site export of an existing Confluence site. This is useful when youre migrating to another database or setting up a test site.

Good to know:

- You can only import sites from the same or earlier Confluence version.
- The system administrator account and all other user data and content will be imported from your previous installation.

In the setup wizard:

- Upload a backup file use this option if your site export file is small (25mb or less).
- Restore a backup file from the file system use this option if your backup file is large. Drop the file into your <confluence-home>/restoredirectory then follow the prompts to restore the backup.
- **Build Index** well need to build an index before your imported content is searchable. This can take a long time for large sites, so deselect this option if you would rather build the index later. Your content won't be searchable until the index is built.

7. Choose where to manage users

Choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Management:

Choose this option if you're happy to manage users in Confluence, or don't have a Jira application installed.

Good to know:

- If you do plan to manage users in a Jira application, but have not yet installed it, we recommend installing Jira first, and then returning to the Confluence setup.
- You can add external user management (for example LDAP, Crowd or Jira) later if you choose.

Choose this option if you have a Jira application installed and want to manage users across both applications.

Good to know:

- 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.
- You'll be able to specify exactly which groups in your Jira app should also be allowed to log in to Confluence. Your license tiers do not need to be the same for each application.
- You'll need either Jira 4.3 or later, Jira Core 7.0 or later, Jira Software 7.0 or later, or Jira Service Management 3.0 or later.

In the setup wizard:

- Jira Base URL the address of your Jira server, such as http://www.example.com:8080/jira/orhttp://jira.example.com/
- Jira Administrator Login this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application. Confluence will also use this username and password to create a local administrator account which will let you access Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.
- Confluence Base URL this is the URL Jira will use to access your Confluence server. The URL you give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira application.

- **User Groups** these are the Jira groups whose members should be allowed to use Confluence. Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
 - Jira 6.4 and earlier: jira-users.
 - Jira Software 7.x and later: jira-software-users
 - Jira Core 7.x and later: jira-core-users
 - Jira Service Management (formerly Jira Service Desk) 3.x and later: jira-servicedesk-users
- Admin Groups provide one or more Jira groups whose members should have administrative access to Confluence. The default group is jira-administrators. These groups will get the system administrator and Confluence administrator global permissions in Confluence.

8. Create your administrator account

Enter details for the administrator account.

Skip this step if you chose to manage users in a Jira application or you imported data from an existing site.

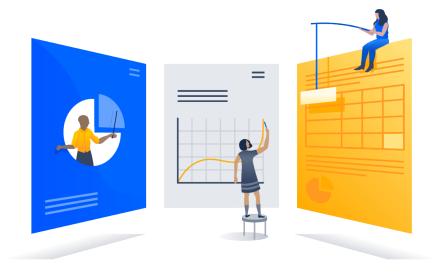
9. Start using Confluence

That's it!Your Confluence site is accessible from a URL like this: http://<computer_name_or_IP_address>: <port>

If you plan to run Confluence behind a reverse proxy, check outProxy and SSL considerationsbefore you go any further.

Here's a few things that will help you get your team up and running:

- Set the server base URL this is the URL people will use to access Confluence.
- Set up a mail server this allows Confluence to send people notification about content.
- Add and invite users get your team on board!
- Start and stop Confluence find out how to start and stop Confluence.



Troubleshooting

- 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 theConfluence installation.
- Can't start Confluence? See Confluence does not start due to Spring Application context has not been set
- Collaborative editing errors? See Troubleshooting Collaborative Editing.

Head to Installation Troubleshooting in our Knowledge Base for more help.

Installing Confluence on Windows from Zip File

In this guide we'll run you through installing Confluence in a production environment, with an external database, manually using a zip file.

This method gives you the most control of the installation process.



Other ways to install Confluence:

- Evaluation-get your free trial up and running in no time.
- Installer install Confluence using the Windows installer.
- Linux install Confluence on a Linux operating system.

On this page:

Before you begin Install Confluence

- 1.Download Confluence
- 2. Create the installation directory
- 3. Create the home directory
- 4. Check the ports
- 5. Start Confluence

Set up Confluence

- 6. Choose installation type
- 7. Enter your license
- 8. Connect to your database
- 9. Populate your new site with content
- 10. Choose where to manage users
- 11. Create your administrator account
- 12. Start using Confluence

Troubleshooting

Before you begin

Before you install Confluence, there's a few questions you need to answer.

Are you using a supported operating system and Java version?

Check the Supported Platforms page for the version of Confluence you are installing. This will give you info on supported operating systems, databases and browsers.

Good to know:

- We don't support installing Confluence on OS X or mac OS for production environments.
- You'll need to install either Adoptium OpenJDK (formerly AdoptOpenJDK)or Oracle JDK. We don't support other OpenJDK binaries.
- You can use either the JDK (Java Development Kit) or JRE (Java Runtime Environment).
- We only support the version of Apache Tomcat that is bundled with Confluence.

Do you want to run Confluence as a Windows Service? Running Confluence as a service in Windows means that Confluence will automatically start up when Windows is started.

You should use the Windows installer if you want to run Confluence as a Service.

If you choose not to run Confluence as a service:

- You will start and stop Confluence by running the start-confluence.bat file in your Confluence installation directory.
- Confluence will be run as the Windows user account that was used to install
 Confluence, or you can choose to run as a dedicated user (this user must have full
 read and write access to the installation directory and home directory).
- Confluence will need to be restarted manually if your server is restarted.

Are ports 8090 and 8091 available? Confluence runs on port 8090 by default. If this port is already in use, the installer will prompt you to choose a different port.

Synchrony, which is required for collaborative editing, runs on port 8091 by default. If this port is already in use, you will need to change the port that Synchrony runs on after your Confluence installation is complete. See Administering Collaborative Editing to find out how to change the port Synchrony runs on. You won't be able to edit pages until Synchrony has an available port.

SeePorts used by Atlassian Applicationsfor a summary of all the ports used.

What database do you plan to use?

To run Confluence you'll need an external database. Check the Supported Platforms page for the version you're installing for the list of databases we currently support. If you don't already have a database, PostgreSQL is free and easy to set up.

Good to know:

- Set up your database before you begin. Step-by-step guides are available for Postg reSQL, Oracle, MySQL, and SQL Server.
- If you're using Oracle or MySQL you'll need to download the driver for your database.
- To use a datasource seeConfiguring a datasource connection as there are some steps you need to perform before running the setup wizard.

Do you have a Confluence license? You'll need a valid license to use Confluence.

Good to know:

- If you have not yet purchased a Confluence license you'll be able to create an evaluation license during setup.
- If you already have a license key you'll be prompted to log in to my.atlassian.com to retrieve it, or you can enter the key manually during setup.
- If you're migrating from Confluence Cloud, you'll need a new license.
- Weve ended sales for new server licenses and will end support for server on February 2, 2024. Were continuing our investment in Data Center.Learn more

Is your JRE_HOME variable set correctly?

Before you install Confluence, check that you're running a supported Java version and that the JRE_HOME (or JAVA_HOME) environment variable is set correctly.

To check the JRE HOME variable:

Open a command prompt and type echo %JRE_HOME% and hit Enter.

- If you see a path to your Java installation directory, the JRE_Home environment variable has been set correctly.
- If nothing is displayed, or only %JRE_HOME% is returned, you'll need to set the JRE_ HOME environment variable manually. See Setting the JAVA_HOME Variable in Windows for a step by step guide.



⚠ Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Install Confluence

1.Download Confluence

Download the zip file for your operating system https://www.atlassian.com/software/confluence/download.

2. Create the installation directory

- 1. Create your installation directory (with full control permission for the dedicated Windows administrator accountyou'll use to run Confluence) this is where Confluence will be installed. Avoid using spaces or special characters in the path. We'll refer to this directory as your < installation-directory >.
- 2. Extract the Confluence zip file to your<installation-directory>. We recommend using7ziporWi nzip.

3. Create the home directory

- 1. Create your home directory (with full control permission) this is where Confluence data like logs, search indexes and files will be stored. This should be seperate to your installation directory. We'll refer to this directory as your<home-directory>.
- 2. Edit<installation-directory>\confluence\WEB-INF\classes\confluence-init. properties.
- 3. At the bottom of the file, enter the path to your<home directory>.

You can edit the confluence-init.properties file in Notepad or any other text editor.

a. Scroll to the bottom of the text and find this line:

confluence.home=c:/confluence/data

b. Remove the '#' and the space at the beginning of this line (so Confluence doesn't regard the line as a comment)

confluence.home=c:/data/confluence-home

- c. If you decide to use a different directory as the home directory you should:
 - Avoid spaces in the directory path or file name.
 - Use forward slashes '/' to define the path in this file.

4. Check the ports

By default Confluence listens on port8090. If you have another application running on your server that uses the same ports, you'll need to tell Confluence to use a different port.

To change the ports:

- 1. Edit <installation-directory>\conf\server.xml
- 2. Change the **Server** port (8000) and the **Connector** port (8090) to free ports on your server.

In the example below we've changed the **Server** port to 5000 and the **Connector** port to 5050.

```
Server port="5000" shutdown="SHUTDOWN" debug="0">

<Service name="Tomcat-Standalone">

<Connector port="5050" connectionTimeout="20000" redirectPort="8443"

maxThreads="48" minSpareThreads="10"

enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"

protocol="org.apache.coyote.httpl1.Httpl1NioProtocol" />
```

5. Start Confluence

Run<installation-directory>/bin/start-confluence.batto start the install process. We
recommend using a dedicated user account.

A command prompt will open. Closing this window will stop Confluence.

- 2. Go to http://localhost:8090/ to launch Confluence in your browser (change the port if you've updated the Connector port).
 - If the command prompt window closes immediately, your JAVA_HOME variable may not be set correctly. See Setting the JAVA_HOME Variable in Windows.
 - If you see an error, see Confluence does not start due to Spring Application context has not been set for troubleshooting options.

Set up Confluence

6. Choose installation type

- 1. ChooseProduction installation.
- 2. Choose any appsyou'd also like to install.

7. Enter your license

Follow the prompts to log in tomy.atlassian.comto retrieve your license, or enter a license key.

8. Connect to your database

- 1. If you've not already done so, it's time to create your database. See the 'Before you begin' section of this page for details and connection options.
- 2. For MySQL and Oracle, follow the prompts to download and install the required driver.
- 3. Enter your database details. Use**test connection** to check your database is set up correctly.

If you want to specify particular parameters, you can choose to connect **By connection string**. You'll be prompted to enter:

- **Database URL** the JDBC URL for your database. If you're not sure, check the documentation for your database.
- **Username and Password** A valid username and password that Confluence can use to access your database.

9. Populate your new site with content

Choose whether you'd like Confluence to populate your site with content:

This option will create a space that you and your users can use to get to know Confluence. You can delete this space at any time.

Use this option if you have a full site export of an existing Confluence site. This is useful when youre migrating to another database or setting up a test site.

Good to know:

- You can only import sites from the **same** or **earlier** Confluence version.
- The system administrator account and all other user data and content will be imported from your previous installation.

In the setup wizard:

- Upload a backup file use this option if your site export file is small (25mb or less).
- Restore a backup file from the file system use this option if your backup file is large. Drop the file into your <confluence-home>/restoredirectory then follow the prompts to restore the backup.
- **Build Index** well need to build an index before your imported content is searchable. This can take a long time for large sites, so deselect this option if you would rather build the index later. Your content won't be searchable until the index is built.

10. Choose where to manage users

Choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Management:

Choose this option if you're happy to manage users in Confluence, or don't have a Jira application installed.

Good to know:

- If you do plan to manage users in a Jira application, but have not yet installed it, we recommend installing Jira first, and then returning to the Confluence setup.
- You can add external user management (for example LDAP, Crowd or Jira) later if you choose.

Choose this option if you have a Jira application installed and want to manage users across both applications.

Good to know:

- 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.
- You'll be able to specify exactly which groups in your Jira app should also be allowed to log in to Confluence. Your license tiers do not need to be the same for each application.
- You'll need either Jira 4.3 or later, Jira Core 7.0 or later, Jira Software 7.0 or later, or Jira Service Management 3.0 or later.

In the setup wizard:

• Jira Base URL the address of your Jira server, such as http://www.example.com:8080/jira/orhttp://jira.example.com/

- **Jira Administrator Login** this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application. Confluence will also use this username and password to create a local administrator account which will let you access Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.
- Confluence Base URL this is the URL Jira will use to access your Confluence server. The URL you give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira application.
- User Groups these are the Jira groups whose members should be allowed to use Confluence.
 Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
 - Jira 6.4 and earlier: jira-users.
 - Jira Software 7.x and later: jira-software-users
 - Jira Core 7.x and later: jira-core-users
 - Jira Service Management (formerly Jira Service Desk) 3.x and later: jira-servicedeskusers
- Admin Groups provide one or more Jira groups whose members should have administrative
 access to Confluence. The default group is jira-administrators. These groups will get the
 system administrator and Confluence administrator global permissions in Confluence.

11. Create your administrator account

Enter details for the administrator account.

Skip this step if you chose to manage users in a Jira application or you imported data from an existing site.

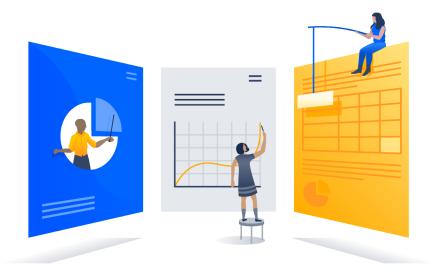
12. Start using Confluence

That's it!Your Confluence site is accessible from a URL like this: http://<computer_name_or_IP_address>: <port>

If you plan to run Confluence behind a reverse proxy, check outProxy and SSL considerationsbefore you go any further.

Here's a few things that will help you get your team up and running:

- Set the server base URL this is the URL people will use to access Confluence.
- Set up a mail server this allows Confluence to send people notification about content.
- Add and invite users get your team on board!
- Start and stop Confluence find out how to start and stop Confluence.



Troubleshooting

• If your web browser window shows an error the first time you try to access Confluence, wait for 30 seconds or so and then refresh the page.

- If the command prompt window closes immediately, your JAVA_HOME variable may not be set correctly. See Setting the JAVA_HOME Variable in Windows.
- If you see an error, see Confluence does not start due to Spring Application context has not been set for troubleshooting options.
- Collaborative editing errors? See Troubleshooting Collaborative Editing.

Head to Installation Troubleshooting in our Knowledge Base for more help.

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.
 - 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.

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

In this guide we'll run you through installing Confluence in a production environment, with an external database, using the Linux installer.

This is the most straightforward way to get your production site up and running on a Linux server.



Other ways to install Confluence:

- Evaluation-get your free trial up and running in no time.
- TAR.GZ install Confluence manually from an archive file.
- Windows install Confluence on a Windows server.

On this page:

Before you begin Install Confluence

- 1.Download Confluence
- 2. Run the installer

Set up Confluence

- 3. Choose installation type
- 4. Enter your license
- 5. Connect to your database
- 6. Populate your new site with content
- 7. Choose where to manage users
- 8. Create your administrator account
- 9. Start using Confluence

Troubleshooting

Before you begin

Before you install Confluence, there are a few questions you need to answer.

Are you using a supported operating system?	Check the Supported Platforms page for the version of Confluence you are installing. This will give you info on supported operating systems, databases and browsers. Good to know: We don't support installing Confluence on OSX for production sites. The Confluence installer includes Java (JRE) and Tomcat, so you don't need to install these separately. Confluence can only run on Oracle JDK or AdoptOpenJDK.
Does your Linux server have a font	Many Linux distributions don't include a suitable font config package by default, so you will need to install one before you can run the Confluence installer.
config package installed?	See Confluence Server 6.13 or later fails with FontConfiguration error when installing on Linux operating systems for commands to install a suitable package on several popular Linux distributions.

Do you want to run Confluence as a service?

Running Confluence as a service means that Confluence will automatically start up when Linux is started.

If you choose to run Confluence as a service:

- You must use sudo to run the installer to be able to install Confluence as a service.
- The installer will create a dedicated user account, confluence, that will run the service.

If you choose not to run Confluence as a service:

- You will start and stop Confluence by running the start-confluence.sh file in your Confluence installation directory.
- Confluence will be run as the user account that was used to install Confluence, or you can choose to run as a dedicated user.
- Confluence will need to be restarted manually if your server is restarted.

Are ports 8090 and 8091 available?

Confluence runs on port 8090 by default. If this port is already in use, the installer will prompt you to choose a different port.

Synchrony, which is required for collaborative editing, runs on port 8091 by default. If this port is already in use, you will need to change the port that Synchrony runs on after your Confluence installation is complete. See Administering Collaborative Editing to find out how to change the port Synchrony runs on. You won't be able to edit pages until Synchrony has an available port.

SeePorts used by Atlassian Applicationsfor a summary of all the ports used.

Is your database set up and ready to use?

To run Confluence you'll need an external database. Check the Supported Platforms page for the version you're installing for the list of databases we currently support. If you don't already have a database, PostgreSQL is free and easy to set up.

Good to know:

- Set up your database before you begin. Step-by-step guides are available for Postgr eSQL, Oracle, MySQL, and SQL Server.
- If you're using Oracle or MySQL you'll need to download the driver for your database.
- To use a datasource seeConfiguring a datasource connection as there are some steps you need to perform before running the setup wizard.

Do you have a Confluence license?

You'll need a valid license to use Confluence.

Good to know:

- If you have not yet purchased a Confluence license you'll be able to create an evaluation license during setup.
- If you already have a license key you'll be prompted to log in to my.atlassian.com to retrieve it, or you can enter the key manually during setup.
- If you're migrating from Confluence Cloud, you'll need a new license.
- Weve ended sales for new server licenses and will end support for server on February 2, 2024. Were continuing our investment in Data Center. Learn more



Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Install Confluence

1.Download Confluence

Download the installer for your operating system https://www.atlassian.com/software/confluence/download

2. Run the installer

1. Make the installer executable.

Change to the directory where you downloaded Confluence then execute this command:

```
$ chmod a+x atlassian-confluence-X.X.X-x64.bin
```

Where x.x.x is is the Confluence version you downloaded.

2. Run the installer we recommend usingsudoto run the installer as this will create a dedicated account to run Confluence and allow you to run Confluence as a service.

To use sudo to run the installer execute this command:

```
$ sudo ./atlassian-confluence-X.X.X-x64.bin
```

Where x.x.x is is the Confluence version you downloaded.

You can also choose to run the installer as with root user privileges.

- 3. Follow the prompts to install Confluence. You'll be asked for the following info:
 - Install type choose option 2 (custom) for the most control.
 - **Destination directory** this is whereConfluencewill be installed.
 - Home directory this is where Confluence data like logs, search indexes and files will be stored.
 - TCP ports these arethe HTTP connector port and control port Confluence will run on. Stick with the default unless you're running another application on the same port.
 - Install as service this option is only available if you ran the installer assudo.
- 4. Once installation is complete head tohttp://localhost:8090/in your browser to begin the setup process. (Replace 8090 if you chose a different port during installation).

If you're installing Confluence on a fresh Linux installation seeConfluence throws a Confluence is vacant error on install for troubleshooting options.

FontConfiguration error? See Confluence Server 6.13 or later fails with FontConfiguration error when installing on Linux operating systems to find out how to install a suitable font configuration package.

Set up Confluence

3. Choose installation type

- 1. ChooseProduction installation.
- 2. Choose any apps you'd also like to install.

4. Enter your license

Follow the prompts to log in tomy.atlassian.comto retrieve your license, or enter a license key.

5. Connect to your database

- 1. If you've not already done so, it's time to create your database. See the 'Before you begin' section of this page for details and connection options.
- 2. For MySQL and Oracle, follow the prompts to download and install the required driver.
- 3. Enter your database details. Use**test connection** to check your database is set up correctly.

If you want to specify particular parameters, you can choose to connect **By connection string**. You'll be prompted to enter:

- **Database URL** the JDBC URL for your database. If you're not sure, check the documentation for your database.
- Username and Password A valid username and password that Confluence can use to access your database.

6. Populate your new site with content

Choose whether you'd like Confluence to populate your site with content:

This option will create a space that you and your users can use to get to know Confluence. You can delete this space at any time.

Use this option if you have a full site export of an existing Confluence site. This is useful when youre migrating to another database or setting up a test site.

Good to know:

- You can only import sites from the **same** or **earlier** Confluence version.
- The system administrator account and all other user data and content will be imported from your previous installation.

In the setup wizard:

- Upload a backup file use this option if your site export file is small (25mb or less).
- Restore a backup file from the file system use this option if your backup file is large. Drop the file into your <confluence-home>/restoredirectory then follow the prompts to restore the backup.
- **Build Index** well need to build an index before your imported content is searchable. This can take a long time for large sites, so deselect this option if you would rather build the index later. Your content won't be searchable until the index is built.

7. Choose where to manage users

Choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Management:

Choose this option if you're happy to manage users in Confluence, or don't have a Jira application installed.

Good to know:

• If you do plan to manage users in a Jira application, but have not yet installed it, we recommend installing Jira first, and then returning to the Confluence setup.

You can add external user management (for example LDAP, Crowd or Jira) later if you choose.

Choose this option if you have a Jira application installed and want to manage users across both applications.

Good to know:

- 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.
- You'll be able to specify exactly which groups in your Jira app should also be allowed to log in to Confluence. Your license tiers do not need to be the same for each application.
- You'll need either Jira 4.3 or later, Jira Core 7.0 or later, Jira Software 7.0 or later, or Jira Service Management 3.0 or later.

In the setup wizard:

- Jira Base URL the address of your Jira server, such as http://www.example.com:8080/jira/orhttp://jira.example.com/
- Jira Administrator Login this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application. Confluence will also use this username and password to create a local administrator account which will let you access Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.
- Confluence Base URL this is the URL Jira will use to access your Confluence server. The URL
 you give here overrides the base URL specified in Confluence, for the purposes of connecting to
 the Jira application.
- User Groups these are the Jira groups whose members should be allowed to use Confluence.
 Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
 - Jira 6.4 and earlier: jira-users.
 - Jira Software 7.x and later: jira-software-users
 - Jira Core 7.x and later: jira-core-users
 - Jira Service Management (formerly Jira Service Desk) 3.x and later: jira-servicedeskusers
- Admin Groups provide one or more Jira groups whose members should have administrative access to Confluence. The default group is jira-administrators. These groups will get the system administrator and Confluence administrator global permissions in Confluence.

8. Create your administrator account

Enter details for the administrator account.

Skip this step if you chose to manage users in a Jira application or you imported data from an existing site.

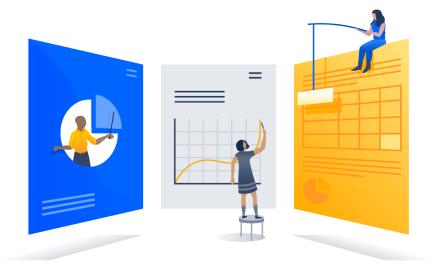
9. Start using Confluence

That's it!Your Confluence site is accessible from a URL like this: http://<computer_name_or_IP_address>: <port>

If you plan to run Confluence behind a reverse proxy, check outProxy and SSL considerationsbefore you go any further.

Here's a few things that will help you get your team up and running:

- Set the server base URL this is the URL people will use to access Confluence.
- Set up a mail server this allows Confluence to send people notification about content.
- Add and invite users get your team on board!
- Start and stop Confluence find out how to start and stop Confluence.



Troubleshooting

- If the installer fails with a FontConfiguration error, you'll need to install a font package. See Conflue nce Server 6.13 or later fails with FontConfiguration error when installing on Linux operating systems for info on how to do this.
- 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 theConfluence installation.
- The Linux OOM Killer can sometimes kill Confluence processes when memory on the server becomes too low. See How to Configure the Linux Out-of-Memory Killer.
- Collaborative editing errors? See Troubleshooting Collaborative Editing.

Head to Installation Troubleshooting in our Knowledge Base for more help.

Installing Confluence on Linux from Archive File

In this guide we'll run you through installing Confluence in a production environment, with an external database, manually using a zip file.

This method gives you the most control over the installation process.



Other ways to install Confluence:

- Evaluation-get your free trial up and running in no time.
- Installer install Confluence using the Linux installer.
- Windows install Confluence on a Windows server.

On this page:

Before you begin Install Confluence

- 1.Download Confluence
- 2. Create the installation directory
- 3. Create the home directory
- 4. Check the ports
- 5. Start Confluence

Set up Confluence

- 6. Choose installation type
- 7. Enter your license
- 8. Connect to your database
- 9. Populate your new site with content
- 10. Choose where to manage users
- 11. Create your administrator account
- 12. Start using Confluence Troubleshooting

Before you begin

Before you install Confluence, there are a few questions you need to answer.

Are you using a supported operating system and Java version?

Check the Supported Platforms page for the version of Confluence you are installing. This will give you info on supported operating systems, databases and browsers.

Good to know:

- We don't support installing Confluence on OS X or mac OS for production environments.
- You'll need to install either Adoptium OpenJDK (formerly AdoptOpenJDK)or Oracle JDK. We don't support other OpenJDK binaries.
- You can use either the JDK (Java Development Kit) or JRE (Java Runtime Environment).
- We only support the version of Apache Tomcat that is bundled with Confluence.

Do you want to run Confluence as a service? Running Confluence as a service means that Confluence will automatically start up when your Linux server is started.

You should use the Linux installer if you want to run Confluence as a service.

If you choose not to run Confluence as a service:

- You will start Confluence by running the start-confluence.sh file in your Confluence installation directory.
- We recommend creating a dedicated user to run Confluence. This user must have full read, write and execute access to the installation directory and home directory.
- Confluence will need to be restarted manually if your server is restarted.

Are ports 8090 and 8091 available?

Confluence runs on port 8090 by default. If this port is already in use, the installer will prompt you to choose a different port.

Synchrony, which is required for collaborative editing, runs on port 8091 by default. If this port is already in use, you will need to change the port that Synchrony runs on after your Confluence installation is complete. See Administering Collaborative Editing to find out how to change the port Synchrony runs on. You won't be able to edit pages until Synchrony has an available port.

SeePorts used by Atlassian Applicationsfor a summary of all the ports used.

What database do you plan to use?

To run Confluence you'll need an external database. Check the Supported Platforms page for the version you're installing for the list of databases we currently support. If you don't already have a database, PostgreSQL is free and easy to set up.

Good to know:

- Set up your database before you begin. Step-by-step guides are available for Postg reSQL, Oracle, MySQL, and SQL Server.
- If you're using Oracle or MySQL you'll need to download the driver for your database.
- To use a datasource seeConfiguring a datasource connection as there are some steps you need to perform before running the setup wizard.

Do you have a Confluence license?

You'll need a valid license to use Confluence.

Good to know:

- If you have not yet purchased a Confluence license you'll be able to create an evaluation license during setup.
- If you already have a license key you'll be prompted to log in to my.atlassian.com to retrieve it, or you can enter the key manually during setup.
- If you're migrating from Confluence Cloud, you'll need a new license.
- Weve ended sales for new server licenses and will end support for server on February 2, 2024. Were continuing our investment in Data Center.Learn more

Is your JAVA_HOM E variable set correctly?

Before you install Confluence, check that you're running a supported Java version and that the JAVA_HOME environment variable is set correctly.

Confluence can only run with Oracle JDK or JRE.

To check your Java version:

\$ java -version

To check your JAVA_HOME variable is set correctly:

\$ echo \$JAVA_HOME

If you see a path to your Java installation directory, the <code>JAVA_Home</code> environment variable has been set correctly. If a path is not returned you'll need to set your <code>JAVA_HOME</code> environment variable manually before installing Confluence.

Have you created a dedicated user to run Confluence?

We strongly recommend running Confluence as a dedicated user.

You should create this user before you begin, so that when creating the installation and home directories, you can give this user appropriate read and write permissions.

In this example, we'll create a user called confluence:

\$ sudo /usr/sbin/useradd --create-home --comment "Account for running Confluence" --shell /bin/bash confluence

See Creating a Dedicated User Account on the Operating System to Run Confluence for more information.



Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Install Confluence

1.Download Confluence

Download thetar, gzfile for your operating system -https://www.atlassian.com/software/confluence /download.

2. Create the installation directory

1. Create your installation directory this is where Confluence will be installed. Avoid using spaces or special characters in the path. We'll refer to this directory as your < installation-directory >.

In this example we'll call our installation directory confluence:

```
$ mkdir confluence
```

2. Extract the Confluencetar.gzfile to your<installation-directory>. We recommend using aGNU version of the archive utility, especially on Solaris.

Change to the directory where you downloaded Confluence then execute these commands:

```
$ tar -xzf atlassian-confluence-X.X.X.tar.gz -C <installation-directory>
$ cd <installation-directory>
$ tar -xf atlassian-confluence-X.X.X.tar
```

Replace x.x.x with your Confluence version and <installation-directory> with the full path to the directory you created in the last step.

3. Give your dedicated Confluence user read, write and execute permission toyour<installationdirectory>.

In this example we're changing ownership of the installation directory and giving the user conflue nce read, write and execute permissions.

```
$ chown -R confluence <installation-directory>
$ chmod -R u=rwx,go-rwx <installation-directory>
```

3. Create the home directory

Create your home directory this is where Confluence application data like logs, search indexes and
files will be stored. This should be separate to your installation directory, with nospaces or special
characters in the path. We'll refer to this directory as your<home-directory>.

In this example we'll call our home directory confluence-home:

```
$ mkdir confluence-home
```

2. Give your dedicated Confluence user read, write and execute permissions to the <home-directory>.

In this example we're changing ownership of the home directory and giving the user confluence read, write and execute permissions.

```
$ chown -R confluence <home-directory>
$ chmod -R u=rwx,go-rwx <home-directory>
```

- Edit<installation-directory>\confluence\WEB-INF\classes\confluence-init.
 properties.
- 4. At the bottom of the file, enter the absolute path to your<home-directory>. This tells Confluence where to find your<home-directory>when it starts up.

You can edit the confluence.init.properties file any text editor.

a. Scroll to the bottom of the text and find this line:

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

b. Remove the # and the space at the beginning of this line (so Confluence doesn't read the line as a comment) and add the absolute path to your home directory (not a symlink). For example:

```
confluence.home=/var/confluence-home
```

4. Check the ports

By default Confluence listens on port8090. If you have another application running on your server that uses the same ports, you'll need to tell Confluence to use a different port.

To change the ports:

- 1. Edit <installation-directory>\conf\server.xml
- 2. Change the **Server** port (8000) and the **Connector** port (8090) to free ports on your server.

In the example below we've changed the **Server** port to 5000 and the **Connector** port to 5050.

Linux won't allow you to bind to ports less than 1024. If you want to run Confluence on port 80, for example, you could use a reverse proxy to redirect traffic from port 80. See Using Apache with mod_proxy.

5. Start Confluence

1. Run<installation-directory>/bin/start-confluence.shto start the setup process.

We recommend running Confluence as your dedicated user.

```
$ su -u <user>
$ ./start-confluence.sh
```

If you're using Ubuntu the command is a little different:

```
$ sudo su <user>
$ ./start-confluence.sh
```

- 2. Go tohttp://localhost:8090/to launch Confluence in your browser (change the port if you've updated the Connector port).
 - Check your JAVA_HOME variable is set correctly.
 - If you see an error, see Confluence does not start due to Spring Application context has not been set for troubleshooting options.

Set up Confluence

6. Choose installation type

- 1. ChooseProduction installation.
- Choose anyappsyou'd also like to install.

7. Enter your license

Follow the prompts to log in tomy.atlassian.comto retrieve your license, or enter a license key.

8. Connect to your database

- 1. If you've not already done so, it's time to create your database. See the 'Before you begin' section of this page for details and connection options.
- 2. For MySQL and Oracle, follow the prompts to download and install the required driver.
- 3. Enter your database details. Use**test connection** to check your database is set up correctly.

If you want to specify particular parameters, you can choose to connect **By connection string**. You'll be prompted to enter:

• **Database URL** the JDBC URL for your database. If you're not sure, check the documentation for your database.

 Username and Password A valid username and password that Confluence can use to access your database.

9. Populate your new site with content

Choose whether you'd like Confluence to populate your site with content:

This option will create a space that you and your users can use to get to know Confluence. You can delete this space at any time.

Use this option if you have a full site export of an existing Confluence site. This is useful when youre migrating to another database or setting up a test site.

Good to know:

- You can only import sites from the same or earlier Confluence version.
- The system administrator account and all other user data and content will be imported from your previous installation.

In the setup wizard:

- Upload a backup file use this option if your site export file is small (25mb or less).
- Restore a backup file from the file system use this option if your backup file is large. Drop the file into your <confluence-home>/restoredirectory then follow the prompts to restore the backup.
- **Build Index** well need to build an index before your imported content is searchable. This can take a long time for large sites, so deselect this option if you would rather build the index later. Your content won't be searchable until the index is built.

10. Choose where to manage users

Choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Management:

Choose this option if you're happy to manage users in Confluence, or don't have a Jira application installed.

Good to know:

- If you do plan to manage users in a Jira application, but have not yet installed it, we recommend installing Jira first, and then returning to the Confluence setup.
- You can add external user management (for example LDAP, Crowd or Jira) later if you choose.

Choose this option if you have a Jira application installed and want to manage users across both applications.

Good to know:

- 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.
- You'll be able to specify exactly which groups in your Jira app should also be allowed to log in to Confluence. Your license tiers do not need to be the same for each application.
- You'll need either Jira 4.3 or later, Jira Core 7.0 or later, Jira Software 7.0 or later, or Jira Service Management 3.0 or later.

In the setup wizard:

- Jira Base URL the address of your Jira server, such as http://www.example.com:8080/jira/Orhttp://jira.example.com/
- Jira Administrator Login this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application. Confluence will also use this username and password to create a local administrator account which will let you access

Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.

- Confluence Base URL this is the URL Jira will use to access your Confluence server. The URL you give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira application.
- User Groups these are the Jira groups whose members should be allowed to use Confluence.
 Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
 - Jira 6.4 and earlier: jira-users.
 - Jira Software 7.x and later: jira-software-users
 - Jira Core 7.x and later: jira-core-users
 - Jira Service Management (formerly Jira Service Desk) 3.x and later: jira-servicedesk-users
- Admin Groups provide one or more Jira groups whose members should have administrative access to Confluence. The default group is jira-administrators. These groups will get the system administrator and Confluence administrator global permissions in Confluence.

11. Create your administrator account

Enter details for the administrator account.

Skip this step if you chose to manage users in a Jira application or you imported data from an existing site.

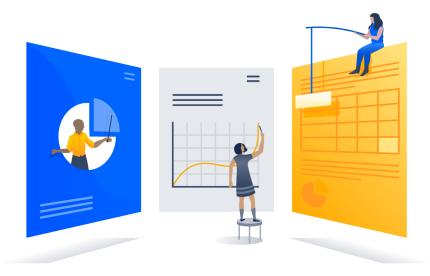
12. Start using Confluence

That's it!Your Confluence site is accessible from a URL like this: http://<computer_name_or_IP_address>: <port>

If you plan to run Confluence behind a reverse proxy, check outProxy and SSL considerationsbefore you go any further.

Here's a few things that will help you get your team up and running:

- Set the server base URL this is the URL people will use to access Confluence.
- Set up a mail server this allows Confluence to send people notification about content.
- Add and invite users get your team on board!
- Start and stop Confluence find out how to start and stop Confluence.



Troubleshooting

- Check your JAVA_HOME is set correctly.
- If you see an error, see Confluence does not start due to Spring Application context has not been set for troubleshooting options.
- Use a GNU version of the unzip utility. There are known issues extracting the tar.gz file on Solaris and AIX. See 'extractBundledPlugins Couldn't find atlassian-bundled-plugins.zip on classpath' Due to Solaris TAR Utility.

• Collaborative editing errors? See Troubleshooting Collaborative Editing.

Head to Installation Troubleshooting in our Knowledge Base for more help.

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.

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.

Unattended installation

If you've previously installed Confluence using the Windows or Linux installer, you can use a configuration file from your existing Confluence installation (response.varfile) to re-install Confluence in unattended mode, no user input required.



This can be useful when you have installed Confluence on a test server and are ready to install on your production server with the same configuration.

Good to know

- Theresponse.varfilefile contains the options specified during the installation wizard steps of your previous Confluence installation. Don't uninstall your previous Confluence installation until after you've copied this file to your new install location.
- If you decide to modify theresponse.varfilefile, make sure all directory paths specified are absolute, for example,sys.installationDir=C\:\\Program Files\\Atlassian\\Confluence (Windows) orsys.installationDir=/opt/atlassian/confluence(Linux).

Unattended installations will fail the file contains relative directory paths.

 It's not possible to automate the database configuration step. This must be done via the setup wizard in your browser.

Install Confluence in unattended mode

These steps cover where you have an existing Confluence installation.

- 1. Downloadthe appropriate installer for your operating system.
- 2. Copy <installation-directory>/.install4j/response.varfilefrom your existing Confluence installation to where you downloaded the installer.
- 3. In command prompt or terminal change directory (cd) to where you downloaded the installer.
- 4. Run the following command to install Confluence:

```
Windows
> atlassian-confluence-X.X.X-x64.exe -q -varfile response.varfile

Linux
$ atlassian-confluence-X.X.X-x64.bin -q -varfile response.varfile
```

Wherex.x.xis the Confluence version you downloaded.

- -qinstructs the installer to run in unattended mode (quietly). -varfilespecifies the location and name of the configuration file containing the options used by the installer.
- 5. Confluence will start automatically once the silent installation finishes.

Once Confluence is installed, you will still need to head to http://localhost:<port>to finish setting up Confluence.

See the **Set up Confluence** section onInstalling Confluence on Windows orInstalling Confluence on Linux for more info.

Create your own response.varfile

It is also possible to create your own response.varfile, rather than one generated by an existing installation, if you are installing Confluence for the first time.

Example response.varfile

app.confHome=/var/atlassian/application-data/confluence6_15_5
app.install.service\$Boolean=false
portChoice=custom
httpPort\$Long=26112
rmiPort\$Long=8001
launch.application\$Boolean=false
sys.adminRights\$Boolean=true
sys.confirmedUpdateInstallationString=false
sys.installationDir=/opt/atlassian/confluence6_15_5
sys.languageId=en

The following parameters can be included in the file.

Parameters	Accepted values	Description
app.confHome		This is the path to your target local home directory.
app.install. service\$Boolean	truefalse	Determines whether Confluence should be installed as a service.
portChoice	customdefault	Determines whether Confluence should be installed with default ports.
httpPort\$Long		If portChoice is custom, this sets the HTTP connector port in Tomcat.
rmiPort\$Long		If portChoice is custom, this sets the Tomcat server port.
launch. application\$Boolean	• true • false	Determines whether the installer should start Confluence once installation is complete
sys. adminRights\$Boolea n=true	• true • false	Indicates whether the user running the installer has admin privileges on the machine.
sys. confirmedUpdateInst allationString	• true • false	Set this to false for a fresh unattended installation. Set to true to perform an unattended upgrade. Always back up your existing site before attempting to upgrade.

sys.installationDir	path to install directory	This is the path to your target installation directory for a new install, or existing installation directory to be upgraded.
sys.languageld		Default application language.

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 <code>conf/server.xml</code> under your Confluence Installation directory. The first four lines of the file look like this:

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':

To access Confluence in this configuration, point your web browser to http://localhost:8099/.

Final Configuration

- If this is the URL your users will use to access Confluence, update your Base URL to point to the new URL.
- 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.

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 your Jira application's distribution, seeChanging JIRA application TCP ports.
- [3] You will need to restart Confluence after editingserver.xmlfor the changes to take effect.

Start and Stop Confluence

How you start and stop Confluence depends on whether you are running Confluence as a Service.

To check whether Confluence is already running you can go tohttp://sbase-url>/status.

Windows

If you installed Confluence as a service, you can **Start Confluence Server** and **Stop Confluence Server** from the Windows Start menu.

You can't start or stop Confluence manually using the start-confluence.bat and stop-confluence.bat file.

If you didn't install Confluence as a service you'll need to start and stop Confluence manually. The way you do this depends on how Confluence was originally installed.

If you installed Confluence manually, and have Java installed on your server:

- To start Confluence run<installation-directory>\bin\start-confluence.bat
- To stop Confluence run<installation-directory>\bin\stop-confluence.bat

We recommend running Confluence with a dedicated user account. To do this, use use therunascommand to executestart-confluence.bat.

```
> runas /env /user:<DOMAIN>\<confluence> start-confluence.bat
```

Where < DOMAIN > is your Windows domain or computer name and < confluence > is the name of your dedicated user.

If you **installed Confluence using the installer**, and don't have Java installed, use the Start and Stop Confluence options in the Start menu, or:

- To start Confluence run <installation-directory>\startup-bundled-jre.bat
- To stop Confluence run <installation-directory>\shutdown-bundled-jre.bat

It is possible to start Confluence Server with user installed apps temporarily disabled. This is useful if you need to troubleshoot problems with your site, particularly if an app may be preventing Confluence from starting up successfully.

To start Confluence with all user installed apps temporarily disabled:

```
> cd <installation-directory>/bin
> start-confluence.bat /disablealladdons
```

To start Confluence with a particular app temporarily disabled:

```
> cd <installation-directory>/bin
> start-confluence.bat /disableaddon=com.atlassian.test.plugin
```

where <code>com.atlassian.test.plugin</code> is the app key. To disable multiple apps, use a colon separated list. Regex/wildcards are not permitted, the full key of the plugin must be provided.

These parameters are applied at startup only, they do not persist. If you want to permanently disable an app, go to



to do this via UPM.

Notes

- If the app key contains a space, disabling the app using this method will not work, you need to manuall
 y deal with that app.
- This feature does not work for Confluence Data Center.
- replace /bin/start-confluence.bat with startup-bundled-jre.bat if you installed Confluence using the installer, and are using the bundled JRE (Java Runtime Engine).

Linux

If you installed Confluence as a service, use one of the following commands to **start**, **stop** or **restart** Confluence.

```
$ sudo /etc/init.d/confluence start
$ sudo /etc/init.d/confluence stop
$ sudo /etc/init.d/confluence restart
```

You can't start or stop Confluence manually using the start-confluence.sh and stop-confluence.sh files.

If you didn't install Confluence as a service you'll need to start and stop Confluence manually.

- To start Confluence run<installation-directory>\bin\start-confluence.sh
- $\bullet \ \ \textbf{To stop Confluence run} < \texttt{installation-directory} \\ \verb|\bin\stop-confluence.sh| \\$

We recommend running Confluence with a dedicated user account:

```
$ su -u <user>
$ ./start-confluence.sh
```

Where <user> is the name of your dedicated user.

If you're using Ubuntu the command is a little different:

```
$ sudo su <user>
$ ./start-confluence.sh
```

It is possible to start Confluence with user installed apps temporarily disabled. This is useful if you need to troubleshoot problems with your site, particularly if an app may be preventing Confluence from starting up successfully.

To start Confluence with **all** user installed apps temporarily disabled:

```
$ cd <installation-directory>/bin
$ ./start-confluence.sh --disable-all-addons
```

To start Confluence with a particular app temporarily disabled:

```
$ cd <installation-directory>/bin
$ ./start-confluence.sh --disable-addons=com.atlassian.test.plugin
```

where com.atlassian.test.plugin is the app key.

To disable multiple apps, use a colon separated list, for example, com.atlassian.test.plugin:com.atlassian.another.plugin. Regex/wildcards are not permitted, the full key of the plugin must be provided.

These parameters are applied at startup only, they do not persist. If you want to permanently disable an app, go to

O > Manage apps

to do this via UPM.

Notes

- If the app key contains a space, disabling the app using this method will not work, you need to manuall y deal with that app.
- This feature does **not** work for Confluence Data Center.

Installing Confluence Data Center

In this guide we'll run you through installing Confluence Data Center in a Windows or Linux Environment. You can run Data Center as a standalone installation, or in a cluster, depending on your organisation's needs.

This guide covers installing for the first time, with no existing data. If you already have a Confluence Server instance, see Migrate from Server to Data Center.



On this page:

Before you begin Supported platforms

Requirements

Install Confluence Data Center

Install Confluence Data Center in a cluster

Terminology

Install and set up Confluence Add more Confluence nodes

Security

Troubleshooting

Upgrading a cluster

Other ways to install Confluence Data Center:

- Kubernetes install on a Kubernetes cluster using our Helm charts
- AWS- hassle free deployment in AWS using our Quick Start
- Azure reference templates for Microsoft Azure deployment
- Move to Data Center- for existing Confluence Server sites

Interested in learning more about Data Center? Find out more about the benefits of Confluence Data Center.

Before you begin

Supported platforms

See our Supported Platformspage for information on the database, Java, and operating systems you'll be able to use. These requirements are the same for Server and Data Center deployments.

Requirements

To use Confluence Data Center you must:

- Have a Data Center license (you canpurchase a Data Center licenseor create an evaluation license at my.atlassian.com)
- Use asupportedexternal database, operating system and Java version
- Use OAuth authentication if you haveapplication linksto other Atlassian products (such as Jira)

To run Confluence in a cluster you must also:

- Use a load balancer with session affinity in front of the Confluence cluster. WebSockets support is also recommended for collaborative editing.
- Have a shared directory accessible to all cluster nodes in the same path (this will be your shared home directory). This must be a separate directory, and not located within the local home or install directory.



Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Install Confluence Data Center

If your organization doesn't need high availability or disaster recovery capabilities right now, you can install Confluence Data Center without setting up a cluster.

To install Confluence Data Center, without setting up a cluster, follow the instructions for Confluence Server:

- Installing Confluence on Windows
- Installing Confluence on Linux

The process is almost identical to an ordinary Confluence Server installation, just be sure to choose Standal oneafter you've entered your Data Center license.

Install Confluence Data Center in a cluster

If your organization requirescontinuous uptime, scalability, and performance under heavy load, you'll want to run Confluence Data Center in a cluster.

SeeClustering with Confluence Data Centerfor a complete overview of hardware and infrastructure considerations.

Terminology

In this guide we'll use the following terminology:

- Installation directory The directory where you installed Confluence.
- Local home directory The home or data directory stored locally on each cluster node (if Confluence is not running in a cluster, this is simply known as the home directory).
- Shared home directory The directory you created that is accessible to all nodes in the cluster via the same path.

At the end of the installation process, you'll have an installation and local home directory on each node, and a single shared home directory (a total of 5 directories in a two node cluster) for Confluence plus directories for Synchrony.

Install and set up Confluence

1. Install Confluence on the first node

- 1. Install Confluence on node 1 See Installing Confluence on Windows from Zip Fileor Installing Confluence on Linux from Archive File for more information.
- 2. Start Confluence on Node 1
- 3. Follow the prompts to enter your Data Center license then choose Clustered as the deployment type.
- 4. The setup wizard will prompt you to configure the cluster, by entering:
 - A name for your cluster
 - The path to the shared home directory you created earlier
 - The network interface Confluence will use to communicate between nodes
 - How you want Confluence to discover cluster nodes:
 - Multicast enter your own multicast address or automatically generate one.

- TCP/IP enter the IP address of each cluster node
- AWS enter your IAM Role or secret key, and region.

We recommend using our Quick Start or Cloud Formation Template to deploy Confluence Data Center in AWS, as it will automatically provision, configure and connect everything you need.

If you do decide to do your own custom deployment, you can provide the following information to allow Confluence to auto-discover cluster nodes:

Field	Description
IAM Role or Secret Key	This is your authentication method. You can choose to authenticate by IAM Role or Secret Key.
Region	This is the region your cluster nodes (EC2 instances) will be running in.
Host header	Optional. This is the AWS endpoint for Confluence to use (the address where the EC2 API can be found, for example 'ec2.amazonaws.com'). Leave blank to use the default endpoint.
Securit y group name	Optional. Use to narrow the members of your cluster to only resources in a particular security group (specified in the EC2 console).
Tag key and Ta g value	Optional. Use to narrow the members of your cluster to only resources with particular tags (specified in the EC2 console).

- 5. Follow the prompts to set up your database and administrator account.
- Confirm that you can log in to Confluence and everything is working as expected, then stop Confluence on Node 1.

Add more Confluence nodes

2.Copy Confluence to second node

To copy Confluence to the second node:

- 1. Shut down Confluence on node 1.
- 2. Copy the installation directory from node 1 to node 2.
- 3. Copy the local home directory from node 1 to node 2.

Copying the local home directory ensures the Confluence search index, the database and cluster configuration, and any other settings are copied to node 2.

3. Configure load balancer

Configure your load balancer for Confluence. You can use the load balancer of your choice, but it needs to support session affinity and WebSockets.

You can verify that your load balancer is sending requests correctly to your existing Confluence server by accessing Confluence through the load balancer and creating a page, then checking that this page can be viewed/edited by another machine through the load balancer.

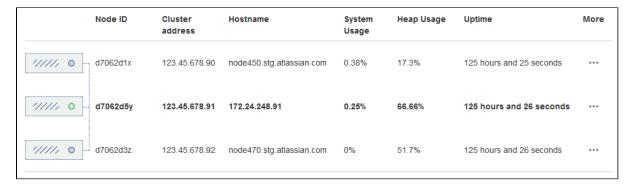
4. Start Confluence one node at a time

You must only start Confluence **one node at a time**. The first node must be up and available before starting the next one.

- 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.

TheCluster monitoring console (Seneral Configuration > Clustering) shows information about the active cluster.

When the cluster is running properly, this page displays the details of each node, including system usage and uptime. Use the *** menu to see more information about each node in the cluster.



5. Test your Confluence cluster

To test creating content you'll need to access Confluence via your load balancer URL. You can't create or edit pages when accessing a node directly.

A simple process to ensure your cluster is working correctly is:

- 1. Access a node via your load balancer URL, and create a new document on this node.
- 2. Ensure the new document is visible by accessing it directly on a different node.
- 3. Search for the new document on the original node, and ensure it appears.
- 4. Search for the new document on another node, and ensure it appears.
- (i) 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.

6. Set up your Synchrony cluster (optional)

Synchrony is required for collaborative editing. You have two options for running Synchrony with a Data Center license:

- managed by Confluence(recommended)
 - This is the default setup. Confluence will automatically launch a Synchrony process on the same node, and manage it for you. No manual steps are required.
- Standalone Synchrony cluster (managed by you)
 You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need.
 Significant setup is required. See Set up a Synchrony cluster for Confluence Data Center for a step-by-step guide.

Head to Administering Collaborative Editing to find out more about collaborative editing.

Security

Ensure that only permittedcluster nodes are allowed to connect to the following portsthrough the use of a firewall and / or network segregation:

- 5801 -Hazelcastport for Confluence
- 5701 -Hazelcastport for Synchrony
- 25500 Cluster base port for Synchrony

If you use multicast for cluster discovery:

54327- Multicast port for Synchrony (only required if running Synchrony standalone cluster)

Troubleshooting

If you have problems with the above procedure, please see our Cluster Troubleshooting guide.

If you're testing Confluence Data Center by running the cluster on a single machine, please refer to our developer instructions on Starting a Confluence cluster on a single machine.

Upgrading a cluster

It's important that upgrades follow the procedure for Upgrading Confluence Data Center.

Upgrading Confluence Data Center

This page contains instructions for upgrading an existing Confluence cluster.

If you are not running Confluence in a cluster, follow the instructions in Upgrading Confluence.

If you're running Confluence in a cluster in AWS, follow the instructions inRunning Confluence Data Center in AWS.

If you are upgrading to the next bug fix update (for example, 7.9.0 to 7.9.3), you can do so with no downtime. Follow the instructions in Upgrade Confluence without downtime.

In this guide we'll use the following terminology:

- Installation directory The directory where you installed Confluence.
- Local home directory The home or data directory stored locally on each cluster node (if Confluence is not running in a cluster, this is simply known as the home directory).
- Shared home directory The directory you created that is accessible to all nodes in the cluster via the same path.

Currently using Confluence Server? Learn more about thebenefits of Confluence Data Center.

On this page:

- 1. Back up
- 2. Download Confluence
- 3. Stop the cluster
- 4. Upgrade the first node
- 5. Upgrade Synchrony (optional)
- 6. Copy Confluence to remaining nodes
- 7. Start Confluence and check cluster connectivity

1. Back up

We strongly recommend that you backup your Confluence home and install directories and your database before proceeding.

More information on specific files and directories to backup can be found in Upgrading Confluence.

2. Download Confluence

Download the appropriate file for your operating system fromhttps://www.atlassian.com/software/confluence/download

3. Stop the cluster

You must stop all the nodes in the cluster before upgrading.

We recommend configuring your load balancer to redirect traffic away from Confluence until the upgrade is complete on all nodes.

4. Upgrade the first node

To upgrade the first node:

- 1. Extract (unzip) the files to adirectory (this will be your new installation directory, and must be different to your existing installation directory)
- 2. Update the following line in the<Installation-Directory>\confluence\WEB-INF\classes\confluence-init.propertiesfile to point to theexistinglocal homedirectory on that node.

- 3. If your deployment uses a MySQL database, copy the jdbc driver jar file from your existing Confluence installation directory to confluence/WEB-INF/libin your new installation directory. The jdbc driver will be located in either the<Install-Directory>/common/libor<Installation-Directory>/confluence/WEB-INF/libdirectories. SeeDatabase Setup For MySQL for more details.
- 4. If you run Confluence as a service:
 - On Windows, delete the existing service then re-install the serviceby running<install-directory>/bin/service.bat.
 - On Linux, update the service to point to the new installation directory (or use symbolic links to do this).
- 5. Copy any other immediately required customizations from the old version to the new one(for example if you are not running Confluence on the default ports or if you manage users externally, you'll need to update / copy the relevant files find out more in Upgrading Confluence Manually).
 - (i) If you configured Confluence to run as a Windows or Linux service, don't forget to update its service configuration as well. For related information, see Start Confluence Automatically on Windows as a Service or Run Confluence as a systemd service on linux.
- 6. Start Confluence, and confirm that you can log in and view pages before continuing to the next step.

You should now stop Confluence, and reapply any additional customizations from the old version to the new version, before upgrading the remaining nodes.

5. Upgrade Synchrony (optional)

If you've chosen to let Confluence manage Synchrony for you (recommended), you don't need to do anything. Synchrony was automatically upgraded with Confluence.

If you're running your own Synchrony cluster, you should:

- 1. Grab the newsynchrony-standalone.jarfrom the<local-home>directory on your upgraded Confluence node.
- 2. Copy the newsynchrony-standalone.jarto each of your Synchrony nodes, and start Synchrony as normal.

6. Copy Confluence to remaining nodes

The next step is to replicate your upgraded Confluence directories to other nodes in the cluster.

- 1. Copy the installation directory and local home directory from the first node to the next node.
- 2. If the path to the local home directory is different on this node, edit the confluence-init. properties to point to the correct location.
- 3. Start Confluence, and and confirm that you can log in and view pages on this node.

Stop Confluence on this node, then repeat this process for each remaining node.

7. Start Confluence and check cluster connectivity

Once all nodes have been upgraded you can start Confluence Data Center on each node, **one at a time** (starting up multiple nodes simultaneously can lead to serious failures).

TheCluster monitoring console (Seneral Configuration > Clustering) includes information about the active cluster nodes. When the cluster is running properly, you should be able to see the details of each node.

Adding and Removing Data Center Nodes

Your Data Center license is based on the number of users in your cluster, rather than the number of nodes. This meansyou can add and remove nodes from your Data Center cluster at any time.

If you deployed Confluence Data Center on AWS using the Quick Start, your Confluence and Synchrony nodes will be inauto-scaling groups. You will add and remove nodes in the AWS console either by changing the minimum and maximum size of each group or using a scaling plan.

On this page:

- Adding a node
- Removing a node
- Changing the node identifier
- Moving to a non-clustered installation

Adding a node

To add a node:

- 1. Copy the installation directory and local home directory from the stopped node to your new node.
- Start Confluence on your new node.
 During the startup process Confluence will recover indexes from a running node to bring the new node up to date.
- 3. Go to General Configuration > Clustering and check that the new node is visible.

You should only start one node at a time. Starting up multiple nodes simultaneously can cause serious failures.

If the discovery mode is set to TCP/IP, you'll need to update the <code>confluence.cluster.peers</code> property in the <code>confluence.cfg.xml</code> file for each node so the file lists all nodes in your cluster:

<property name="confluence.cluster.peers">[node 1 IP],[node 2 IP],[node 3 IP]/property> <!-- A commaseparated list of node IP addresses, without spaces -->

Removing a node

To remove a node, stop Confluence on that node. You can then remove the installation and local home directory as required.

To see the number of nodes remaining go to General Configuration>Clustering.

Changing the node identifier

Confluence generates an identifier for each node in your cluster. You can use the confluence.cluster. node.namesystem propertyto set the node identifier on each node so that it's easier for your users and administrators to read.

SeeConfiguring System Properties for more information on how to set the system property.

Moving to a non-clustered installation

If you no longer need clustering, and want to avoid the overhead that comes from running a cluster with just one node, you can go back to a non-clustered Data Center installation. You'll need to make some infrastructure changes as part of the switch.

SeeMove to a non-clustered installation to find out how to do this.

Change Node Discovery from Multicast to TCP/IP or AWS

On this page:

- To change from multicast to TCP/IP
- To change from multicast to AWS
- To change from TCP/IP to AWS
- To change from TCP/IP to multicast
- Reference of properties in the confluence.cfg.xml file

If you're setting up Confluence Data Center for the first time, it'll step you through the process of choosing your discovery mode and adding cluster nodes. If you decide to change the node discovery for the cluster, you'll need to edit theconfluence.cfg.xmlfile in the local home directory of each cluster node.



- Before you make any changes, shut down all nodes in your cluster
- Make sure the discovery configuration is exactly the same for each node (make the same changes to theconfluence.cfg.xml filein each local home directory)
- Always perform a safety backup before making manual edits to these files

The changes you need to make may differ slightly, depending on whether you've upgraded from an older version of Confluence Data Center or if you've started with version 5.9. We've detailed both methods, below.

To change from multicast to TCP/IP

Look for the following two lines in theconfluence.cfg.xmlfile:

If both lines exist in the file, change them to the lines below; where theconfluence.cluster.addresspropert y exists, butthere's no reference to theconfluence.cluster.join.typeproperty, update the first line and add the second line as shown below.

```
<property name="confluence.cluster.peers">[node 1 IP],[node 2 IP],[node 3 IP]/property> <!-- A comma-
separated list of node IP addresses, without spaces -->
<property name="confluence.cluster.join.type">tcp_ip/property> <!-- accepted values are multicast or
tcp_ip -->
```

Enter the address of each node, and separate each address with a comma. Please, make sure to remove the brackets from around the IP addresses.

You can now restart your cluster nodes.

To change from multicast to AWS

Look for the following two lines in theconfluence.cfg.xmlfile and remove them:

Depending on which type of credentials you are passing to Confluence, you will addone of the following two blocks with your AWS configuration.

Option 1: For Access Key/Secret Key based credentials:

Option 2: For IAM role based credentials:

```
<property name="confluence.cluster.join.type">aws</property>
<property name="confluence.cluster.aws.host.header">[---VALUE---]</property>
<property name="confluence.cluster.aws.region">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.value">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.value">[---VALUE---]</property>
<property name="confluence.cluster.aws.iam.role">[---VALUE---]</property></property></property>
```

To change from TCP/IP to AWS

Look for the following two lines in theconfluence.cfg.xmlfile and remove them:

Depending on which type of credentials you are passing to Confluence, you will addone of the following two blocks with your AWS configuration.

Option 1: For Access Key/Secret Key based credentials:

```
<property name="confluence.cluster.join.type">aws</property>
<property name="confluence.cluster.aws.host.header">[---VALUE---]</property>
<property name="confluence.cluster.aws.region">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.value">[---VALUE---]</property>
<property name="confluence.cluster.aws.access.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.access.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.secret.key">[---VALUE---]</property>
```

Option 2: For IAM role based credentials:

```
<property name="confluence.cluster.join.type">aws</property>
<property name="confluence.cluster.aws.host.header">[---VALUE---]</property>
<property name="confluence.cluster.aws.region">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.value">[---VALUE---]</property>
<property name="confluence.cluster.aws.iam.role">[---VALUE---]</property>
```

You can now restart your cluster nodes.

Note that if you're using a CloudFormation YAML template you need to make sure you have these appropriate values as a minimum and they should be reflected on the AWS side as well. If you switch to AWS mode cluster type, please also reviewRunning Confluence Data Center in AWS and make sure you have the following set up in your YAML:

```
Key: Cluster
Value: !Ref AWS::StackName
PropagateAtLaunch: true
```

To change from TCP/IP to multicast

To switch from TCP/IP to multicast, just perform the reverse of the changes outlined above.

Reference of properties in the confluence.cfg.xml file

key	valid values	notes
confluence. cluster.join. type	'multicast' or 'tcp_ip' or 'aws'	Pre-5.9 Data Center installations won't have this key. By default, if the key is missing, Confluence will choose multicast
confluence. cluster. address	a single multicast IP address	This key is only used by confluence if confluence. cluster.join.type is set to multicast
confluence. cluster.peers	a comma-separated string of IP addresses (no spaces)	There must be at least one address here. The addresses are the IP address of each node in the cluster, for example <pre></pre>

Running Confluence Data Center in AWS

If you decide to deploy Confluence Data Center in a clustered environment, consider usingAmazon Web Services (AWS). AWS allows you to scale your deployment elastically by resizing and quickly launching additional nodes, and provides a number of managed services that work out of the box with Confluence Data Center. These services make it easier to configure, manage, and maintain your deployment's clustered infrastructure.

Interested in learning more about the benefits of Data Center? Check out our overview of Confluence Data Center.

On this page:

- Non-clustered VS clustered environment
- Deploying Confluence Data Center in a cluster using the AWS Quick Start
 - Advanced customizations
 - Launching the Quick Start from your own S3 bucket (recommended)
 - Amazon Aurora database for high availability
 - Synchrony setup
 - Amazon CloudWatch for basic monitoring and centralized logging
 - Auto Scaling groups
 - EC2 sizing recommendations
 - Customizing the AWS Quick Start's CloudFormation templates
- Supported AWS regions
- Internal domain name routing with Route53 Private Hosted Zones
 - Step 1: Create a new hosted zone
 - Step 2: Configure your stack to use the hosted zone
 - Step 3: Link your DNS server to the Confluence sites VPC
- Scaling up and down
- Connecting to your nodes over SSH
- Upgrading
 - Step 1: Terminate all running Confluence Data Center application nodes
 - Step 2: Update the version used by your Confluence Data Center stack
 - Step 3: Scale up the number of application nodes
- Backing up
- Migrating your existing Confluence site to AWS

Non-clustered VS clustered environment

A single node is adequate for most Small or Medium size deployments, unless you need specific features that require clustering (for example, high availability). If you have an existing Server installation, you can still use its infrastructure when you upgrade to Data Center. Many features exclusive to Data Center (likeSAML single sign-on, self-protection via rate limiting, andCDN support) don't require clustered infrastructure. You can start using these Data Center features by simply upgrading your Server installations license. For more information on whether clustering is right for you, check outAtlassian Data Center architecture and infrastructure options.

Deploying Confluence Data Center in a cluster using the AWS Quick Start

The simplest way to deploy your entire Data Center cluster in AWS is by using the Quick Start. The Quick Start launches, configures, and runs the AWS compute, network, storage, and other services required to deploy a specific workload on AWS, using AWS best practices for security and availability.

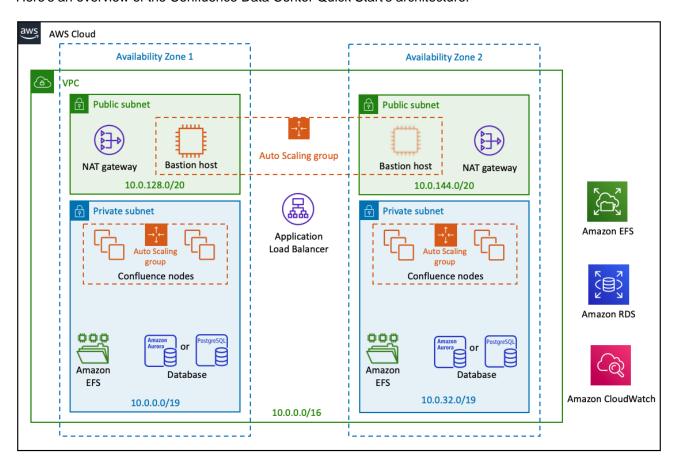
The Quick Start provides two deployment options, each with its own template. The first option deploys the Atlassian Standard Infrastructure (ASI) and then provisions Confluence Data Center into this ASI. The second option only provisions Confluence Data Center on an existing ASI.



Atlassian Standard Infrastructure (ASI)

The ASI is a virtual private cloud (VPC) that contains the components required by all Atlassian Data Center applications. For more information, see Atlassian Standard Infrastructure (ASI) on AWS.

Here's an overview of the Confluence Data Center Quick Start's architecture:



The deployment consists of the following components:

- Instances/nodes: One or more Amazon Elastic Cloud (EC2) instances as cluster nodes, running Confluence.
- Load balancer: An Application Load Balancer (ALB), which works both as load balancer and SSLterminating reverse proxy.
- Amazon EFS: A shared file system for storing artifacts in a common location, accessible to multiple Confluence nodes. The Quick Start architecture implements the shared file system using the highly available Amazon Elastic File System (Amazon EFS) service.
- Database: Your choice of shared database instance Amazon RDS or Amazon Aurora.
- Amazon CloudWatch: Basic monitoring and centralized logging through Amazon's native CloudWatch service.

For more information on the architecture, components and deployment process, see our Quick Start Guide.



Confluence will use the Java Runtime Engine (JRE) that is bundled with Confluence (/opt/atlassian /confluence/jre/), and not the JRE that is installed on the EC2 instances (/usr/lib/jvm/jre/).

Advanced customizations

To get you up and running as quickly as possible, the Quick Start doesn't allow the same level of customization as a manual installation. You can, however, further customize your deployment through the variables in the Ansible playbooks we use.

All of our AWS Quick Starts use Ansible playbooks to configure specific components of your deployment. These playbooks are available publicly on this repository:

https://bitbucket.org/atlassian/dc-deployments-automation

You can override these configurations by using Ansible variables. Refer to the repositorys README file for more information.

Launching the Quick Start from your own S3 bucket (recommended)

The fastest way to launch the Quick Start is directly from its AWS S3 bucket. However, when you do, any updates we make to the Quick Start templates will propagate directly to your deployment. These updates sometimes involve adding or removing parameters from the templates. This could introduce unexpected (and possibly breaking) changes to your deployment.

For production environments, we recommend that you copy the Quick Start templates into your own S3 bucket. Then, launch them directly from there. Doing this gives you control over when to propagate Quick Start updates to your deployment.

 Clone the Quick Start templates (including all of its submodules) to your local machine. From the command line, run:

```
git clone --recurse-submoduleshttps://github.com/aws-quickstart
/quickstart-atlassian-confluence.git
```

2. (Optional)TheQuick Start templates repository uses the directory structure required by the Quick Start interface. If needed (for example, to minimize storage costs), you can remove all other files except the following:

```
quickstart-atlassian-confluence
submodules
quickstart-atlassian-services
templates
quickstart-vpc-for-atlassian-services.yaml
templates
quickstart-confluence-master-with-vpc.template.yaml
quickstart-confluence-master.template.yaml
```

- Install and set up the AWS Command Line Interface. This tool will allow you to create an S3 bucket and upload content to it.
- 4. Create an S3 bucket in your region:

```
aws s3 mb s3://<bucket-name> --region <AWS_REGION>
```

At this point, you can now upload the Quick Start templates to your own S3 bucket. Before you do, you'll have to choose which Quick Start template youll be using:

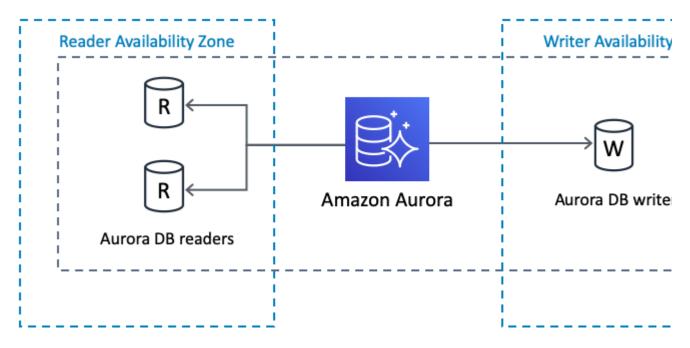
- quickstart-confluence-master-with-vpc.template.yaml: use this for deploying into a new ASI (end-to-end deployment).
- quickstart-confluence-master.template.yaml: use this for deploying into an existing ASI.
- 1. In the template youve chosen, the QSS3BucketNamedefault value is set to aws-quickstart. Replace this default with the name of your S3 bucket.
- 2. Go into the parent directory of yourlocalclone of the Quick Start templates. From there, upload all the files in local clone to your S3 bucket:

aws s3 cp quickstart-atlassian-confluence s3://<bucket-name> --recursive --acl public-read

3. Once youve uploaded everything, youre ready to deploy your production stack from your S3 bucket. Go to Cloudformation Create Stack. When specifying a template, paste in the Object URL of the Quick Start template youll be using.

Amazon Aurora database for high availability

The Quick Start also allows you to deploy Confluence Data Center with an Amazon Aurora clustered database (instead of RDS). This cluster will be PostgreSQL-compatible, featuring a primary database writer that replicates to two database readers. You can also set up the writers and readers in separate availability zones for better resiliency.



If the writer fails, Aurora automatically promotes one of the readers to take its place. For more information, seeAmazon Aurora Features: PostgreSQL-Compatible Edition.



⚠ If you want to set up an existing Confluence Data Center instance with Amazon Aurora, youll need. to perform some extra steps. See Configuring Confluence Data Center to work with Amazon Aurora for detailed instructions.

Synchrony setup

If you have a Confluence Data Center license, two methods are available for running Synchrony:

- managed by Confluence(recommended) Confluence will automatically launch a Synchrony process on the same node, and manage it for you. No manual setup is required.
- Standalone Synchrony cluster (managed by you) You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need. Significant setup is required. During a rolling upgrade, you'll need to upgrade the Synchrony separately from the Confluence cluster.

If you want simple setup and maintenance, we recommend allowing Confluence to manage Synchrony for you. If you want full control, or if making sure the editor is highly available is essential, then managing Synchrony in its own cluster may be the right solution for your organisation.

By default, the Quick Start will configure Synchrony to be managed by Confluence. However, you can use the Quick Start to configure standalone Synchrony. When you do, the Quick Start creates an Auto Scaling group containing one or more Amazon EC2 instances as cluster nodes, running Synchrony.

For more information about Synchrony configuration, seePossible Confluence and Synchrony Configurations.



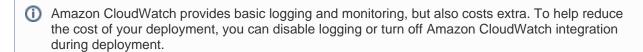
Managed mode is only available in 6.12 and later

If you plan to deploy a Confluence Data Center version earlier than 6.12, you can only use Standalone mode. In the Quick Start, this means you should set your Collaborative editing mode to synchrony-separate-nodes.

Amazon CloudWatch for basic monitoring and centralized logging

The Quick Start can also provide you with node monitoring through Amazon Cloud Watch. This will allow you to track each node's CPU, disk, and network activity, all from a pre-configuredCloudWatch dashboard. The dashboard will be configured to display the latest log output, and you can customize the dashboard later on with additional monitoring and metrics.

By default, Amazon CloudWatch will also collect and store logs from each node into a single, central source. This centralized logging allows you to search and analyze your deployment's log data more easily and effectively. See Analyzing Log Data with CloudWatch Logs Insights and Search Log Data Using Filter Patterns for more information.





✓ To download your log data (for example, to archive it or analyze it outside of AWS), youll have to export it first to S3. From there, you can download it. See Exporting Log Data to Amazon S3 for details.

Auto Scaling groups

This Quick Start uses Auto Scaling groups, but only to statically control the number of its cluster nodes. We don't recommend that you use Auto Scaling to dynamically scale the size of your cluster. Adding an application node to the cluster usually takes more than 20 minutes, which isn't fast enough to address sudden load spikes.

If you can identify any periods of high and low load, you can schedule the application node cluster to scale accordingly. SeeScheduled Scaling for Amazon EC2 Auto Scalingfor more information.

To study trends in your organization's load, you'll need to monitor the performance of your deployment. Refer toConfluence Data Center sample deployment and monitoring strategyfor tips on how to do so.

EC2 sizing recommendations

For Large or XLarge deployments, check out our AWS infrastructure recommendations for application, Synchrony, and database sizing advice. For smaller deployments, you can use instances that meet Confluence's system requirements. Smaller instance types (micro, small, medium) are generally not adequate for running Confluence.

Customizing the AWS Quick Start's CloudFormation templates

To get you up and running as guickly as possible, the Quick Start doesn't allow the same level of customization as a manual installation. Alternatively, you can customize the CloudFormation templates used by the Quick Start to fit your needs. These templates are available from the following repository:

https://github.com/aws-quickstart/quickstart-atlassian-confluence

Supported AWS regions

Not all regions offer the services required to run Confluence. You'll need to choose a region that supports Amazon Elastic File System (EFS). You can currently deploy Confluence using the Quick Start in the following regions:

- Americas
 - Northern Virginia
 - Ohio
 - Oregon
 - Northern California
 - Montreal
- Europe/Middle East/Africa
 - Ireland
 - Frankfurt
 - London
 - Paris
- Asia Pacific
 - Singapore
 - Tokyo
 - Sydney
 - Seoul
 - Mumbai

This list was last updated on 20 Jun 2019.

The services offered in each region change from time to time. If your preferred region isn't on this list, check the Regional Product Services table in the AWS documentation to see if it already supports EFS.



AWS GovCloud is not for production

Our Quick Starts are also available on AWS GovCloudregions, but only for testing purposes. We do not support any deployments on GovCloud regions resulting from this Quick Start.

There is an additional dependency for Confluence versions earlier than 6.3.2. Synchrony (which is required for collaborative editing) uses a third party library to interact with the Amazon API, and the correct endpoints are not available in all regions. This means you can't run Synchrony in the following regions:

- US East (Ohio)
- EU (London)¹
- Asia Pacific (Mumbai) ¹
- Asia Pacific (Seoul) ¹
- Canada (Central) ¹

Internal domain name routing with Route53 Private Hosted Zones

Even if your Confluence site is hosted on AWS, you can still link its DNS with an internal, on-premise DNS server (if you have one). You can do this through Amazon Route 53, creating a link between the public DNS and internal DNS. This will make it easier to access your infrastructure resources (database, shared home, and the like) through friendly domain names. You can make those domain names accessible externally or internally, depending on your DNS preferences.

Step 1: Create a new hosted zone

Create a Private hosted zone in Services > Route 53. The Domain Name is your preferred domain. For the VPC, use the existing Atlassian Standard Infrastructure.

Step 2: Configure your stack to use the hosted zone

¹ At the time of writing, these regions did not yet support EFS, so also can't be used to run Confluence.

Use your deployments Quick Start template to point your stack to the hosted zone from Step 1. If youre setting up Confluence for the first time, follow the Quick Start template as below:

- 1. Under DNS (Optional), enter the name of your hosted zone in the Route 53 Hosted Zone field.
- 2. Enter your preferred domain sub-domain in the Sub-domain for Hosted Zone field. If you leave it blank, we'll use your stack name as the sub-domain.
- 3. Follow the prompts to deploy the stack.

If you already have an existing Confluence site, you can also configure your stack through the Quick Start template. To access this template:

- 1. Go to to Services > CloudFormation in the AWS console
- 2. Select the stack, and click **Update Stack**.
- 3. Under DNS (Optional), enter the name of your hosted zone in the Route 53 Hosted Zone field.
- 4. Enter your preferred domain sub-domain in the Sub-domain for Hosted Zone field. If you leave it blank, we'll use your stack name as the sub-domain.
- 5. Follow the prompts to update the stack.

In either case, AWS will generate URLs and Route 53 records for the load balancer, EFS, and database. For example, if your hosted zone is my.hostedzone.com and your stack is named mystack, you can access the database through the URL mystack.db.my.hostedzone.com.

Step 3: Link your DNS server to the Confluence sites VPC

If you use a DNS server outside of AWS, then you need to link it to your deployments VPC (in this case, the Atlassian Standard Infrastructure). This means your DNS server should use Route 53 to resolve all queries to the hosted zones preferred domain (in Step 1).

For instructions on how to set this up, see Resolving DNS Queries Between VPCs and Your Network.

If you want to deploy an internal facing Confluence site, using your own DNS server, you can useAmazon Route 53to create a link between the public DNS and internal DNS.

- 1. In Route 53, create a **Private**hosted zone. For the VPC, you can use the existing Atlassian Services VPC. The domain name is your preferred domain.
- 2. If you've already set up Confluence, go to Services > CloudFormation in the AWS console, select the stack, and click Update Stack. (If you're setting up Confluence for the first time, follow the Quick Start template as below).
- 3. UnderOther Parameters, enter the name of your hosted zone in theRoute 53 Hosted Zonefield.
- 4. Enter your preferred sub-domain or leave the Sub-domain for Hosted Zone field blank and we'll use your stack name as the sub-domain.
- 5. Follow the prompts to update the stack. We'll then generate the load balancer and EFS url, and create a record in Route 53 for each.
- 6. In Confluence, go to General Configuration and update the Confluence base URL to your Route 53 domain.
- 7. Set up DNS resolution between your on-premises network and the VPC with the private hosted zone. You can do this with:
 - a. an Active Directory (eitherAmazon Directory Service or Microsoft Active Directory)
 - b. a DNS forwarder on EC2 using bind9 or Unbound.
- 8. Finally, terminate and re-provision each Confluence and Synchrony node to pick up the changes.



For related information on configuring Confluence's base URL, see Configuring the Server Base URL.

Scaling up and down

To increase or decrease the number of Confluence or Synchrony cluster nodes:

- 1. Sign in to the AWS Management Console, use the region selector in the navigation bar to choose the AWS Region for your deployment, and open the AWS CloudFormation console athttps://console.aws. amazon.com/cloudformation/.
- 2. Click the Stack name of your deployment. This will display your deployment's Stack info. From there, click Update.
- 3. On the Select Template page, leave Use current template selected, and then choose Next.
- 4. On the Specify Detailspage, go to the Cluster nodes section of Parameters. From there, set your desired number of application nodesin the following parameters:
 - a. Minimum number of cluster nodes
 - b. Maximum number of cluster nodes
- 5. Click through to update the stack.



Disabled Auto Scaling

Since your cluster has the same minimum and maximum number of nodes, Auto Scalingis effectively disabled.

Setting different values for the minimum and maximum number of cluster nodes enables Auto Scaling. This dynamically scale the size of your cluster based on system load.

However, we recommend that you keep Auto Scaling disabled. At present, Auto Scaling can't effectively address sudden spikes in your deployment's system load. This means that you'll have to manually re-scale your cluster depending on the load.

Vertical VS Horizontal scaling

Adding new cluster nodes, especially automatically in response to load spikes, is a great way to increase capacity of a cluster temporarily. Beyond a certain point, adding very large numbers of cluster nodeswill bring diminishing returns. In general, increasing the size of each node (i.e., "vertical" scaling) will be able to handle a greater sustained capacity than increasing the number of nodes (i.e., "horizontal" scaling), especially if the nodes themselves are small. SeeInfrastructure recommendations for enterprise Confluence instances on AWS for more details.

See the AWS documentation for more information on auto scaling groups.

Connecting to your nodes over SSH

You can perform node-level configuration or maintenance tasks on your deployment through the AWS Systems Manager Sessions Manager. This browser-based terminal lets you access your nodes without any SSH Keys or a Bastion host. For more information, seeGetting started with Session Manager.



Access via Bastion host

You can also access your nodes via a Bastion host (if you deployed one). To do this, you'll need your SSH private key file (the PEM file you specified for the Key Nameparameter). Remember, this key can access all nodes in your deployment, so keep this key in a safe place.

The Bastion host acts as your "jump box" to any instance in your deployment's internal subnets. That is, access the Bastion host first, and from there access any instance in your deployment.

The Bastion host's public IP is the Bastion Publipoutput of your deployment's ATL-Bastion Stacks tack. This stack is nested in your deployment's Atlassian Standard Infrastructure (ASI). To access the Bastion host, useec2-useras the user name, for example:

ssh -i keyfile.pem ec2-user@<BastionPubIp>

Theec2-userhassudoaccess. SSH access is byrootis not allowed.

Upgrading

Consider upgrading to aLong Term Support release(if you're not on one already). Enterprise releases get fixes for critical bugs and security issues throughout its two-year support window. This gives you the option to keep a slower upgrade cadence without sacrificing security or stability. Long Term Support releases are suitable for companies who can't keep up with the frequency at which we ship feature releases.

Here's some useful advice for upgrading your deployment:

- 1. Before upgrading to a later version of Confluence Data Center, check if your apps are compatible with that version. Update your appsif needed. For more information about managing apps, see Using the Universal Plugin Manager.
- 2. If you need to keep Confluence Data Center running during your upgrade, we recommendusing readonly mode for site maintenance. Your users will be able to view pages, but not create or change them.
- 3. We strongly recommend that you perform the upgrade first in a stagingen vironment before upgrading your production instance. Create a staging environment for upgrading Confluence provides helpful tips on doing so.



Rolling upgrades

As of Confluence Data Center 7.9, you can now upgrade to the next bug fix version (for example, 7.9.0 to 7.9.3) with no downtime. Follow the instructions in Upgrade Confluence without downtime.

When the time comes to upgrade your deployment, perform the following steps:

Step 1: Terminate all running Confluence Data Center application nodes

Set the number of application nodes used by the Confluence Data Center stack to 0. Then, update the stack.



If your deployment uses standalone Synchrony, scale the number of Synchrony nodes to 0 at the same time.

- 1. In the AWS console, go to Services > CloudFormation. Select your deployments stack to view its Stack Details.
- 2. In the Stack Details screen, click Update Stack.
- 3. From the Select Templatescreen, select Use current template and click Next.
- 4. Youll need to terminate all running nodes. To do that, set the following parameters to 0:
 - a. Maximum number of cluster nodes
 - b. Minimum number of cluster nodes
- 5. Click Next. Click through the next pages, and then to apply the change using the Update button.
- 6. Once the update is complete, check that all application nodes have been terminated.

Step 2: Update the version used by your Confluence Data Center stack

Set the number of application nodes used by Confluence Data Center to 1. Configure it to use the version you want. Then, update the stack again.



If your deployment uses standalone Synchrony, scale the number of Synchrony nodes to 1 at the same time.

- 1. From your deployments Stack Details screen, click **Update Stack** again.
- 2. From the Select Templatescreen, select Use current template and click Next.

- 3. Set the **Version** parameter to the version youre updating to.
- 4. Configure your stack to use one node. To do that, set the following parameters to 1:
 - a. Maximum number of cluster nodes
 - b. Minimum number of cluster nodes
- 5. ClickNext. Click through the next pages, and then to apply the change using the Updatebutton.

Step 3: Scale up the number of application nodes

You can now scale up your deployment to your original number of application nodes. You can do so for your Synchrony nodes as well, if you have standalone Synchrony. Refer back to Step 1 for instructions on how to re-configure the number of nodes used by your cluster.

Confluence Data Center in AWS currently doesn't allow upgrading an instance without some downtime in between the last cluster node of the old version shutting down and the first cluster node on the new version starting up.Make sure all existing nodes are terminated before launching new nodes on the new version.

Backing up

We recommend you use the AWS native backup facility, which utilizes snap-shots to back up your Confluence Data Center. For more information, seeAWS Backup.

Migrating your existing Confluence site to AWS

After deploying Confluence on AWS, you might want to migrate your old deployment to it. To do so:

- 1. Upgrade your existing site to the version you have deployed to AWS (Confluence 6.1 or later).
- 2. (Optional) If your old database isn't PostgreSQL, you'll need to migrate it. SeeMigrating to Another Databasefor instructions.
- ${\it 3. } \ \ Back\ up\ your\ PostgreSQL\ database\ and\ your\ existing \verb|<shared-home>/attachmentsdirectory.$
- 4. Copy your backup files to /media/atl/confluence/shared-homein your EC2 instance.
- Restore your PostgreSQL database dump to your RDS instance with pg_restore.
 SeeImporting Data into PostgreSQL on Amazon RDSin Amazon documentation for more information on how to do this.

(i) Important notes

- When you create a cluster using the CloudFormation template, the database name isconfluence. Youmust maintain this database namewhen you restore, or there will be problems when new nodes are provisioned. You will need to drop the new database and replace it with your backup.
- You don't need to copy indexes or anything from your existing local home or installation directories, just the attachments from your existing shared home directory.
- If you've modified the <shared-home > /config/cache-settings-overrides. properties file you may want to reapply your changes in your new environment.
- The_copymethod described in this AWS page, Importing Data into PostgreSQL on Amazon RDS, is not suitable for migrating Confluence.

Getting started with Confluence Data Center on Azure

On this page:

- Non-clustered VS clustered environment
- How it works
- Limitations
- Preparing for your deployment
- · Migrating an existing site to Azure
- DeployingConfluenceData Center to Azure via Azure marketplace
- Deploying Confluence Data Center to Azure using the CLI
- Securing your Azure deployment
- Monitoring

If you decide to deploy Confluence Data Center in a clustered environment, consider using Microsoft Azure. This platform allows you to scale your deployment elastically by resizing and quickly launching additional nodes, and provides a number of managed services that work out of the box with Confluence Data Center. These services make it easier to configure, manaAutoscaling is not yet available, due to a problem with Hazelcast, which Confluence uses to discover nodes.ge, and maintain your deployment's clustered infrastructure.

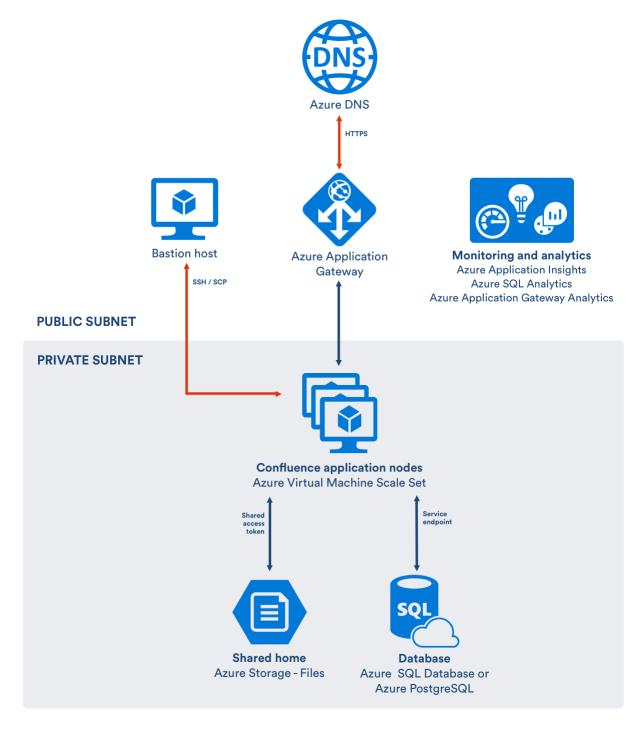
We strongly recommend you set up user management, central logging storage, a backup strategy, and monitoring, just as you would for aConfluenceData Center installation running on your own hardware.

Non-clustered VS clustered environment

A single node is adequate for most Small or Medium size deployments, unless you need specific features that require clustering (for example, high availability). If you have an existing Server installation, you can still use its infrastructure when you upgrade to Data Center. Many features exclusive to Data Center (likeSAML single signon,self-protection via rate limiting,andCDN support) don't require clustered infrastructure. You can start using these Data Center features by simply upgrading your Server installations license. For more information on whether clustering is right for you, check outAtlassian Data Center architecture and infrastructure options.

How it works

Here's an architectural overview of what you'll get when deployingConfluenceData Center using the template:



The deployment contains one or more Azure standard VM instances as cluster nodes in a scale set. Each cluster node runs Confluence Data Center and Synchrony. This way, you don't need to provision extra nodes to enable collaborative editing.

The template also provisions an Azure Files storage account for the shared home. This shared home stores attachments and other files accessible to the application cluster nodes. It's mounted as a SAN drive on each node, and treated normally like any other file.

Standardized infrastructure

The Jira Data Center, Confluence Data Center, Bitbucket Data Center, and Crowd Data Centertemplates deploy the following infrastructure components identically:

Component	Configuration	
-----------	---------------	--

Bastion host	This is a lightweight but highly secure Azure Linux VM that controls SSH access to the application cluster nodes.
Application Gateway	By default, this gateway is composed of two instances for high availability. It acts as a HTTP /HTTPS load balancer for your scale set of application cluster nodes.
Monitoring	The ARM templates configure Azure Monitoring to perform basic health and availability monitoring to cluster nodes and database.
Database	You can choose between Azure SQL Database (MS SQL Server-compatible) or Azure PostgreSQL database. Either way, the database will be configured as service endpoints to only allow traffic from the private network that the cluster nodes are in. This restricted traffic setup helps enhance security.

Limitations

There are some limitations you should be aware of before deciding to deploy to Azure:

- Changing the size of the cluster after creation is not possible, due to a limitation in Hazelcast, which Confluence uses to discover nodes.
- You can't use the deployment template toupgrade an existing Confluence deployment, or to provision new nodes running a different version to the rest of your cluster.
- If a node is deleted manually, it can't be redeployed without first removing the cluster. The existing database, and the existing shared home directory won't be removed when redeploying.

Preparing for your deployment

Before you begin, you should use the Confluence Data Center load profiles to determine the size of your site. This information will help you choose the right infrastructure size during deployment.

You should also decide which Azure region is best for your site. Some services, such as such as Application Insights and Azure SQL Analytics, may not be available in all regions. You can check this athttps://azure.microsoft.com/en-gb/global-infrastructure/regions/.

During the deployment you'll need:

- Your database details, if you want to use an existing Azure database service. You'll need the database URL, port, username, and password.
- A Base64 encoded PFX certificate from a trusted Certificate Authority.
- Details of your existing CNAME, if you don't want Azure to generate a random domain for you.

Migrating an existing site to Azure

To migrate, you will need to set up a new Confluence Data Center site in Azure, and then import content from your old site. This approach ensures that your new site is created with optimum settings for Azure.

Here's a high level overview of the steps:

- 1. Back up your existing site, including your database and home directories.
- 2. Make a list of any Marketplace or other user-installed apps
- 3. Perform a full site export, excluding attachments if you have a large site. You can also turn on read-only mode, to prevent users from making changes in your old site.
- 4. Deploy Confluence Data Center in Azure via the Azure Portal, or CLI, and test that Confluence is working as expected.
- 5. Import your site export file. Make sure you know the administrator password for your existing site, as you'll be logged out during the import.
- 6. Copy the contents of your /attachmentsdirectory to the equivalent directory in your shared home.
- 7. Install any apps.
- 8. Test your site.

At this point you can make the site available to your users, and tear down your old site.

Tips for a successful migration:

- Do a trial run first -export your existing site, and import it into Azure to iron out any issues.
- Because you're setting up your new site in parallel, your current Confluence site can remain accessible throughout the process. If you're already running Confluence Data Center, use read-only mode to prevent people making changes after you've exported the site.
- Unless your existing site is small, exporting the site without attachments will keep the export file smaller.

DeployingConfluenceData Center to Azure via Azure marketplace

This method uses the Azure Marketplace to deploy Confluence Data Center using our deployment templates as a reference.

To deploy Confluence Data Center to Azure using our Marketplace app:

- 1. Log in toAzure Portal.
- 2. Choose Create a resource to start a new deployment
- 3. Search for Atlassian then select Confluence Data Center from the list of Marketplace apps
- 4. Choose Create to start configuring the deployment
- 5. Follow the prompts in the wizard to configure your deployment. Refer to the parameters table below for more information.
- Confirm all the details are correct then click Create to purchase the subscription. Deployment will take about 30 minutes.
- 7. Once deployment is complete, go to the Confluence URL (APPENDPOINT) listed in the deployment outputs to complete onboarding and start using Confluence.

Confluence-specific parameters

Parameters	Description		
Confluence Version	Specify the version of Confluence you'd like to install in full (for example, 6.14.0). Head to Confluence Release Notes for a list of all releases.		
Confluence admin credentials	Provide a name and password for the initial Confluence administrator on your instance.		
Confluence Cluster	Select the expected size of your site - trial, small, medium, large, extra large. This will determine the number of Confluence application nodes, and the size of VMs to be provisioned. Choose Change Size to override the defaults.		

Standardized infrastructure parameters

The Jira Data Center, Confluence Data Center, Bitbucket Data Center, and Crowd Data Centertemplates all share the same parameters:

Parameter	Description
Subscription	Your Microsoft Azure subscription type.
Resource group	If you have an existing resource group, you can use it, or create a new one.
Location	This is the region where Azure will house your deployment.
SSH Access	Provide an SSH public key to be used to SSH into the instance that will act as bastion host, and a username and password for SSH access to the Bitbucket nodes. SeeCreate and use an SSH public-private key pair for Linux VMs in Azurein the Microsoft Azure documentation.

Database configuration	Choose between an Azure SQL Database, or Azure Database for PostgreSQL. Provide a username and password for the database admin user.		
	★ Existing databases If you want to integrate with an existing database, you'll have to deploy to Azure using the CLI.		
CNAME	This is the Canonical Name record (CNAME) for your organization. If you don't provide one, Azure will generate a random sub domain for your instance.		
HTTP/SSL	Provide the certificate and password to be used for SSL termination on the Azure Application Gateway.		
Monitoring	Choose the monitoring and analytics services that you would like to enable. Subject to availability in your location. SeeMonitoringfor related information.		

Deploying Confluence Data Center to Azure using the CLI

This method uses the Azure command line interface to deploy Confluence Data Center using our deployment templates as a reference. You'll need to install the Azure CLI to do this.

Using the deployment templates directly allows for greater configuration granularity. All hardware choices such as the number of cluster nodes, size, disk size, and OS type are configurable as parameters.

Head tohttps://bitbucket.org/atlassian/atlassian-azure-deploymentand check out the README to find out how to to deploy using the CLI.

Required parameters

The deployment template requires a number of values to be provided in order to deploy your Confluence Data Center instance.

Parameter	Description
confCluster Size	To use recommended hardware options for the Confluence installation choose a size. Allowed values: • trial • small • medium • large • enterprise If set, all further Gateway, VM, DB size parameters will be ignored.
clusterSshP assword	This is the SSH password you'll use to access your Confluence nodes.
dbPassword	This the password for your dedicated database user. The password must meet a strong password requirement (imposed by AzureSQL Server): it must be between 16 and 41 characters long, and must contain at least one uppercase letter, one lowercase letter, one number (0-9), and one non-alphanumeric character (.,!, \$, #, %, etc). See the Azure SQL password documentation for details.
confAdmin UserPassw ord	This is the password for your Confluence administrator's account.

Optional parameters

The following parameters are optional. If you don't provide a value in the parameter file, we'll use the default values listed below.

Parameter	Default value	Description
confluence Version	Latest	This is the version of Confluence you want to install on your cluster nodes. Enter the Confluence version number in full, for example "6.14.0".
		We don't recommend using versions prior to 6.12, as they don't support managed Synchrony.
customDow nloadUrl	empty	Use this URL to override standard Atlassian download url, for example to specify beta, release candidate or EAP versions. Used in conjunction with the confluenceVersion parameter.
dbCreateNew	true	Create a new database or attempt to use an existing specified database. Note that this has to be in same resource group and location as the target deployment.
dbType	Azure SQL DB	Choose between Azure SQL Server and Azure DB for PostgreSQL.
dbHost	auto- generated	The hostname of database server to be used if an external database is being used. This will be autogenerated if a new database is to be created.
dbPort	1433	The database port to use if an external database is being used. This will be autogenerated if a new database is to be created.
dbDatabase	confdata base	The database name to use if an external database is being used. This will be autogenerated if a new database is to be created.
dbSchema	auto- generated	The database schema to use if an external database is being used. This will be autogenerated if a new database is to be created.
dbUsername	confluen cedbuser	The username for the dedicated database user.
cname	auto- generated	This is the Canonical Name record (CNAME) for your organization. If you don't provide one, Azure will generate a random domain.
		If you do use a custom domain, you must also update your Domain Registrar's settings to add the Azure DNS Name Servers. Consult your domain registry's documentation on how to configure cname records.
sslBase64E ncodedPfx Certificate		The certificate to be used for SSL termination on the Azure Application Gateway.
sslPfxCertifi catePasswo rd		The certificate password to be used for SSL termination on the Azure Application Gateway.
jumpboxSs hKey		The SSH public key to use to access the bastion host (jumpbox)
confAdmin UserName	admin	The username for the Confluence Administrator's account. Must be lowercase.
confAdmin UserFullNa me	Admin Admin	The full name of the Confluence Administrator's account.

confAdmin UserEmail	admin@ example. com	The email address of the Confluence Administrator user.
confAppTitle	Atlassian Confluen ce	The name of your Confluence site.
jumpboxSs hUser	confluen ceadmin	This is the SSH user you'll use to access the bastion host (jumpbox).
clusterSshU ser	confluen ceadmin	The SSH username to use to access the Confluence nodes from the bastion host (jumpbox). This is the only way you can access Confluence nodes.
enableEmai IAlerts	true	Enable email alerts.
enableAppli cationInsigh ts	true	Enable Azure Application Insights.
enableAnal ytics	true	Enable Azure Operational Insights.

Overriding the recommended hardware options

The confClusterSizeparameter allows you to select the size of your deployment, and then use our recommendations for all resources to be created.

If you choose not to set the confClusterSizeparameter, you can choose to define your own values for things like dbTier, dbTierSize, clusterVmSize, LinuxOsType, and appGtwyTier.

These parameters are all listed in theazuredeploy.jsontemplate file, with a description and allowed values. You should also check out the Developing guide in the template repository to learn more about developing your own template.

Securing your Azure deployment

We recommend deploying Confluence with SSL. Our template will prompt you for a certificate and password.

Good to know:

- HTTPS is terminated at the application gateway.
- Your certificate should be from a trusted Certificate Authority. You should avoid self-signed certificates.

Monitoring

As a number of the resources we provision are managed by Azure, a number of options are available for monitoring. For example:

- A number of default alerts are available, such as cluster nodes going offline, CPU, or Db DTU exceeding 80%. These alerts will be emailed to the Confluence Administrator email address specified in the deployment.
- Application Insights can be used to see the overall system health, and dig into particular areas of interest Application Insights in the Azure documentation.
- Azure SQL Analytics is available for more granular monitoring of your SQL Server database. Monitor Azure SQL Database using Azure SQL Analytics in the Microsoft Azure documentation.

Note that some of these resources are still in Preview, so may not be available in your location yet.

Administering Confluence Data Center on Azure

Once you've deployed Confluence Data Center to Azure using the deployment template, administering the application is similar to managing an application on your own hardware, with the exception that you'll need to go via the bastion host (jumpbox) to access your nodes.

To access your jumpbox and nodes you'll need:

- the SSH credentials you provided during setup,
- the Confluence node credentials you provided during setup
- the public DNS name or IP address of your jumpbox (you can obtain this through the Azure portal viaMen u>Resource groups><your resource group>>confluencenat), and
- the node IP addresses, listed against theconfluencecluster (instance n)row inConnected devices. (You can obtain this through the Azure portal viaMenu>Resource groups><your resource group>>confluencevnet).

Connecting to your Azure jumpbox over SSH

You can SSH into yourConfluencecluster nodes, Synchrony nodes and shared home directory to perform configuration or maintenance tasks. Note that you must keep your SSH public key filein a safe place. This is the key to your jumpbox, and therefore all the nodes in your instance.

Access the jumpbox via a terminal or command line using:

```
$ ssh JUMPBOX_USERNAME@DNS_NAME_OR_IP_ADDRESS
```

You can find the SSH URL in the outputs section of your deployment.

Once you've accessed the jumpbox, we can jump to any of the nodes in the cluster, using:

```
$ ssh NODE_USERNAME@NODE_IP_ADDRESS
```

You'll then be asked for your node password - after providing this, you should be connected to the node.

Accessing your configuration files

For your Azure deployment, you may need to make changes to some configuration files, just as you would for a deployment on your own hardware:

- yourserver.xmllives in/opt/atlassian/confluence/conf
- yoursetenv.shlivesin/opt/atlassian/confluence/bin
- your local homeconfluence.cfg.xmllives in/var/atlassian/application-data/confluence
- your shared homeconfluence.cfg.xmllives in/media/atl/confluence/shared

These files are only accessible from the existing nodes. The shared home is mounted (think of it as a network hard disk) on each node under/media/atl/confluence/shared. So from an existing node (when you're logged in through SSH), you can go to/media/atl/confluence/shared.

If modifications to these files are made manually, new nodes will not pick up those modifications. You can either repeat the modifications on each node, or change the templates in the/media/atl/confluence/shareddire ctory from which those files are derived. The mappings are:

- theserver.xmlfile is derived from/media/atl/confluence/shared/server.xml
- thesetenv.shfile is derived from/media/atl/confluence/shared/setenv.sh
- the local homeconfluence.cfg.xmlis derived from/media/atl/confluence/shared/homeconfluence.cfg.xml
- the shared homeconfluence.cfg.xmlis derived from/media/atl/confluence/shared/shared-confluence.cfg.xml

These template files contain placeholders for values that are injected via the deployment script. Removing or changing them may cause breakages with the deployment. In most cases, these files should not be modified, as a lot of these settings are produced from the Azure Resource Manager templates automatically.

Upgrading

Consider upgrading to aLong Term Support release(if you're not on one already). Enterprise releases get fixes for critical bugs and security issues throughout its two-year support window. This gives you the option to keep a slower upgrade cadence without sacrificing security or stability. Long Term Support releases are suitable for companies who can't keep up with the frequency at which we ship feature releases.

Here's some useful advice for upgrading your deployment:

- 1. Before upgrading to a later version of Confluence Data Center, check if your apps are compatible with that version. Update your appsif needed. For more information about managing apps, see Using the Universal Plugin Manager.
- 2. If you need to keep Confluence Data Center running during your upgrade, we recommendusing read-only mode for site maintenance. Your users will be able to view pages, but not create or change them.
- 3. We strongly recommend that you perform the upgrade first in a stagingenvironment before upgrading your production instance. Create a staging environment for upgrading Confluence provides helpful tips on doing



Rolling upgrades

As of Confluence Data Center 7.9, you can now upgrade to the next bug fix version (for example, 7.9.0 to 7.9.3) with no downtime. Follow the instructions in Upgrade Confluence without downtime.

Upgrading Confluence in Azure

The process of upgrading Confluence is the same as if you were running the cluster on your own hardware. You will stop Confluence on all nodes, upgrade one node, stop that node then copy the installation directory across to each remaining node in the cluster, before restarting each node, one at a time.

SeeUpgrading Confluence Data Centerfor more details.



You can't use the confluenceVersion parameter in the deployment template to upgrade an existing Confluence deployment, or to provision new nodes running a different version to the rest of your cluster.

You also can't do a rolling upgrade. You will need to bring all nodes down before upgrading.

Upgrading your operating system

If you need to upgrade the operating system running on your Confluence nodes, you will need to SSH into each node, perform asudo apt dist-upgrade(Ubuntu) and reboot each node.

As Confluence is running as a service it will be automatically restarted on reboot.

You can't simply reimage an instance, as you might do in Jira, due to the way Hazelcast discovers cluster nodes.

Backing up and recovering from failures

We recommend you use the Azure native backup facilities where possible to make sure your data is backed up, and you can easily recover in the case of a failure.

Database backups

We use Azure-managed database instances with high availability. Azure provides several excellent options for backing up your database, so you should take some time to work out which will be the best, and most cost effective option for your needs. See the following Azure documentation for your chosen database:

SQL Database: Automated backups
 SQL Database: Backup retention
 PostGreSQL: Backup concepts

Shared home backups

The shared home stores your your attachments, profile pictures, and export files. We create a general purpose Azure storage account, configured with local redundant storage (LRS), which means there are multiple copies of the data at any one time.

LRS provides a basic redundancy strategy for your shared home. As such, it shouldn't be necessary to take regular backups yourself. If you need to take point-in-time backups, usesnapshots.

Application nodes

The application nodes are VMs in an Azure Virtual Machine Scale Set. Each application node has a Confluence installation directory and a local home directory containing things like logs and search indexes.

Like the shared home, application nodes are configured with local redundant storage. This means there are multiple copies of the data at any one time.

If you've manually customised any configuration files in the installation directory (for example velocity templates), you may also want to manually back these up as a reference.

Bastion host

As this VM acts as a jumpbox, and doesn't store any data it doesn't need to be backed up. If the VM becomes unresponsive it can be restarted from the Azure Portal.

Application gateway

The application gateway is highly available. We deploy 2 instances by default. As with the bastion host, it doesn't need to be backed up.

Disaster recovery

SeeConfluence Data Center disaster recovery to learn about how you can develop a disaster recovery strategy. See also information in the Azure documentation about recovering from a region-wide failureAzure resiliency technical guidance: recovery from a region-wide service disruption.

Running Confluence Data Center on a Kubernetes cluster

If you're running self-managed environments and looking to adopt modern infrastructures, Atlassian Data Center products can now be deployed on Kubernetes clusters. By leveraging Kubernetes, you can drive greater agility amongst your teams while experiencing a simplified administrative experience at scale, without compromising your organizations regulatory requirements.

We offer Helm charts on GitHubfor installing and operating Atlassian products on a Kubernetes cluster of your choice. The Helm charts we offer are for:

- Jira
- Confluence
- Bitbucket
- Bamboo

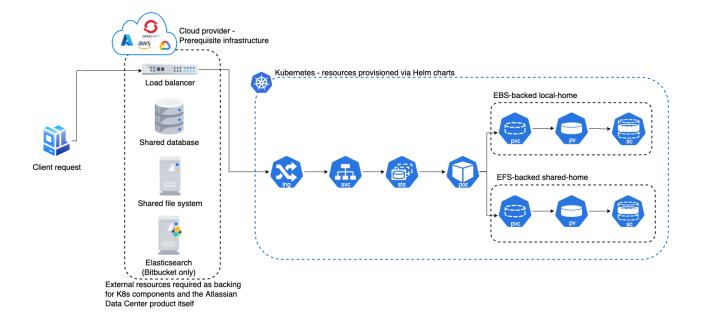
Helmis a package manager forKubernetesthat allows you to package, configure, and deployapplications and services ontoKubernetesclusters. Helm uses a packaging format called *charts*, which are collections of files that describe a related set of Kubernetes resources.

The Kubernetes cluster can be a managed environment, such as Amazon EKS, Azure Kubernetes Service, Google Kubernetes Engine, or a custom on-premise system.

We strongly recommend you set up user management, central logging storage, a backup strategy, and monitoring, just as you would for a Data Center installation running on your own hardware.

How it works

Here's an architectural overview of what you'll get when deployingyour Data Center application on a Kubernetes cluster using the Helm charts:



Kubernetes entities required for product deployment:

- Ingress and Ingress controller (ing) the Ingress defines the rules for traffic routing, which indicate where a request will go in the Kubernetes cluster. The Ingress controller is the component responsible for fulfilling those rules.
- Service (svc) provides a single address for a set of pods to enable load-balancing between application nodes
- StatefulSets (sts) manages the deployment and scaling of a set ofpods.

- Pod a group of one or more containers, with shared storage and network resources, and a specification
 for how to run the containers. Podsare the smallest deployable units of computing that you can create
 and manage in Kubernetes.
- PersistentVolumeClaim (pvc) reserves the Persistent Volume (PV) to be used by a pod.
- PersistentVolume (pv) is the "physical"volumeon the host machine that stores your persistent data.
- StorageClass (sc) provides a way for administrators to describe the "classes" of storage they offer.

Installingyour Data Center application on a Kubernetes cluster

In addition to the Helm charts we provide for installing the Atlassian Data Center applications, we also providedo cumentation and examples on a dedicated documentation site.

Before you begin, you need to have an understanding of Kubernetes and Helm concepts.

To install and operate your Data Center application on a Kubernetes cluster using our Helm charts:

- 1. Follow the requirements and set up your environment according to the Prerequisites guide.
- 2. Perform the installation steps described in the Installation guide.
- 3. Learn how to upgrade applications, scale your cluster, and update resources using the Operation guide.

Installing Java for Confluence

This page contains instructions for installing the Java Development Kit (JDK). This is a manual step that's only required if you're installing Confluence from a zip or archive file.

If you're using the Confluenceinstaller, you don't need to install Java manually, but you can choose to use a different Java vendor.

Check the Supported Platforms page to find out which Java versions and vendors can be used with Confluence.

Installing Java

The JDK (Java Development Kit) needs to be installed on the same server that will have Confluence installed. We support running Confluence with the JDK or JRE (Java Runtime Environment). These instructions will just cover installing the JDK.

Before you start, go to **Control Panel** > **Programs and Features** to check whether a JDK is already installed.

To install the JDK on Windows:

- Download the appropriate Adoptium OpenJDK or Oracle JDK version.
 Check the Supported Platforms page to find out which JDK / JRE versions and vendors are supported for your version of Confluence. Be sure to download the right one for your operating system.
- 2. Run the Java installer. Make a note of the installation directory, as you'll need this later.
- 3. Once the Java installation is complete, check that the <code>JAVA_HOME</code> environment variable has been set correctly.

Open a command prompt and type echo %JAVA_HOME% and hit Enter.

- If you see a path to your Java installation directory, the JAVA_Home environment variable has been set correctly.
- If nothing is displayed, or only %JAVA_HOME% is returned, you'll need to set the JAVA_HOME environment variable manually. See Setting the JAVA_HOME Variable in Windows for a step by step guide.

Before you start, check whether a JDK is already installed. Open a shell console and type echo \$JAVA_HOME and hit Enter.

- If it returns something like/opt/JDK8 or /usr/lib/jvm/java8, then your JDK is installed and properly configured.
- If nothing is displayed, you'll 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.

To install the JDK on Linux:

- Download the appropriate Adoptium OpenJDK or Oracle JDK version.
 Check theSupported Platforms page to find out which JDK / JRE versions are supported for your version of Confluence. Be sure to download the right one for your operating system.
- 2. Run the Java installer.
- 3. Open a shell console and typeecho \$JAVA_HOME and hit Enter to check that it has installed correctly (see notes above).

Note: Any Java or JDK version numbers on this page are **examples only**. Please refer to the **Supported Platforms** page for supported versions of Java.

Setting the JAVA_HOME Variable in Windows

To install Confluence manually on Windows, you will need to set an environment variable to point Confluence to the your Java installation directory.



This information is only relevant if you're installing Confluence manually on a Windows server. If you're using the installer, you don't need to do this.

Related pages

- Starting Tomcat as a Windows Service
- Installing Confluence in Linux

In most cases you should set the JRE HOME environment variable, but if it is not set, Confluence will use JAVA_HOME.

Set the JAVA HOME Variable

To set the JRE HOME or JAVA HOME variable:

1. Locate your Java installation directory



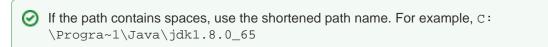
If you didn't change the path during installation, it'll be something like C:\Program Files\Java\jdk1.8.0 65

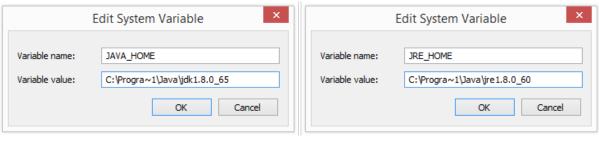
You can also type where java at the command prompt.

2. Do one of the following:

Windows 7 Right click My Computer and select Properties > Advanced Windows 8 Go to Control Panel > System > Advanced System Settings Windows 10 Search for Environment Variables then select Edit the system environment variables

- 3. Click the **Environment Variables** button.
- 4. Under System Variables, click New.
- 5. In the **Variable Name** field, enter either:
 - JAVA_HOME if you installed the JDK (Java Development Kit)
 - JRE_HOME if you installed the JRE (Java Runtime Environment)
- 6. In the Variable Value field, enter your JDK or JRE installation path.





Note for Windows users on 64-bit systems Progra~1 = 'Program Files' Progra~2 = 'Program Files(x86)'

7. Click OK and Apply Changes as prompted

You'll need to close and re-open any command windows that were open before you made these changes, as there's no way to reload environment variables from an active command prompt. If the changes don't take effect after reopening the command window, restart Windows.

Set the JAVA_HOME variable via the command line

If you would prefer to set the JAVA_HOME (or JRE_HOME) variable via the command line:

- Open Command Prompt (make sure you Run as administrator so you're able to add a system environment variable).
- 2. Set the value of the environment variable to your JDK (or JRE) installation path as follows:

```
setx -m JAVA_HOME "C:\Progra~1\Java\jdk1.8.0_XX"
```

If the path contains spaces, use the shortened path name.

3. Restart Command Prompt to reload the environment variables then use the following command to check the it's been added correctly.

```
echo %JAVA_HOME%
```

You should see the path to your JDK (or JRE) installation.

Change the Java vendor or version Confluence uses

When you install Confluence Server using the installer, it will run Confluence with the Java Runtime Engine (JRE) that was bundled with that Confluence release.

If you want to use a different Java vendor, version, or you want to install the full JDK, you can tell Confluence to use the version of Java installed on your server.

Not all vendors and versions are supported, and some versions have known issues, so always check the Supported Platforms page, as using an unsupported version can cause problems in Confluence.

On this page:

- Check your current setup
- Installer method Windows
- Installer method Linux
- Environment variable method Windows and Linux
- How Confluence determines which Java to use
- Which Java vendor can I use with my Confluence version?
- Known issues
- Upgrading Java

Check your current setup

How you change Confluence's Java path depends on whether you originally installed Confluence using the installer, or manually from a .zip or .tar.gz file.

The easiest way to check how Confluence is currently finding your Java is to:

- 1. Go to <install-directory>/bin/setjre.shfile (Linux) or setjre.bat(Windows) file.
- 2. Scroll to the bottom of the file and look for a line similar to the following. The file path may be different in your file.

In Linux:

```
JRE_HOME="/opt/atlassian/confluence/jre/"; export JRE_HOME
```

In Windows:

```
SET "JRE_HOME=C:\Program Files\Atlassian\Confluence\jre"
```

If a line similar to the one above is present, then JRE_HOME **is set** in this file by the installer, and you should use the **installer method**for Windows or Linux below.

If this line isn't present, JRE_HOME **is not set** in this file (because Confluence was installed manually), and you should use the **environment variable** method below.

Installer method - Windows

The way you do this depends on whether you run Confluence manually using the start-confluence.batfile, or as a Windows service.

In these examples we're going to point Confluence to the AdoptOpenJDK JRE, which is installed on our Windows server atC:\Program Files\AdoptOpenJDK\jdk8u192-b12\jre. The location of your JRE will be different, but the steps are the same for any supported Java vendor and version.

If you start Confluence manually

To change the Java that Confluence uses if you start Confluence manually in Windows:

1. In Command Prompt, use the following command tocheck that Java is installed and has been added to your path correctly.

```
> java -version
```

This will return your Java version. If nothing is returned, or it returns the wrong version, check the installation instructions for your Java vendor.

- 2. Stop Confluence.
- 3. In the Confluence installation directory edit the <install-directory > /bin/setjre.batfile and change the last line to your local Java installation, as in the example below.

```
SET "JRE_HOME=C:\Progra~1\AdoptOpenJDK\jdk8u192-b12\jre"
```

If this line isn't present, exit this file and use the environment variable method below.

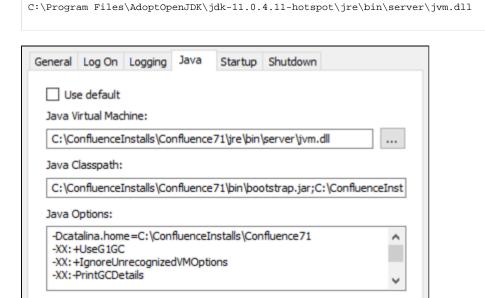
- 4. Start Confluence.
- 5. Go to Seneral Configuration System Information and check that Confluence is using the expected Java version.

Remember, when you next upgrade Confluence this file will be overwritten, so you will need to re-apply this change to the newsetjre.batfile.

If you run Confluence as a Windows service

To change the Java that Confluence uses if you run Confluence as a Windows service:

- 1. Open the Tomcat properties dialog. SeeHow to set system properties for Confluence running as a service on Windowsfor a step-by-step guide to locating your service and launching the Tomcat dialog.
- 2. Choose the Javatab.
- 3. Update the Java Virtual Machineline to point to the Adopt Open JDK jvm.dll, as in the example below. The path to your Java installation will be different to our example.



- 4. Restart the Confluence Windows Service.
- 5. Go to Seneral Configuration System Information and check that Confluence is using the expected Java version.

Remember, when you next upgrade Confluence this file will be overwritten, so you will need to re-apply this change to the service.

Installer method - Linux

In this example we're going to point Confluence to the AdoptOpenJDK JRE, which is installed on our Linus server at /opt/java/adoptopenjdk/jdk-11.0.4.11-hotspot/. The location of your JRE will be different, but the steps are the same for any supported Java vendor and version.

To change the Java that Confluence uses in Linux:

1. In Terminal, use the following command to check that Java is installed and added to your path correctly.

```
$ java -version
```

This will return your Java version. If nothing is returned, or it returns the wrong version, seeInstalling Java for Confluenceor check the installation instructions for your Java vendor.

- 2. Stop Confluence.
- 3. In the Confluence installation directory edit the <install-directory>/bin/setjre.shfile andchange the last lineto point to your local Java installation, as in the example below.

The path to your Java installation will be different to our example.

```
JRE_HOME="/opt/java/adoptopenjdk/jdk-11.0.4.11-hotspot/"; export JRE_HOME
```

If this line isn't present, exit this file and use the environment variable methodbelow.

- 4. Start Confluence.
- 5. Go to General Configuration>System Information and check that Confluence is using the expected Java version.

Remember, when you next upgrade Confluence this file will be overwritten, so you will need to re-apply this change to the new setjre.shfile.

Environment variable method - Windows and Linux

If you installed Confluence manually (the path to the bundled JRE was not automatically set in the setjrefile), Confluence will use the path set in the JRE_HOME environment variable. If JRE_HOME is not set, it will use the path set in JAVA_HOME.

See Setting JAVA_HOME variable for Confluence to find out how to set this environment variable in Windows.

Refer to the documentation for your Linux distribution to find out how to set an environment variable in Linux.

You won't need to update the JRE_HOME environment variable when you upgrade Confluence, but you will need to update the path if you upgrade Java.

How Confluence determines which Java to use

The JRE_HOME set in the set jrefile takes precedence. If you installed Confluence using the installer, this will be automatically set to the Java version bundled with Confluence.

If JRE_HOME is not set in thesetjre.batorsetjre.shfile, Confluence will use the JRE_HOME defined in your environment or service. If it can't find JRE HOME, it will use the JAVA HOME environment variable.

Which Java vendor can I use with my Confluence version?

The following table lists the supported Java vendors, and whether Oracle or AdoptOpenJDK is bundled with Confluence.

Confluence version	Supported Java vendors	Bundled Java vendor
6.6.12 and earlier	Oracle JRE	Oracle JRE
6.7.0 to 6.13.1, and 6.14.0	Oracle JRE	Oracle JRE
6.13.2 to 6.13.x, and 6.14.1 and later	Oracle JDK/JRE Adoptium OpenJDK (formerly known as AdoptOpenJDK)	AdoptOpenJDK

Known issues

- You may find that Oracle is still listed as the vendor in System Information. This is a known issue in Confluence which we hope to have resolved soon. The Java version will be reported correctly, so you can use that to make sure Confluence is pointing to the right version.
- AdoptOpenJDK does not include a required font configuration package, which may cause issues when
 installing in Linux. SeeConfluence Server 6.13 or later fails with FontConfiguration error when installing
 on Linux operating systemsfor information on how to install the required package manually.
- AdoptOpenJDK is now known as Adoptium OpenJDK.

Upgrading Java

If you choose not to use the bundled Java version, you will need to manually update Java from time to time, to get access to new security fixes and other improvements.

Always check the Supported Platforms page before upgrading, for any known issues affecting particular Java versions.

If upgrading to a major version, for example from Java 8 to Java 11, be aware that some Java arguments will not be recognised in later versions. When you upgrade, make sure you apply your customisations manually, don't simply copy over your old setenv.sh/setenv.batfile, or existing Java options if you run Confluence as a service.

Creating a Dedicated User Account on the Operating System to Run Confluence

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.

This is optional if you're evaluating Confluence, but is required for production installations. If you used the Confluence installer on Linux, the installer created this user automatically.

Create a dedicated user account

Linux

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
```

Windows

If your operating system is Windows create the dedicated user account bytyping 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.)

Alternatively, openthe Windows 'Computer Management' console to add your 'confluence' user with its own password.

Next, 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).

Allow the account to write to specific Confluence directories

Ensure that the following directories can be read and written to by this dedicated user account (e.g. 'confluence'):

- The sub-directories of theConfluence Installation Directory:
 - logs
 - temp
 - work
- The entireConfluence Home directory.

Set who can access Confluence directories in Linux

To achieve this in Linux run the following commands:

```
sudo chown -R confluence <confluence-home-folder>/
sudo chown -R confluence <confluence-install-folder>/logs
sudo chown -R confluence <confluence-install-folder>/work
sudo chown -R confluence <confluence-install-folder>/temp
```

The other install directories should be left as root as those are controlled by the installer and allow for future upgrades:

```
sudo chmod -R u=rwx,g=rx,o=rx <confluence-install-folder>
sudo chmod -R u=rwx,g=rx,o=rx <confluence-home-folder>
```

See alsoBest Practices for Configuring Confluence Security.

Confirm who can access Confluence directories in Windows

After installing Confluence youshould check the permissions assigned to the installation directory, and make sure there are no unnecessary permissions being inherited. You can also repeat this process for the home directory.

To check the permissions for the install directory:

- 1. Right click your installation directory and select **Properties**.
- 2. In the Security tab, select Advanced.
- 3. Select **Disable inheritance**, and when prompted choose **Convert inherited permissions into explicit** permissions on this object.
- 4. Select OK.
- Select any group or user account that should not have access and choose Remove.
 We recommend limiting access to only the dedicated 'confluence' user and system administrator groups.
- 6. Select **OK** to apply changes to your install directory (and all sub-directories).

To confirm your changes, log in to Windows with a normal user account, and check that you can't access the contents of the install directory.

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

- Start Confluence (if it is not already running)
 For Windows, go toStart > 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 Linuxbased systems.
- 2. Go http://localhost:8090/in your browser If you chose a different port during installation, change '8090' to the port you specified If you see an error, check you are using the port 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 your Jira application
- 10. Set up system administrator account
- 11. Setup is Complete

2. Choose your installation type

In this step, you'll choose whether you want a trial or a production installation.

Trial installation

Choose this option if you don't have a license, and want to try Confluence for the first time. You'll need an external database.

Production installation

Set up Confluence with your own external database. This option is recommended for setting up Confluence in a production environment.

3.Enter your license key

Follow the prompts to generate an evaluation license, or enter an existing license key. To retrieve an existing license key head tomy. atlassian.com, or to purchase a new commercial license go towww.atlassian.com/buy.

If you selected a **Trial installation** in the previous step, Confluence will generate your license. This will take a few minutes. Once complete, go to step 8 below.

If you selected a **Production installation**, go to the next step to set up your external database.

4. Production installation: database configuration

Next it's time to set up your database. Some things to consider:

- Check thesupported platformslist to confirm that your chosen database and version is supported.
- Seedatabase configuration for information on setting up your database, including UTF-8 character encoding requirements.

- If you are using Confluence as a production system youmustuse an external database.
- The embedded H2 database isonlysupported for testing and app development purposes onnonclustered (single node)Confluence Data Center installations.

Screenshot: Database configuration

Choose a Database Configuration Choose where Confluence should store its data. 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 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

5. Production installation: external database

and reliability. A guide on how to migrate your data can be found here.

B

Before you Start

Embedded Database

- 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'.

Choose how you want Confluence to connect to your database eithervia a direct JDBC connection or via a server-managed datasource connection.

Screenshot: Connection options

Configure Database

Choose how you wish Confluence to connect to your database.

Direct JDBC Connection

Connect directly to the database

Direct JDBC

Connect via a Datasource

Connect to the database via a datasource configured in the application server

Datasource

Direct JDBC

This uses a standard JDBC database connection. Connection pooling is handled within Confluence.

- 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. SeeDatabase JDBC Driversfor 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 and Password A valid username and password that Confluence can use to access your database.

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.

Datasource

This asks your application server (Tomcat) for a database connection. You will need to have configured a datasource in your application server. For information about configuring an external database, seeDatabase Configuration.

• 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 like java: comp/env/jdbc/datasourcename. Check 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

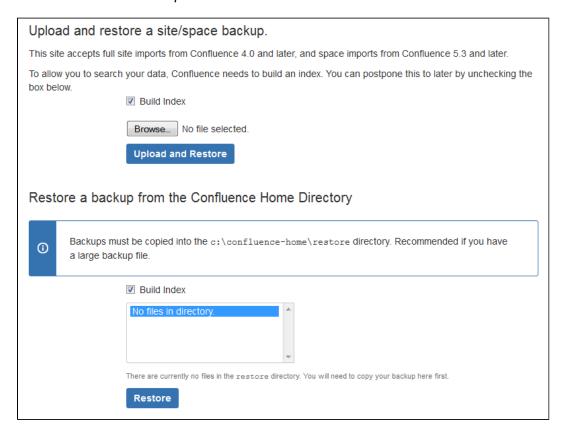
We can help you get your new Confluence site started with some demonstration content (which you can remove once you're up and running), or you can choose to proceed with an empty site. You'll need to create a space in your new site before you can start adding content.

If you're migrating from another Confluence installation choose **Restore from backup**to import your existing Confluence data.

7. Production Installation: restore data from backup

This option allows you to import data from an existing Confluence installation as part of the setup process. You'll need amanual backupfile from your existing Confluence installation to do this (go to **Backup and Restore** in the administration console of your existing Confluence site).

Screenshot: restore data options



There are two ways to restore your data - upload the file, or restore from a location on your file system.

Upload a backup file

This option will load the data from a zipped backup file. If your backup file is very large, restoring from the file system is a better option. Follow the prompts to browse for your backup file. Ensure select **Build Index** is selected so the search index is generated.

Restore a backup file from the file system

This option is recommended if your backup file is very large (100mb or more), or your backup file is already on the same server.

Copy your XML backup file into the <confluence-home>/restoredirectory. Yourbackup file will appear in the list. Follow the prompts to restore the backup. Ensure select**Build Index**is selected so the search index is generated.

When the restore process has you'll be ready to log in to Confluence. The system administrator account and all other user data and content has been imported from your previous installation.

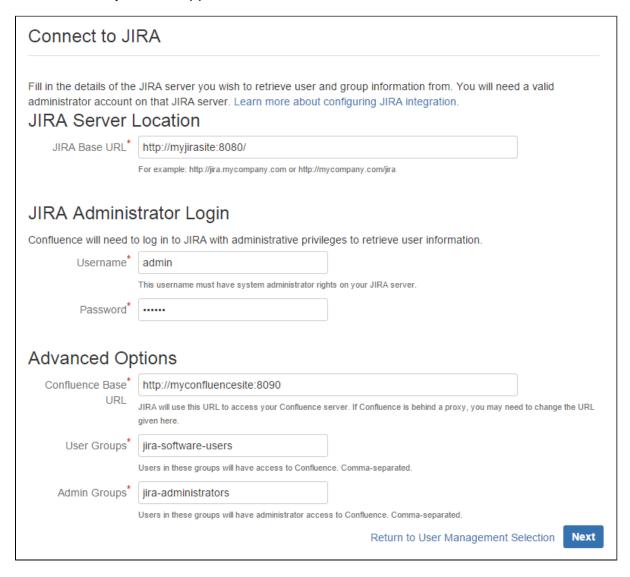
8. Set up user management

You can choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Management.

- If you do not have a Jira application installed, or if you would prefer to set up external user management later, chooseManage users and groups within Confluence.
- If you have a Jira application 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 your Jira application



Enter the following information:

- Jira Base URL the address of your Jira server, such as http://www.example.com:8080/jira/ or http://jira.example.com
- Jira Administrator Login- this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application.

Confluence will also use this username and password to create a local administrator account which will let you access Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.

- Confluence Base URL this is the URL Jira will use to access your Confluence server. The URL you
 give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira
 application.
- User Groups these are the Jira groups whose members should be allowed to use Confluence.
 Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
 - Jira 6.4 and earlier: jira-users.
 - Jira Software 7.x and later: jira-software-users
 - Jira Core 7.x and later: jira-core-users

- Jira Service Management (formerly Jira Service Desk) 3.x and later: jira-servicedesk-
- 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 get the system administrator and Confluence administrator global permissions in Confluence.

For full details and a troubleshooting guide, see Configuring Jira Integration in the Setup Wizard.

10. Set up system administrator account

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.

If you've delegated user management to a Jira application, we'll use the Jira system administrator account you specified as Confluence's system administrator account.

11. Setup is Complete

That's it, Confluence is ready to go. Click Startto jump straight in to Confluence.

Choose 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.

Configuring Jira Integration in the Setup Wizard

This page describes the **Connect to Jira** step in the Confluence setup wizard.

If you are already using a Jira application, you can choose to delegate user management to Jira, instead of separately maintaining your users in Confluence.

You'll be able to specify exactly which groups in your Jira app should also be allowed to log in to Confluence. Your license tiers do not need to be the same for each application.

It's possible to connect Confluence to Jira after completing the setup process, but it's much quicker and easier to set it up at this stage.

You can delegate Confluence's user management to:

- Jira 4.3 or later
- Jira Core 7.0 or later
- Jira Software 7.0 or later
- Jira Service Management (formerly Jira Service Desk) 3.0 or later.

On this page:

- Connecting to a Jira application in the Setup Wizard
- Troubleshooting

Related pages:

- User Management Limitations and Recommendations
- Connecting to Crowd or Jira for User Management
- Confluence Setup Guide

Connecting to a Jira application in the Setup Wizard

Connect to JIRA			
	e JIRA server you wish to retrieve user and group information from. You will need a valid on that JIRA server. Learn more about configuring JIRA integration. Location		
JIRA Base URL*	http://myjirasite:8080/		
	For example: http://jira.mycompany.com or http://mycompany.com/jira		
JIRA Adminis	strator Login		
Confluence will need t	to log in to JIRA with administrative privileges to retrieve user information.		
Username*	admin		
	This username must have system administrator rights on your JIRA server.		
Password*			
Advanced Options			
Confluence Base*	http://myconfluencesite:8090		
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*	jira-software-users		
	Users in these groups will have access to Confluence. Comma-separated.		
Admin Groups*	jira-administrators		
	Users in these groups will have administrator access to Confluence. Comma-separated.		
	Return to User Management Selection Next		

Enter the following information:

- Jira Base URL the address of your Jira server, such as http://www.example.com:8080/jira/ or http://jira.example.com
- **Jira Administrator Login** this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application.

Confluence will also use this username and password to create a local administrator account which will let you access Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.

- Confluence Base URL this is the URL Jira will use to access your Confluence server. The URL you
 give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira
 application.
- **User Groups** these are the Jira groups whose members should be allowed to use Confluence. Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
 - Jira 6.4 and earlier: jira-users.
 - Jira Software 7.x and later: jira-software-users
 - Jira Core 7.x and later: jira-core-users
 - Jira Service Management (formerly Jira Service Desk) 3.x and later: jira-servicedesk-users
- 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 get the system administrator and Confluence administrator global permissions in Confluence.

Troubleshooting

If you have trouble connecting Confluence to Jira, the following troubleshooting information should help you get up and running.

If no users can log in to Confluence after you've completed the setup process, check that the people are members of the Jira groups you specified. Only members of these groups will get the 'Can Use' Confluence permission.

Error in the setup wizard	Cause	Solution
Failed to create application link, or Failed to authenticate application link	The setup wizard failed to complete registration of the peer-to-peer application link with Jira. Jira integration is only partially configured.	Follow the steps below to remove the partial configuration then try the Connect to Jira step again.
Failed to register Confluence configuration in Jira for shared user management	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.	Follow the steps below to remove the partial configuration then try the Connect to Jira step again.
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 yourconfig.xmlfile. This may be caused by a problem in your environment, such as a full disk.	Fix the problem that prevented the application from saving the configuration file to disk then follow the steps below to remove the partial configuration before trying the Connect to Jira step again.
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 Confluence. You should be able to continue with the setup wizard. If this does not work, contact Atlassian Support for help.
java.lang. IllegalStateExce ption: Could not create the application in Jira/Crowd (code: 500)	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.</application>	Follow the steps below to remove the partial configurationand resolve any conflict with existing links then try the Connect to Jira step again.

Removing a partial configuration

If you hit a roadblock, you'll need to log in to Jira and remove the partial integration before you can try again. The specific steps will differ depending on your Jira application and version, but the essentials are the same for all versions:

- Log in to Jira as a user with system administrator permissions.
- In the Administrator screens, go to **Application Links**.
- Remove the application link that matches the base URL of your Confluence server.
- In the User Management screens, go to Jira User Server.
- Remove the link that matches the name and base URL of your Confluence server from the list of applications that can use Jira for user management.

If you're unable to tell which link matches your Confluence server because you have multiple servers of the same type running on the same host you can check the application ID, which is listed beside each server.

To find out the application ID of your new Confluence site, go to <code><baseUrl>/rest/applinks /1.0/manifest</code> (where <code><baseurl></code> is the base URL of your Confluence site). The application ID will be listed in the <code><ID></code> element.

• Return to the Confluence setup wizard and try the **Connect to Jira** step again.

If you're still unable to connect Jira and Confluence using the setup wizard, you may need to skip this step and set up the links between Jira and Confluence manually once you've completed the Confluence setup process. SeeConnecting to Crowd or Jira for User Management.

Upgrading Confluence

In this guide we'll run you through using the installer to upgrade your Confluence site to the latest Confluence version on Windows or Linux.

Upgrading to any later version is free if you have current software maintenance. See ourLicensing FAQto find out more.



Other ways to upgrade Confluence:

- Manuallyupgrade Server or single-node Data Center without using the installer.
- Cluster with downtime upgrade your Data Center cluster.
- Cluster without downtime rolling upgrade to a compatible bug fix version, with no downtime.

XML backups should**not**be used to upgrade Confluence.

Before you begin

Before you upgrade Confluence, there's a few questions you need to answer.

Which upgrade method is the best option?

You can choose to upgrade using the installer, or manually using a zip or tar.gz file. In most cases the installer is the easiest way to upgrade your Confluence instance.

You will need to upgrade manually if you are:

- moving to another operating system or file location as part of this upgrade.
- upgrading from Confluence 3.5 or earlier
- upgrading from Confluence 5.6 or earlier and previously used the EAR/WAR distribution to deploy Confluence into an existing application server.
- performing a rolling upgrade, and you need to upgrade each node individually.

On this page:

Before you begin Plan your upgrade

- 1. Determine your upgrade path
- 2. Complete the pre-upgrade checks
- 3. Upgrade Confluence in a test environment

Upgrade Confluence

- 4. Back up
- 5. Download Confluence
- 6. Run the installer

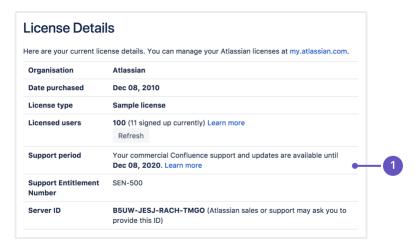
After the upgrade

- 7. Copy your database driver
- 8. Reinstall the service if required (Windows only)
- 9. Re-apply any modifications
- 10. Update your apps (add-ons)
- 11. Update your reverse proxy and check you can access Confluence

Troubleshooting

Are you eligible to upgrade?

To check if software maintenance is current for your license, go to Seneral Configuration > Troubleshooting and support tools and make sure the license support period has not expired.



1. **Software maintenance**: upgrade at any time during this period.

If your support period has expired, follow the prompts to renew your license and reapply it before upgrading.

Have our supported platforms changed?

Check the Supported Platforms page for the version of Confluence you are upgrading to. This will give you info on supported operating systems, databases and browsers.

Good to know:

- The Confluence installer includes Java (JRE) and Tomcat, so you won't need to upgrade these separately.
- If you need to upgrade your database, be sure to read the upgrade notes for the Confluence version you plan to upgrade to (and any in-between) to check for any database configuration changes that you may need to make.

Do you need to make changes to your environment?

Newer Confluence versions sometimes require changes to your environment, such as providing more memory or adjusting your reverse proxy settings.

Good to know:

We use Upgrade Notesto communicate changes that will impact you, such as:

- Changes to supported databases, memory requirements or other changes that will impact your environment.
- Features that have significantly changed or been removed in this release.
- Actions you may need to take in your instance or environment immediately after the upgrade.

It's important to read the notes for the version you're upgrading to and those inbetween. For example, if you are upgrading from 5.8 to 5.10 you should read the upgrade notes for 5.9 and 5.10.

Plan your upgrade



Create a custom upgrade plan

Planning an upgrade? You can instantly generate a tailored upgrade plan from within Confluence. Head to Seneral Configuration > Plan your upgrade

You'll need to have a compatible version of the Troubleshooting and Support tools app installed. Lear n more

1. Determine your upgrade path

Use the table below to determine the most efficient upgrade path from your current version to the latest versions of Confluence.

Your Version	Recommended upgrade path to Confluence 7	
2.7 or earlier	Upgrade to 2.7.4 then upgrade to 3.5.17, 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, and follow paths below.	
4.0 to 4.3	Upgrade to 5.10.x , and follow paths below.	
5.0 to 7.x	Upgrade directly to the latest version of Confluence 7.	

If you are upgrading to the next bug fix update (for example, from 7.9.0 to 7.9.4), you can upgrade with no downtime.



Confluence 7 is a major upgrade

Be sure to check the Confluence Upgrade Matrix, take a full backup, and test your upgrade in a non-production environment before upgrading your production site.

Enterprise releases

A Long Term Support release is a feature release that gets backported critical security updates and critical bug fixes during its entire two-year support window. If you can only upgrade once a year, consider upgrading to a Long Term Support release.Learn more

Long Term Support releases were originally referred to as Enterprise Releases.

2. Complete the pre-upgrade checks

- 1. Check the Upgrade Notes for the version you plan to upgrade to (and any in between).
- 2. Go to Seneral Configuration Plan your upgrade then select the version you want to upgrade to. This will run somepre-upgrade checks.
- 3. Go to > General Configuration > Troubleshooting and support tools to run the health check.

If the software maintenance period included in your license has expired you can keep using Confluence, but you'll need to renew before you can upgrade.

Go to General Configuration > License Details and follow the prompts to renew your license.

If you are using the **embedded (trial) database**you should migrate to a different database **before** up grading. See Embedded H2 Database for more information.

Database character encoding must be set to UTF+8 (or AL32UTF8 for Oracle databases). You will not be able to upgrade to current Confluence versions unless you have the correct character encoding.

- 4. Go to > Manage apps and scroll down to the Confluence Update Checkto check the compatibility of your Marketplace apps.
- 5. Choose the version you plan to upgrade to then hit Check.

If your users rely on particular Marketplace apps, you may want to wait until they are compatible before upgrading Confluence. Vendors generally update their apps very soon after a major release.

Good to know:

- You can disable an app temporarily while you upgrade if it is not yet compatible.
- Compatibility information for Atlassian Labs and other free apps is often not available immediatley after a new release. In many cases the app will still work, so give it a try in a test site before upgrading your production site.

3. Upgrade Confluence in a test environment

- Create a staging copy of your current production environment.
 See Create a staging environment for upgrading Confluence for help creating an environment to test your upgrade in.
- 2. Follow the steps below to upgrade your test environment.
- 3. Test any unsupported user-installed apps, customizations (such as custom theme or layouts) and proxy configuration (if possible) before upgrading your production environment.

Upgrade Confluence

4. Back up

Back up your database and confirm the backup was created properly.
 If your database does not support online backups you'll need to stop Confluence first.

Once you've confirmed your database backup was successful, you can choose to disable the automatic generation of an upgrade recovery file, as this process can take a long time for sites that are medium sized or larger.

2. Back up your installation directory

The installer will completley replace this directory, so any files you've added (such as a keystore or SSL certificate) won't be retained. The installation wizard will back up this directory before starting the upgrade, but you should also back it up manually first.

3. Back up yourhome directory.

The installation wizard gives you the option to also back up your home directory as part of the installation process, but you should also back up this directory manually before starting the upgrade.

You can find the location of your home directory in the <installation-directory>/confluence/WEB-INF/classes/confluence-init.properties file.

This is where your search indexes and attachments are stored. If you store attachments outside the Confluence Home directory, you should also backup your attachments directory.

5.Download Confluence

Download the installer for your operating system.

- Latest versionhttps://www.atlassian.com/software/confluence/download
- Older versionshttps://www.atlassian.com/software/confluence/download-archives

6. Run the installer

1. Run the installer.

Run the .exe file. We recommend using a Windows administrator account.

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.

Change to the directory where you downloaded Confluence then execute this command to make the installer executable:

```
$ chmod a+x atlassian-confluence-X.X.X-x64.bin
```

Where x.x.x is is the Confluence version you downloaded.

Next, run the installer we recommend usingsudoto run the installer:

```
$ sudo ./atlassian-confluence-X.X.X-x64.bin
```

You can also choose to run the installer with root user privileges.

- 2. Follow the prompts to upgrade Confluence:
 - a. When prompted choose **Upgrade an existing Confluence installation** (for Linux users this is option 3).
 - b. Make sure the **Existing Confluence installation directory** suggested by the wizard is correct (especially important if you have multiple Confluence installations on the same machine).
 - c. **Back up Confluence home**is strongly recommended. This will create a .zip backup of the Confluence home and installation directories.
 - d. The installation wizard notifies you of customizations in the Confluence Installation directory. Make a note of these as you'll need to reapply them later.

The installation wizard's ability to notify you about customizations 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 theconf luencesubdirectory 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 thebindirectory or modifications to theserver.xmlfile (such as anSSL configuration).

You won't be notified about files you've added to the installation directory, so be sure to back them up first.

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. Depending on the size of your instance and the number of upgrade tasks to be run, this step may take a few minutes or several hours.

After the upgrade

7. Copy your database driver

If you're using an Oracle or MySQL database, you'll need to copy the jdbc driver jar file from your existing Confluence installation directory toconfluence/WEB-INF/libin your new installation directory.

Microsoft SQL and Postgres users can skip this step.

8. Reinstall the service if required (Windows only)

If you run Confluence as a service on Windows you should delete the existing service then re-install the service by running<install-directory>/bin/service.bat.

This makes sure the service gets the most recent JVM options.

9. Re-apply any modifications

During the upgrade the wizard migrated the following from your existingConfluence installation:

- TCP port values in your <install-directory>/conf/server.xmlfile.
- Location of your Confluence home directory in<install-directory>/confluence/WEB-INF /classes/confluence-init.properties.

All other customizations, including CATALINA_OPTS parameters in your <install-directory>/bin /setenv.sh/setenv.bat files, need to be reapplied manually.

Any other configurations, customizations (including any other modifications in the <installdirectory>/conf/server.xmlfile), the path to your own Java installation in <installdirectory>/bin/setjre.sh, or setjre.bat, or additional files added to the installation directory are **not migrated** during the upgrade and need to be reapplied manually.

- 1. Stop your upgraded Confluence instance.
- 2. Edit each file, and reapply the customizations in your upgraded Confluence Installation directory.
- 3. Copy over any additional files (such as keystore or SSL certificate)
- 4. Restart the upgraded Confluence instance.

Westrongly recommendyou test your customizations in a test instance prior to upgrading your production instance as changes may have been made to Confluence that make your customizations unusable.



 Edit the new file manually, rather than copying over the old file, as the default configuration in these files may have changed between Confluence versions.

System properties can change from time to time. Be sure to check you're using the correctRecognize d System Properties.

10. Update your apps (add-ons)

You can update any apps that are compatible with the new version of Confluence.

1. Go to

> Manage apps

2. Update your apps to the supported versions.



At this stage, it can be useful to clear your plugin cache. Learn how to do this

This is optional, but can be useful to avoid any issues with third-party apps and plugins.

11. Update your reverse proxy and check you can access Confluence

If you are upgrading from Confluence 5.x to Confluence 6.xyou will need to modify your reverse proxy (if used) to add Synchrony, which is required for collaborative editing. SeeProxy and SSL considerationsfor more information on the changes you'll need to make to your proxy config.

Once your upgrade is complete, you should access Confluence (via your reverse proxy, not directly) and:

- Head to Opening General Configuration Collaborative editing and check the Synchrony status is running
- Edit any page to check that your browser can connect to Synchrony.

SeeTroubleshooting Collaborative Editingfor suggested next steps if Synchrony is not running or you see an error in the editor, as you may have a misconfigured reverse proxy.

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.

Can't proceed with upgrade because license has expired

If your license has expired and was not renewed and reapplied before upgrading you will receive errors during the upgrade process. Seeupgrading beyond current license periodfor information on how to resolve this problem.

Can't 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.enabledsystem propertyto falsewill 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.

Confluence doesn't start

Incompatible Marketplace apps can occasionally prevent Confluence from starting successfully. You can troubleshoot the problem by starting Confluence with all user installed apps temporarily disabled. See Start and Stop Confluence for more info.

Collaborative editing errors

If Synchrony is not running or you see an error, head to Troubleshooting Collaborative Editingfor info on how to get collaborative editing up and running in your environment. The most common problems are a misconfigured reverse proxy or port 8091 not being available for Synchrony.

Space directory is empty after the upgrade

If you are upgrading from Confluence 6.3 or earlier, there's a known issue where spaces do not appear in the space directory. You'll need to reindex your site after upgrading to fix this.

You can also refer to the Upgrade Troubleshootingguide in the Confluence Knowledge Base, or check for answers from the community at Atlassian Answers.

Upgrading Beyond Current Licensed Period

This page covers what to do ifyou have upgraded Confluence to a version beyond your current license entitlement.

Related pages:

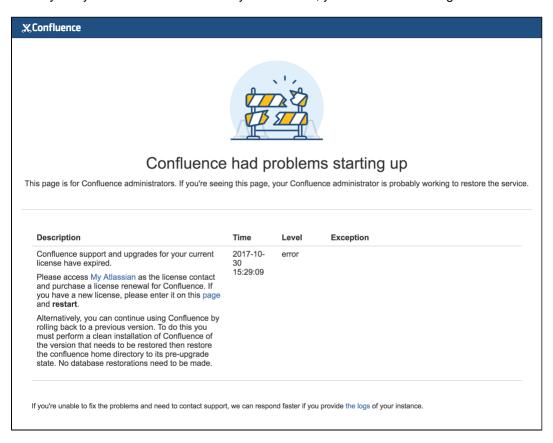
- Upgrading Confluence
- Working with Confluence Logs

License warnings

During the upgrade you will see an error similar to the following in yourapplication logs.

```
ERROR [confluence.upgrade.impl.DefaultUpgradeManager] runUpgradePrerequisites
Current license is not valid: SUPPORT_EXPIRED
```

When you try to access Confluence in your browser, you'll see this warning:



Updating the Confluence license

- 1. Head to my.atlassian.com to renew your license or purchase a new license.
- 2. Follow the prompts on the warning screen to enter your new license key.

Update Confluence License				
You do not have	a license installed for this Confluence installation.			
You can enter your license key below, or obtain an evaluation key online.				
License	Save			
	Save			

3. Restart Confluence to pick up the license change. You should now be able to log in to Confluence as normal.

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] entireupgradeFinished Upgrade completed successfully
```

If you do not see the line in your log similar to the one above, this means that your upgrade may not have completed successfully. Please check our Upgrade Troubleshooting documentation to check for a suitable recommendation or fix.

Upgrade Checklist

Here's a recommended list of things to check after completing an upgrade

1.The editor

Edit a page to check your browser can connect to Synchrony, which is required for collaborative editing. SeeTro ubleshooting Collaborative Editing if you are not able to edit a page.

2. 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.

3. Search

Try searching for content, for example pages, attachments or user names. Check that the expected results are returned. If you notice any problems, you may want to take advantage of the maintenance window to rebuild the indexes from scratch. See Content Index Administration.

4. 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.

5. Attachments

Confirm that attachments are accessible and searchable.

6. Marketplace apps

Outdated third-party apps can cause upgrade failure. Quite often, they will just be incompatible and simply do not work anymore. If you discover that your app is no longer working, please check for the latest version for your app in the The Atlassian Marketplaceor check for compatibility in the Universal Plugin Manager.

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 earler) 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 customized 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.
2 ContentEntityObjects failed.
2 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 new Confluence Editor.

However, in some cases a batch of content may actually have completely failed to migrated. 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 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'
```

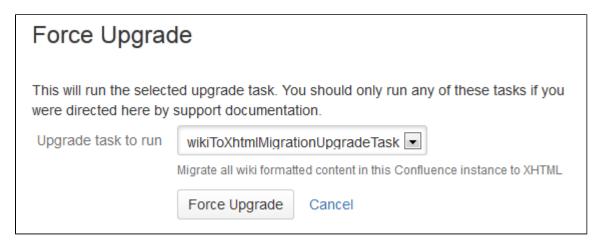
The solution is to ensure you manually re-run the site migration after the restart.

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 remigration 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
- Select wikiToXhtmlMigrationUpgradeTask in the Upgrade task to run dropdown list.
- 3. ChooseForce Upgrade.



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:

- 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



(i) 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 recommed 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
- Page 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.

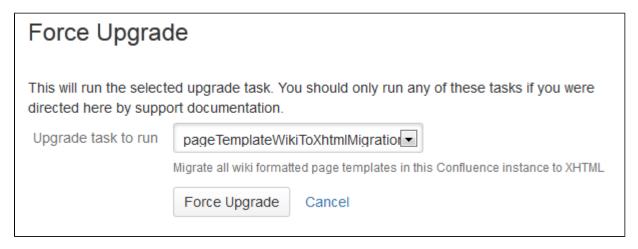
Re-running the migration

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
- Select pageTemplateWikiToXhtmlMigrationUpgradeTask in the Upgrade task to run dropdown list.
- 3. ChooseForce 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

In this guide we'll run you through upgrading your Confluence site to the latest Confluence version on Windows or Linux using the zip / tar.gz file.

Upgrading to any later version is free if you have current software maintenance. See ourLicensing FAQto find out more.



Other ways to upgrade Confluence:

- Installer the simplest way to upgrade Confluence.
- Data Center upgrade your Data Center cluster.
- Rolling upgrade- upgrade your Data Center cluster to the latest available bug fix version, with no downtime.

XML backups should**not**be used to upgrade Confluence.

Before you begin

Before you upgrade Confluence, there's a few questions you need to answer.

Is manual the right upgrade method for you?

You can choose to upgrade using the installer, or manually using a zip or tar.gz file. In most cases the installer is the easiest way to upgrade your Confluence instance.

You will need to upgrade manually if you are:

- moving to another operating system or file location as part of this upgrade.
- upgrading from Confluence 3.5 or earlier
- upgrading from Confluence 5.6 or earlier and previously used the EAR/WAR distribution to deploy Confluence into an existing application server.
- performing a rolling upgrade, and you need to upgrade each node individually.

On this page:

Before you begin Plan your upgrade

- 1. Determine your upgrade path
- 2. Complete the pre-upgrade checks
- 3. Upgrade Confluence in a test environment

Upgrade Confluence

- 4. Back up
- 5. Download Confluence
- 6. Extract the file and upgrade Confluence

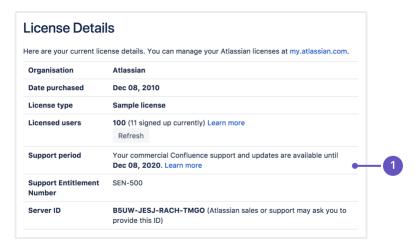
After the upgrade

- 7. Reinstall the service (Windows only)
- 8. Re-apply any modifications
- 9. Update your reverse proxy and check you can access Confluence

Troubleshooting

Are you eligible to upgrade?

To check if software maintenance is current for your license, go to Seneral Configuration > Troubleshooting and support toolsand make sure the license support period has not expired.



1. **Software maintenance**: upgrade at any time during this period.

If your support period has expired, follow the prompts to renew your license and reapply it before upgrading.

Have our supported platforms changed?

Check the Supported Platforms page for the version of Confluence you are upgrading to. This will give you info on supported operating systems, databases and browsers.

Good to know:

- If you need to upgrade Java, remember to update your JAVA_HOME variable to the new version.
- The Confluence installer includes Tomcat, so you won't need to upgrade it separately.
- If you need to upgrade your database, be sure to read the upgrade notes for the Confluence version you plan to upgrade to (and any in-between) to check for any database configuration changes that you may need to make.

Do you need to make changes to your environment?

Newer Confluence versions sometimes require changes to your environment, such as providing more memory or adjusting your reverse proxy settings.

Good to know:

We use Upgrade Notesto communicate changes that will impact you, such as:

- Changes to supported databases, memory requirements or other changes that will impact your environment.
- Features that have significantly changed or been removed in this release.
- Actions you may need to take in your instance or environment immediately after the upgrade.

It's important to read the notes for the version you're upgrading to and those inbetween. For example, if you are upgrading from 5.8 to 5.10 you should read the upgrade notes for 5.9 and 5.10.

Plan your upgrade

1. Determine your upgrade path

Use the table below to determine the most efficient upgrade path from your current version to the latest versions of Confluence.

Your Version	Recommended upgrade path to Confluence 7	
2.7 or earlier	Upgrade to 2.7.4 then upgrade to 3.5.17, 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 , and follow paths below.	
4.0 to 4.3	Upgrade to 5.10.x , and follow paths below.	
5.0 to 7.x Upgrade directly to the latest version of Confluence 7.		

If you are upgrading to the next bug fix update (for example, from 7.9.0 to 7.9.4), you can upgrade with no downtime.



Confluence 7 is a major upgrade

Be sure to check the Confluence Upgrade Matrix, take a full backup, and test your upgrade in a nonproduction environment before upgrading your production site.

2. Complete the pre-upgrade checks

- 1. Check the Upgrade Notes for the version you plan to upgrade to (and any in between).
- 2. Go to General Configuration Plan your upgrade then select the version you want to upgrade to. This will run somepre-upgrade checks.
- 3. Go to O-> General Configuration > Troubleshooting and support tools to run the health check.

If the software maintenance period included in your license has expired you can keep using Confluence, but you'll need to renew before you can upgrade.

Go to General Configuration > License Details and follow the prompts to renew your license.

If you are using the embedded (trial) database you should migrate to a different database before up grading. SeeEmbedded H2 Databasefor more information.

Database character encoding must be set to UTF+8 (or AL32UTF8 for Oracle databases). You will not be able to upgrade to current Confluence versions unless you have the correct character encoding.

- 4. Go to O > Manage appsand scroll down to the Confluence Update Checkto check the compatibility of your Marketplace apps.
- 5. Choose the version you plan to upgrade to then hit **Check**.

If your users rely on particular Marketplace apps, you may want to wait until they are compatible before upgrading Confluence. Vendors generally update their apps very soon after a major release.

Good to know:

- You can disable an app temporarily while you upgrade if it is not yet compatible.
- Compatibility information for Atlassian Labs and other free apps is often not available immediately after a new release. In many cases the app will still work, so give it a try in a test site before upgrading your production site.

3. Upgrade Confluence in a test environment

- Create a staging copy of your current production environment.
 See Create a staging environment for upgrading Confluence for help creating an environment to test your upgrade in.
- 2. Follow the steps below to upgrade your test environment.
- 3. Test any unsupported user-installed apps, customizations (such as custom theme or layouts) and proxy configuration (if possible) before upgrading your production environment.

Upgrade Confluence

4. Back up

Back up your database and confirm the backup was created properly.
 If your database does not support online backups you'll need to stop Confluence first.

Once you've confirmed your database backup was successful, you can choose to disable the automatic generation of an upgrade recovery file, as this process can take a long time for sites that are medium sized or larger.

2. Back up yourinstallation directoryandhome directory.

You can find the location of your home directory in the <installation-directory> /confluence/WEB-INF/classes/confluence-init.properties file.

This is where your search indexes and attachments are stored. If you store attachments outside the Confluence Home directory, you should also backup your attachments directory.

5.Download Confluence

7Zip or Winzip.

Download the appropriate file for your operating system -https://www.atlassian.com/software/confluence/download

6. Extract the file and upgrade Confluence

- Stop Confluence.
 SeeUsing read-only mode for site maintenance if you need to provide uninterrupted access.
- 2. Extract (unzip) the files to adirectory (this is your new installation directory, and must be different to your existing installation directory)
 Note: There are some known issues with unzipping the archive on Windows. We recommend using
- 3. Edit <Installation-Directory>\confluence\WEB-INF\classes\confluence-init. propertiesfile to point to yourexistingConfluence home directory.
- 4. If you're using an Oracle or MySQL database, you'll need to copyyour jdbc driver jar file from your existing Confluence installation directory toconfluence/WEB-INF/libin your new installation directory.
- 5. There are some additional steps you make need to take if:
 - you are running Confluence as aWindows Service

If you are running Confluence as a Windows service, go to 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 Start Confluence Automatically on Windows as a Service.

⚠To remove the service installed by the Confluence installer, you'll need to run<confluence auto installer installation folder>\UninstallService.bat.

You are running Confluence on a different port (not the default 8090)

If you are not running Confluence on port 8090 update <Installation-Directory>\conf\server.xmlfile to include your ports.

6. Start your new Confluence. You should not see the setup wizard.

After the upgrade

7. Reinstall the service (Windows only)

If you run Confluence as a service on Windows you should delete the existing service then re-install the service by running<install-directory>/bin/service.bat.

This makes sure the service gets the most recent JVM options.

8. Re-apply any modifications

If you have customized Confluence (such as an SSL configuration in theserver.xmlfile, or CATALINA_OPT SorJAVA_OPTSparameters in yourconfluence-init.propertiesfile), you'll need to perform the following steps after the upgrade is complete:

- 1. Stop your upgraded Confluence instance.
- Reapply the customizations to the relevant files in the newly upgraded Confluence Installation directory.
- 3. Restart the upgraded Confluence instance.

We**strongly recommend**you test your customizations in a test instance prior to upgrading your production instance as changes may have been made to Confluence that make your customizations unsuable.

9. Update your reverse proxy and check you can access Confluence

If you are upgrading from **Confluence 5.x to Confluence 6.x**you will need to modify your reverse proxy (if used) to add Synchrony, which is required for collaborative editing. SeeProxy and SSL considerations for more information on the changes you'll need to make to your proxy config.

Once your upgrade is complete, you should access Confluence (via your reverse proxy, not directly) and:

- Head to General Configuration>Collaborative editing and check the Synchrony status is running
- Edit any page to check that your browser can connect to Synchrony.

SeeTroubleshooting Collaborative Editingfor suggested next steps if Synchrony is not running or you see an error in the editor, as you may have a misconfigured reverse proxy.

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.

• Can't proceed with upgrade because license has expired

If your license has expired and was not renewed and reapplied before upgrading you will receive errors during the upgrade process. Seeupgrading beyond current license periodfor information on how to resolve this problem.

• Collaborative editing errors

If Synchrony is not running or you see an error, head to Troubleshooting Collaborative Editingfor info on how to get collaborative editing up and running in your environment. The most common problems are a misconfigured reverse proxy or port 8091 not being available for Synchrony.

• 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 theconfluence.upgrade.recovery.file.enabledsystem propertyto falsewill 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 Troubleshootingguide in the Confluence Knowledge Base, or check for answers from the community at Atlassian Answers.

Create a staging environment for upgrading Confluence

When you upgrade Confluence we strongly recommend performing the upgrade in a test environment before upgrading your production site. In this guide we'll refer to this test environment as *st aging*.

Most Confluence licenses include a free developer license for use in a staging environment. SeeHow to get a Confluence Developer licenseto find out how to access your license.

On this page:

- Create a staging environment
- Additional configuration options
- Upgrade your staging environment

Create a staging environment

1. Replicate your environment

Your staging environment should closely replicate your real-live environment (production), including any reverse proxies, SSL configuration, or load balancer (for Data Center). You may decide to use a different physical server or a virtualized solution. The main thing is to make sure it is an appropriate replica of your production environment.

For the purposes of these instructions, we assume your staging environment is physically separate fromyour production environment, and has the same operating system (and Java version if you've installed Confluence manually).

2. Replicate your database

To replicate your database:

- 1. Back up your production database. Refer to the documentation for your database for more info on the best way to do this.
- 2. Install your database on the staging server and restore the backup.

The steps for restoring your database backup will differ depending on your chosen database and backup tool. Make sure:

- Your new staging database has a different name from your production database.
- Your staging database user account has thesameusername and password as your production database user account.
- Character encoding and other configurations are the same as your production database (for example character encoding should be Unicode UTF-8 (or AL32UTF8 for Oracle databases).

3. Replicate Confluence

To replicate Confluence, make a copy of your Confluence installation and point it to your staging database. These instructions are for Confluence Server (for Data Center there are some additional steps before you start Confluence).

- 1. Copy your entire **production installation directory** to your staging server.
- 2. Copy your entire **production home directory** to your staging server.
- 3. Edit <installation-directory>/confluence/WEB-INF/classes/confluence-init. propertiesto point to your staging home directory.
- 4. Edit<home-directory>/confluence.cfg.xmlor<installation-directory>/server.xmlt o point to your staging database.

 $\label{local-post} $$\operatorname{postgresql://localhost:5432/confluencestaging$

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="postgres"
    password="postgres"
    driverClassName="org.postgresql.Driver"
    url="jdbc:postgresql://localhost:5432/confluencestaging"
    maxTotal="60"
    maxIdle="20"
    validationQuery="select 1" />
```

Start Confluence with the followingSystem Properties to make sure your staging site does not send notifications to real users.

```
-Datlassian.notifications.disabled=true
-Datlassian.mail.senddisabled=true
```

- 6. Head to http://localhost:<port>and log in to Confluence on your staging server.
- 7. Go to Some General Configuration and change the base URL of your staging site (for example mysite .staging.com)
- 8. Go to General Configuration> License Details and apply your development license.
- 9. Go to General Configuration System Information and check that Confluence is correctly pointing to your staging database, and staging home directory.



It's essential to check that you are not still connected to your production database.

Additional steps for Data Center

If you have Confluence Data Center, the process is much the same as for Confluence Server described above. You will copy each local home and installation directory to each staging node, and then:

- 1. Copy the production shared home directory to the staging server.
- 2. Edit<local-home-directory>/confluence.cfg.xmlto point to your staging shared home directory. This change **must** be made on every staging node.

Changes to the <installation-directory>/confluence/WEB-INF/classes/confluence-init.propertiesand <home-directory>/confluence.cfg.xml must be made on every staging node.

When it comes time to start Confluence, start one node at a time, as usual.

4. Replicate external user management (optional)

If you're managing users in Jira, Crowd, or in an external LDAP directory you can:

- replicate Jira, Crowd, or your external directory in your staging environment and point your Confluence staging site to your staging external directory (recommended).
- provide your staging server with network or local access to the same hosts as your production server.

Additional configuration options

There are a number of additional things you may want to change in your staging environment, to make sure it does not interact with your production environment, or to clearly differentiate it for users.

Bypass single sign-on

If you've configured single sign-on, you might want to bypass this in your staging environment.

SeeEnable default login page to bypass SAML in Confluence Data Center.

Disable CDN

If you've configured a CDN to cache static assets, you may experience problems with broken resources as the CDN configuration is for your production environment, not the staging environment.

To find out how to disable CDN, seeConfigure your CDN for Confluence Data Center.

Modify application links (recommended)

If you have application links between Confluence and other Atlassian applications you should change the server ID on each staging application. SeeHow to change the server ID of Confluence andChange the server ID for an instance of Jira server for Jira.

If you don't change the server ID and update your application links there is a chance that when you create a new application link in production it will point to your staging server instead.

To review the Application Links manually in the database, use the following following SQL query:

```
select * from bandana where bandanakey like 'applinks%';
```

Modify external gadgets

If you have external gadgets configured, you can update these from the database, using the following SQL query:

```
select * from bandana where bandanakey = 'confluence.ExternalGadgetSpecStore.specs'
```

Change the global color scheme

If can be helpful to use a different color scheme on your staging site, to differentiate it from your production site. See Customizing Color Schemesfor how to do this.

You can also find this data in the database using the following SQL query:

```
select * from bandana where bandanakey = 'atlassian.confluence.colour.scheme';
```

Change the instance name (recommended)

It is a good idea to change the name of your staging site, to differentiate it from your production site. Head to General Configuration and update the Site Title if Confluence is running.

If Confluence is not running, you can do this from the database. You can find the site title using the following SQL query:

```
select * from bandana where bandanakey = 'atlassian.confluence.settings';
```

The attribute you are looking for issetTitle.

Add a banner

It can be useful to add a banner to your staging site, to provide useful information like the date of the last refresh, or who to contact if you want to make changes.

If you have a Confluence Data Center license, you can do this by enabling the banner that is used by readonly mode (you don't need to enable read-only mode to use the banner). If you have a Confluence Server license, you can manually add a banner using HTML. Head to Configuration Custom HTML. Remember to close your tags properly, or Confluence may not display correctly.

If you want to add a banner before starting Confluence, you can do it in the database. You can find the custom HTML using the following SQL query:

```
select * from bandana where bandanakey = 'atlassian.confluence.settings';
```

The attribute you are looking for iscustomHtmlSettingsafterBodyStart

Disable specific plugins

You might want to disable specific plugins or check whether these plugins are already disabled or not. See the How to reset all Confluence plugins back to their default state through the databaseknowledge base article to find how to do this.

You can also disable plugins in Confluence in 6.1+ using Java system properties.

Upgrade your staging environment

Once you have created your staging environment, you can upgrade it in the same way you would your production environment.

Make a note of how long the upgrade takes, as this information will help you plan your production system outage and communicate with your users.

You can also use your staging environment to test any customizations or essential Marketplace apps in your site.

Upgrade Confluence without downtime

If you run Confluence Data Center in a cluster, you may be able to upgrade Confluence without any downtime for your users. This method is known as a rolling upgrade.

In a rolling upgrade, your site is put into upgrade mode, which temporarily allows nodes running different Confluence versions to join the cluster. As you take each node offline to upgrade it, the other active nodes keep your Confluence site available to users. Once all nodes have been upgraded in turn, you finalize the upgrade and turn off upgrade mode.

On this page:

- Can I upgrade without downtime?
- Before you begin
- Prepare for the rolling upgrade
 - 1.Complete pre-upgrade checks
 - 2. Prevent the installation or upgrade of apps during the upgrade period
 - 3. Back up Confluence Data Center
 - 4. Set up a staging environment to test the rolling upgrade
- Perform the rolling upgrade

Can I upgrade without downtime?

Whether you can upgrade your Confluence Data Center cluster without downtime depends on the version you are upgrading from, and the version you are upgrading to.Learn more about the different types of releases.

Upgrading from	Upgrading to		
	Bugfix	Feature	Platform
Confluence 7.8 and earlier	Requires downtime	Requires downtime	Requires downtime
Confluence 7.8 to 7.13	No downtime (for example, from 7.12.0 to 7.12.2)	Requires downtime	Requires downtime
Confluence 7.14 and later	No downtime (for example, from 7.14.0 to 7.14.2)	No downtime when upgrading to the nextfeature version (for example, from 7.14.x to 7.15.x)	Requires downtime
	,	Requires downtime if the upgrade spans <i>more</i> than one feature version (for example, from 7.14.x to 7.17.x)	

Before you begin

Before you start planning a rolling upgrade, theres a few questions you need to answer.

Does my Confluence deployment support rolling upgrades?	You can only perform a rolling upgrade with no downtime on a multi-node Confluence cluster. Clustering is only supported on a Confluence Data Center license. In addition, a rolling upgrade involves enabling upgrade mode, which is only available in Confluence Data Center. Learn more about multi-node clustering in Confluence
Do I have enough nodes to support user requests during the rolling upgrade?	You need to take a node offline to upgrade it. During this time, other active nodes will take over the offline nodes workload. Make sure you have enough active nodes to handle user traffic at any given time. If possible, add a node temporarily to your cluster to compensate for offline nodes.
Is the version compatible with rolling upgrades?	Whether you can upgrade without downtime depends on the version you are upgrading from, and the version you are upgrading to. The pre-upgrade check will confirm whether you can upgrade without downtime.

Prepare for the rolling upgrade

1.Complete pre-upgrade checks

- 1. Check the Upgrade Notes for the version you plan to upgrade to (and any in between).
- 2. Go to General Configuration > Plan your upgrade then select the version you want to upgrade to. This will run somepre-upgrade checks.
- 3. Go to O-> General Configuration > Troubleshooting and support tools to run the health check.

If the software maintenance period included in your license has expired you can keep using Confluence, but you'll need to renew before you can upgrade.

Go to General Configuration > License Details and follow the prompts to renew your license.

If you are using the **embedded (trial) database**you should migrate to a different database **before** up grading. See Embedded H2 Database for more information.

Database character encoding must be set to UTF+8 (or AL32UTF8 for Oracle databases). You will not be able to upgrade to current Confluence versions unless you have the correct character encoding.

- 4. Go to O > Manage appsand scroll down to the Confluence Update Checkto check the compatibility of your Marketplace apps.
- 5. Choose the version you plan to upgrade to then hit **Check**.

If your users rely on particular Marketplace apps, you may want to wait until they are compatible before upgrading Confluence. Vendors generally update their apps very soon after a major release.

Good to know:

- You can disable an app temporarily while you upgrade if it is not yet compatible.
- Compatibility information for Atlassian Labs and other free apps is often not available immediatley after a new release. In many cases the app will still work, so give it a try in a test site before upgrading your production site.
- 2. Prevent the installation or upgrade of apps during the upgrade period

If you manage Confluence with a team of admins, schedule the rolling upgrade with them. Notify them to postpone any app installs or upgrades until after the rolling upgrade. Installing or upgrading apps during a rolling upgrade could result in unexpected errors.

3. Back up Confluence Data Center

Site Backup and Restorelists useful resources, along with recommendations for manual and automated backups. In particular, Production Backup Strategyrecommends specific methods for backing up larger Confluence sites.

If your deployment is hosted on AWS, we recommend that you use the AWS native backup facility, which utilizes snapshots to back up your site. For more information, seeAWS Backup.

4. Set up a staging environment to test the rolling upgrade

We strongly recommend that you perform the rolling upgrade on a staging or test environment first.

- Create a staging copy of your current production environment.
 See Create a staging environment for upgrading Confluence for help creating an environment to test your upgrade in.
- 2. Follow the steps below to upgrade your test environment.
- 3. Test any unsupported user-installed apps, customizations (such as custom theme or layouts) and proxy configuration (if possible) before upgrading your production environment.

Perform the rolling upgrade

There are three methods for performing a rolling upgrade, depending on what orchestration tools your deployment uses.

Method	Description	Instructions
Manual upgrade	A manual upgrade is suitable for deployments that feature minimal orchestration, particularly in node upgrades. If your deployment is based on ourAzure templates, you'll also need to perform a manual upgrade.	Upgrade a Confluence cluster manually without downtime
AWS CloudFo rmation	If your deployment is defined by an AWS CloudFormation template (like ourAWS Quick Start), then you can use the same template to orchestrate your upgrade.	Upgrade a Confluence cluster on AWS without downtime
API- driven	You can orchestrate the entire rolling upgrade process through API calls.	Upgrade a Confluence cluster through the API without downtime

Upgrade a Confluence cluster manually without downtime

This document provides step-by-step instructions on how to perform a rolling upgrade on deployments with little or no automation. These instructions are also suitable for deployments based on ourAzure templates.

For an overview of rolling upgrades (including planning and preparation information), seeUpgrade Confluence without downtime.

Step 1: Download upgrade files

Before you start the upgrade, you'll need to download the right Confluence version. You'll be installing this on each node. Remember, you can only upgrade to a higher bug fix version (for example, from Confluence 7.9.0 to 7.9.4) or to the next feature version (for example, from Confluence 7.14.2 to 7.15.0).

Download Confluence

Alternatively, go to General Configuration Plan your upgrade to run the pre-upgrade checks and download a compatible bug fix version.

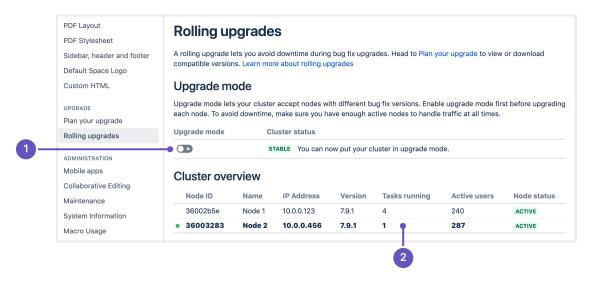
Step 2: Enable upgrade mode

You need System Administrator global permissions to do this.

To enable upgrade mode:

- 1. Go to Co Seneral Configuration Rolling upgrades.
- 2. Select the **Upgrade mode** toggle (1).

Screenshot: The Rolling upgrades screen.



The cluster overview can help you choose which node to upgrade first. The **Tasks running** (2) column shows how many long-running tasks are running on that node, and the **Active users** shows how many users are logged in. When choosing which node to upgrade first, start with the ones with the least number of tasks running and active users.

Upgrade mode allows your cluster to temporarily accept nodes running different Confluence versions. This lets you upgrade a node and let it rejoin the cluster (along with the other non-upgraded nodes). Both upgraded and non-upgraded active nodes work together to keep Confluence available to all users. You can disable upgrade mode as long as you havent upgraded any nodes yet.

Step 3: Upgrade the first node

With upgrade mode enabled, you can now upgrade your first node.



Start with the least busy node

We recommend that you start upgrading the node with the least number of running tasks and active users. You can check this on the Rolling upgrades page.

Start by shutting down Confluence gracefully on the node:

- 1. Access the node through a command line or SSH.
- 2. Shut down Confluence gracefully on the node. To do this, run the stop script corresponding to your operating system and configuration. For example, if you installed Confluence as a service on Linux, run the following command:

\$ sudo /etc/init.d/confluence stop Learn more about graceful Confluence shutdowns

A graceful shutdown allows the Confluence node to finish all of its tasks first before going offline. During shutdown, the node's status will be Terminating, and user requests sent to the node will be redirected by the load balancer to other Active nodes.



For nodes running on Linux or Docker, you can also trigger a graceful shutdown through the kill command (this will send a SIGTERM signal directly to the Confluence process).

3. Wait for the node to go offline. You can monitor its status on the Node status column of the Rolling upgrade pages Cluster overview section.

Once the status of the node is offline, you can start upgrading the node. Copy the Confluence installation file you downloaded to the local file system for that node.

To upgrade the first node:

- 1. Extract (unzip) the files to adirectory (this will be your new installation directory, and must be different to your existing installation directory)
- 2. Update the following line in the < Installation-Directory > \confluence \WEB-INF\classes\confluence-init.propertiesfile to point to the existing local homedirectory on that node.
- 3. If your deployment uses a MySQL database, copy the jdbc driver jar file from your existing Confluence installation directory to confluence/WEB-INF/libin your new installation directory. The jdbc driver will be located in either the < Install - Directory > / common/libor < Installation -Directory>/confluence/WEB-INF/libdirectories. SeeDatabase Setup For MySQL for more details.
- 4. If you run Confluence as a service:
 - On Windows, delete the existing service then re-install the serviceby running<install- directory>/bin/service.bat.
 - On Linux, update the service to point to the new installation directory (or use symbolic links to do this).
- 5. Copy any other immediately required customizations from the old version to the new one(for example if you are not running Confluence on the default ports or if you manage users externally, you'll need to update / copy the relevant files - find out more in Upgrading Confluence Manually).

If you configured Confluence to run as a Windows or Linux service, don't forget to update its service configuration as well. For related information, see Start Confluence Automatically on Windows as a Service or Run Confluence as a systemd service on linux.

6. Start Confluence, and confirm that you can log in and view pages before continuing to the next step.

As soon as the first upgraded node joins the cluster, your cluster status will transition to Mixed. This means that you wont be able to disable Upgrade mode until all nodes are running the same version.

Upgrade Synchrony (optional)

If you've chosen to let Confluence manage Synchrony for you (recommended), you don't need to do anything. Synchrony was automatically upgraded with Confluence.

If you're running your own Synchrony cluster, grab the newsynchrony-standalone.jarfrom theclocal-home>directory on your upgraded Confluence node. Then, perform the following steps on each Synchrony node:

- 1. Stop Synchrony on the node using either the start-synchrony. sh(for Linux) or start-synchrony. bat(for Windows) file from the Synchrony home directory.
- 2. Copy the newsynchrony-standalone. jarto your Synchrony home directory.
- 3. Start Synchrony as normal.

See Set up a Synchrony cluster for Confluence Data Centerfor related information.

Step 4: Upgrade all other nodes individually

After starting the upgraded node, wait for its status to change to Active in the Cluster overview. At this point you should check the application logs for that node, and log in to Confluence on that node to make sure everything is working. It's still possible to roll back the upgrade at this point, so taking some time to test is recommended.

Once you've tested the first node, you can start upgrading another node, following the same steps. Do this for each remaining node as always, we recommend that you upgrade the node with the least number of running tasks each time.

Step 5: Finalize the upgrade

The steps to finalize your upgrade will differ slightly depending on whether you are upgrading to a bugfix version, or to the next feature version which may require upgrade tasks to be run. You should do this soon as possible, as some tasks are put on hold while your cluster is in upgrade mode.

Finalize upgrade to a bugfix version

To finalize the upgrade:

- 1. Wait for the cluster status to change to **Ready to finalize**. This won't happen until all nodes are active, and running the same upgraded version.
- 2. Select the **Finalize upgrade** button.
- 3. Wait for confirmation that the upgrade is complete. The cluster status will change to **Stable**.

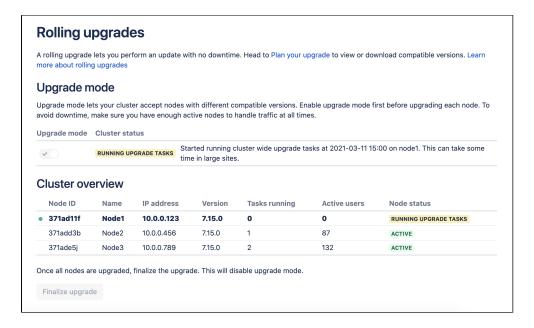
Your upgrade is now complete.

Finalize upgrade to a feature version

To finalize the upgrade:

- 1. Wait for the cluster status to change to **Ready to run upgrade tasks**. This won't happen until all nodes are active, and running the same upgraded version.
- 2. Select the Run upgrade tasks and finalize upgrade button.
- 3. One node will start running upgrade tasks. Tail the logs on this node if you want to monitor the process.
- 4. Wait for confirmation that the upgrade is complete. The cluster status will change to **Stable**.

Your upgrade is now complete.



Screenshot: One cluster node running upgrade tasks for the whole cluster.

Upgrade tasks make any required changes to your database and file system, for example changing the database schema or the way index files are stored in the local home directories.

There are a few things you should know about upgrade tasks:

- One cluster node will run the upgrade tasks on the database and other nodes. If there's a problem, logs will be written to the application log on this node.
- The status of other nodes in the cluster may change to Running upgrade tasks momentarily to
 indicate that an upgrade task is making a change to the file system on that node. The node actually
 running the upgrade tasks does not change.
- Depending on the the size or complexity of your data, some upgrade tasks can take several hours to complete. We generally include a warning in the upgrade notes for the particular version if an upgrade task is likely to take a significant amount of time.
- It's not necessary to direct traffic away from the node running upgrade tasks, but if you know the
 upgrade tasks are likely to be significant, you may want to do this to avoid any performance impact.

Troubleshooting

Node errors during rolling upgrade

If a nodes status transitions to **Error**, it means something went wrong during the upgrade. You cant finish the rolling upgrade if any node has an **Error** status. However, you can still disable Upgrade mode as long as the cluster status is still **Ready to upgrade**.

There are several ways to address this:

- Shut down Confluence gracefully on the node. This should disconnect the node from the cluster, allowing the node to transition to an **Offline** status.
- If you cant shut down Confluence gracefully, shut down the node altogether.

Once all active nodes are upgraded with no nodes in Error, you can finalize the rolling upgrade. You can investigate any problems with the problematic node afterwards and re-connect it to the cluster once you address the error.

Upgrade tasks failed error

If the cluster status changes to **Upgrade tasks failed**, this means that one or more upgrade tasks did not complete successfully and the upgrade has not been finalized. You should:

- 1. Check the application log on the node running the upgrade task for errors. The node identifier is included in the cluster status message.
- 2. Resolve any obvious issues (such as file system permissions, or network connectivity problems)
- 3. Select Re-run upgrade tasks and finalize upgrade to try again.

If upgrade tasks are still failing, and you can't identify a cause, you should contact our Support team for assistance. You may also want to roll back the upgrade at this point. We don't recommend leaving Confluence in upgrade mode for a prolonged period of time.

Roll back a node to its original version

How you roll back depends on the upgrade stage you have reached. See Roll back a rolling upgrade for more information.



Mixed status with Upgrade mode disabled

If a node is in an Error state with Upgrade mode disabled, you can't enable Upgrade mode. Fix the problem or remove the node from the cluster to enable Upgrade mode.

Disconnect a node from the cluster through the load balancer

If a node error prevents you from gracefully shutting down Confluence, try disconnecting it from the cluster through the load balancer. The following table provides guidance how to do so for popular load balancers.

NGINX	NGINX defines groups of cluster nodes through the upstream directive. To prevent the load balancer from connecting to a node, delete the node's entry from its corresponding upstream group. Learn more about the upstream directive in the ngx_http_upstream_module module.
HAProxy	With HAProxy, you candisable all traffic to the node by putting it in a maint state: set server <node hostname="" ip="" or=""> state maint</node>
	Learn more about forcing a server's administrative state.
Apache	You can disable a node (or "worker") by setting its activation member attribute to disabled. Learn more about advanced load balancer worker properties in Apache.
Azure Application Gateway	We provide adeployment template for Confluence Data Center on Azure; this template uses the Azure Application Gateway as its load balancer. The Azure Application Gateway defines each node as a target within a backend pool. Use the Edit backend pool interface to remove your node's corresponding entry. Learn more about adding (and removing) targets from a backend pool.

Traffic is disproportionately distributed during or after upgrade

Some load balancers might use strategies that send a disproportionate amount of active users to a newlyupgraded node. When this happens, the node might become overloaded, slowing down Confluence for all users logged in to the node.

To address this, you can also temporarily disconnect the node from the cluster. This will force the load balancer to re-distribute active users between all other available nodes. Afterwards, you can add the node again to the cluster.

Node won't start up

If a node is Offline or Starting for too long, you may have to troubleshoot Confluence on the node directly. See C onfluence Startup Problems Troubleshooting for related information.

Upgrade a Confluence cluster on AWS without downtime

This document provides step-by-step instructions on performing a rolling upgrade on an AWS deployment orchestrated through CloudFormation. In particular, these instructions are suitable for Confluence Data Center deployments based on our AWS Quick Starts.

For an overview of rolling upgrades (including planning and preparation information), seeUpgrade Confluence without downtime.

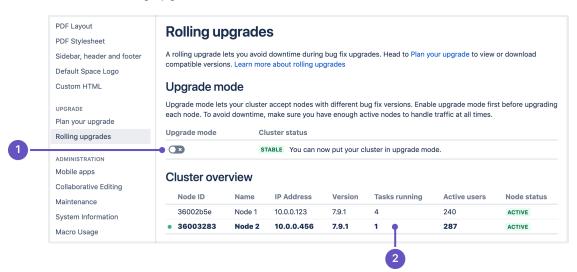
Step 1: Enable upgrade mode

You need System Administrator global permissions to do this.

To enable upgrade mode:

- 1. Go to Co Seneral Configuration Rolling upgrades.
- 2. Select the Upgrade mode toggle (1).

Screenshot: The Rolling upgrades screen.



The cluster overview can help you choose which node to upgrade first. The **Tasks running** (2) column shows how many long-running tasks are running on that node, and the **Active users** shows how many users are logged in. When choosing which node to upgrade first, start with the ones with the least number of tasks running and active users.

Upgrade mode allows your cluster to temporarily accept nodes running different Confluence versions. This lets you upgrade a node and let it rejoin the cluster (along with the other non-upgraded nodes). Both upgraded and non-upgraded active nodes work together to keep Confluence available to all users. You can disable upgrade mode as long as you havent upgraded any nodes yet.

Step 2: Find all the current application nodes in your stack

In AWS, note the Instance IDs of all running application nodes in your stack. These are all the application nodes running your current version. You'll need these IDs for a later step.

- 1. In the AWS console, go to**Services > CloudFormation**. Select your deployments stack to view its Stack Details.
- 2. Expand the **Resources** drop-down. Look for the **ClusterNodeGroup** and click its Physical ID. This will take you to a page showing the Auto Scaling Group details of your application nodes.
- 3. In the Auto Scaling Group details, click on the **Instances** tab. Note all of the Instance IDs listed there; you'll be terminating them at a later step.

Step 3: Update your CloudFormation template

Your deployment uses a CloudFormation template that defines each component of your environment. In this case, upgrading Confluence means updating the version of Confluence used in the template. During the upgrade, we highly recommend that you add a node temporarily to your cluster as well.

- 1. In the AWS console, go to Services > CloudFormation. Select your deployments stack to view its Stack Details.
- 2. In the Stack Details screen, click Update Stack.
- 3. From the Select Templatescreen, select Use current template and click Next.
- 4. Set the Version parameter to the version youre updating to. Since this is a rolling upgrade, you can only set this to a later bug fix version.
- 5. Add an extra node to your cluster. This will help ensure that your cluster won't have a shortage of nodes for user traffic. To do this, increase the value of the following parameters by 1:
 - Maximum number of cluster nodes
 - Minimum number of cluster nodes
- 6. Select **Next**. Click through the next pages, and then to apply the change using the **Update**button.

After updating the stack, you will have one extra node already running the new Confluence version. With Upgrade mode enabled, that node will be allowed to join the cluster and start work. Your other nodes won't be upgraded yet.

As soon as the first upgraded node joins the cluster, your cluster status will transition to Mixed. This means that you wont be able to disable Upgrade mode until all nodes are running the same version.

Once the new upgraded node is running an in an Active state, you shouldcheck the application logs for that node, and log in to Confluence on that node to make sure everything is working. It's still possible to roll back the upgrade at this point, so taking some time to test is recommended.

Once you've tested the first node, you can start upgrading another node. To do that, shut down and terminate the node AWS will then replace the node with a new one running the updated Confluence version.

Step 4: Upgrade another node



Start with the least busy node

We recommend that you start upgrading the node with the least number of running tasks and active users. On the Rolling upgrades page, youll find both in the Cluster overview section.

In Step 2, you noted the instance ID of each node in your cluster. Terminate the node where you gracefully shut down Confluence. To do this:

- 1. In the AWS console, go to Services > EC2. From there, click Running Instances.
- 2. Check the instance of matching the node where you gracefully shut down Confluence.
- 3. From the Actions drop-down, selectinstance State > Terminate.
- 4. Click through to terminate the instance.

Each time you terminate a node, AWS will automatically replace it. The replacement will be running the new version of Confluence. Once the new node's status is Active, you can move on to upgrading another node.

Step 5: Upgrade all other nodes individually

At this point, your cluster should have two nodes running the new version of Confluence. You can now upgrade other nodes. To do so, simply repeat the previous step on another node. As always, we recommend that you upgrade the node with the least number of running tasks each time.



⚠ If your deployment usesstandalone Synchrony, you may need to update the version used by each Synchrony node as well. To do this, terminate each Synchrony node one after the other after you upgrade all nodes to the new version.

Step 6: Finalize the upgrade

The steps to finalize your upgrade will differ slightly depending on whether you are upgrading to a bugfix version, or to the next feature version which may require upgrade tasks to be run. You should do this soon as possible, as some tasks are put on hold while your cluster is in upgrade mode.

Finalize upgrade to a bugfix version

To finalize the upgrade:

- 1. Wait for the cluster status to change to **Ready to finalize**. This won't happen until all nodes are active, and running the same upgraded version.
- 2. Select the **Finalize upgrade** button.
- 3. Wait for confirmation that the upgrade is complete. The cluster status will change to **Stable**.

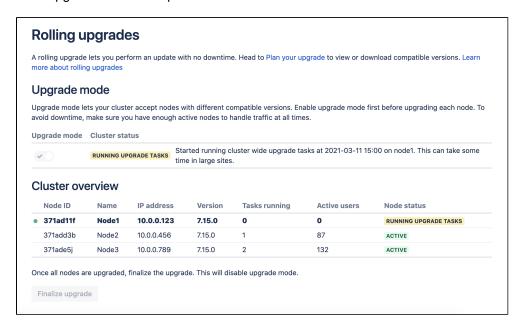
Your upgrade is now complete.

Finalize upgrade to a feature version

To finalize the upgrade:

- 1. Wait for the cluster status to change to **Ready to run upgrade tasks**. This won't happen until all nodes are active, and running the same upgraded version.
- 2. Select the Run upgrade tasks and finalize upgrade button.
- 3. One node will start running upgrade tasks. Tail the logs on this node if you want to monitor the process.
- 4. Wait for confirmation that the upgrade is complete. The cluster status will change to **Stable**.

Your upgrade is now complete.



Screenshot: One cluster node running upgrade tasks for the whole cluster.

Upgrade tasks make any required changes to your database and file system, for example changing the database schema or the way index files are stored in the local home directories.

There are a few things you should know about upgrade tasks:

- One cluster node will run the upgrade tasks on the database and other nodes. If there's a problem, logs will be written to the application log on this node.
- The status of other nodes in the cluster may change to Running upgrade tasks momentarily to
 indicate that an upgrade task is making a change to the file system on that node. The node actually
 running the upgrade tasks does not change.
- Depending on the the size or complexity of your data, some upgrade tasks can take several hours to complete. We generally include a warning in the upgrade notes for the particular version if an upgrade task is likely to take a significant amount of time.

• It's not necessary to direct traffic away from the node running upgrade tasks, but if you know the upgrade tasks are likely to be significant, you may want to do this to avoid any performance impact.

Step 7: Scale down your cluster

In Step 3, we added a node temporarily to the cluster as a replacement for each one we terminated. This was to help ensure we'd have enough nodes to handle normal user traffic. After finalizing the upgrade, you can remove that node:

- In the AWS console, go to Services > CloudFormation. Select your deployments stack to view its Stack Details.
- 2. In the Stack Details screen, click **Update Stack**.
- 3. From the Select Templatescreen, select Use current template and select Next.
- 4. Decrease the value of the following parameters by 1:
 - Maximum number of cluster nodes
 - Minimum number of cluster nodes
- 5. Select **Next**. Click through the next pages, and then to apply the change using the **Update**button.

You can now remove one node from your cluster without AWS replacing it. To do this:

- Choose the node with the least number of running tasks.
- Shut down Confluence gracefully on the node.
- Terminate the node.

Refer to Step 4 for detailed instructions.

Troubleshooting

Disconnect a node from the cluster through the load balancer

If an error prevents you from terminating a node, try disconnecting the node from the cluster through the load balancer. In the AWS Application Load Balancer, each node is registered as a target so to disconnect a node, you'll have to de-register it. For more information on how to do this, seeTarget groups for your Application Load Balancers andRegistered targets.

Traffic is disproportionately distributed during or after upgrade

Some load balancers might use strategies that send a disproportionate amount of active users to a newly-upgraded node. When this happens, the node might become overloaded, slowing down Confluence for all users logged in to the node.

To address this, you can also temporarily disconnect the node from the cluster. This will force the load balancer to re-distribute active users between all other available nodes. Afterwards, you can add the node again to the cluster.

Node errors during rolling upgrade

If a nodes status transitions to **Error**, it means something went wrong during the upgrade. You cant finish the rolling upgrade if any node has an **Error** status. However, you can still disable Upgrade mode as long as the cluster status is still **Ready to upgrade**.

There are several ways to address this:

- Shut down Confluence gracefully on the node. This should disconnect the node from the cluster, allowing the node to transition to an **Offline** status.
- If you cant shut down Confluence gracefully, shut down the node altogether.

Once all active nodes are upgraded with no nodes in Error, you can finalize the rolling upgrade. You can investigate any problems with the problematic node afterwards and re-connect it to the cluster once you address the error.

Roll back to the original version

How you roll back depends on the upgrade stage you have reached. See Roll back a rolling upgrade for more information.



Mixed status with Upgrade mode disabled

If a node is in an Error state with Upgrade mode disabled, you can't enable Upgrade mode. Fix the problem or remove the node from the cluster to enable Upgrade mode.

Node won't start up

If a node is Offline or Starting for too long, you may have to troubleshoot Confluence on the node directly. See C onfluence Startup Problems Troubleshooting for related information.

Upgrade a Confluence cluster through the API without downtime

This document provides guidance on how to initiate and finalize a rolling upgrade through API calls. This upgrade method is suitable for admins with the skills and automation tools to orchestrate maintenance tasks (like upgrades).

For an overview of rolling upgrades (including planning and preparation information), seeUpgrade Confluence without downtime.

API reference

The entire rolling upgrade process is governed by the following API:

http://<host>:<port>/rest/zdu/cluster/zdu/

ThisAPI has the following calls:

/zdu	Get an overview of the cluster's status.	
/zdu/start	Enable upgrade mode.	
/zdu/state	Get the status of the cluster.	
/zdu/nodes/ {nodeld}	Get an overview of a node's status, including the number of running tasks.	
/zdu/cancel	Disable upgrade mode. You can only use this call if the upgrade progress is not MIXED.	
/zdu/approve	Once all nodes are upgraded, finalize the rolling upgrade. This will automatically disable upgrade mode.	

For detailed information about each API call, seeConfluence REST API Documentation.

Initiating a rolling upgrade

To initiate a rolling upgrade, enable rolling upgrade first. To do this, use:

http://<host>:<port>/rest/zdu/cluster/zdu/start

Upgrade mode allows your cluster to temporarily accept nodes running different Confluence versions. This lets you upgrade a node and let it rejoin the cluster (along with the other non-upgraded nodes). Both upgraded and non-upgraded active nodes work together to keep Confluence available to all users. You can disable upgrade mode as long as you havent upgraded any nodes yet.

Upgrading each node individually

Before you upgrade a node, you'll need to gracefully shut down Confluence on it. To do this, run the stop script corresponding to your operating system and configuration. Learn more about graceful Confluence shutdowns.

For example, if you installed Confluence as a service on Linux, run the following command:

\$ sudo /etc/init.d/confluence stop

After upgrading Confluence on the node, wait for it to transition to an Active status first before upgrading another node.

Node statuses

To get the status of a node, use:

http://<host>:<port>/rest/zdu/cluster/zdu/nodes/<nodeID>

ACTIVE	Confluence is connected to the cluster and running with no errors.	
STARTING	Confluence is still loading, and should transition to Active once finished.	
TERMINATING	Confluence was gracefully shut down, and should transition to Offline once finished.	
OFFLINE	Confluence is not responding on the node. This node will be removed from the cluster completely if it is still offline after Upgrade mode is disabled.	
ERROR	Something went wrong with Confluence on the node.	

Cluster statuses

To get the status of the cluster, use:

http://<host>:<port>/rest/zdu/cluster/zdu/state

STABLE	You can turn on Upgrade mode now.
READY_TO_UPGRADE	Upgrade mode is enabled, but no nodes have been upgraded yet. You can start upgrading your first node now.
MIXED	At least one node is upgraded, but you haven't finished upgrading all nodes yet. Your cluster has nodes running different Confluence versions. You need to upgrade all nodes to the same bug fix version to transition to the next status (READY_TO_RUN_UPGRA DE_TASKS).
READY_TO_RUN_UPGRADE_TASKS	All nodes have node been upgraded. You can now finalize the rolling upgrade: http:// <host>:<port>/rest/zdu/cluster/zdu/approve</port></host>



Enable and disable Upgrade mode

How you roll back depends on the upgrade stage you have reached. See Roll back a rolling upgrade for more information.



Mixed status with Upgrade mode disabled

If a node is in an Error state with Upgrade mode disabled, you can't enable Upgrade mode. Fix the problem or remove the node from the cluster to enable Upgrade mode.

Troubleshooting

Node errors during rolling upgrade

If a nodes status transitions to Error, it means something went wrong during the upgrade. You cant finish the rolling upgrade if any node has an Error status. However, you can still disable Upgrade mode as long as the cluster status is still Ready to upgrade.

There are several ways to address this:

- Shut down Confluence gracefully on the node. This should disconnect the node from the cluster, allowing the node to transition to an Offline status.
- If you cant shut down Confluence gracefully, shut down the node altogether.

Once all active nodes are upgraded with no nodes in Error, you can finalize the rolling upgrade. You can investigate any problems with the problematic node afterwards and re-connect it to the cluster once you address the error.

Disconnecting a node from the cluster through the load balancer

If a node error prevents you from gracefully shutting down Confluence, try disconnecting it from the cluster through the load balancer. The following table provides guidance how to do so for popular load balancers.

NGINX	NGINX defines groups of cluster nodes through the upstream directive. To prevent the load balancer from connecting to a node, delete the node's entry from its corresponding upstream group. Learn more about the upstream directive in the ngx_http_upstream_module module.
HAProxy	With HAProxy, you candisable all traffic to the node by putting it in a maint state: set server <node hostname="" ip="" or=""> state maint</node>
	Learn more about forcing a server's administrative state.
Apache	You can disable a node (or "worker") by setting its activation member attribute to disabled. Learn more about advanced load balancer worker properties in Apache.
Azure Application Gateway	We provide adeployment template for Confluence Data Center on Azure; this template uses the Azure Application Gateway as its load balancer. The Azure Application Gateway defines each node as a target within a backend pool. Use the Edit backend pool interface to remove your node's corresponding entry. Learn more about adding (and removing) targets from a backend pool.

Traffic is disproportionately distributed during or after upgrade

Some load balancers might use strategies that send a disproportionate amount of active users to a newly-upgraded node. When this happens, the node might become overloaded, slowing down Confluence for all users logged in to the node.

To address this, you can also temporarily disconnect the node from the cluster. This will force the load balancer to re-distribute active users between all other available nodes. Afterwards, you can add the node again to the cluster.

Node won't start up

If a node is Offline or Starting for too long, you may have to troubleshoot Confluence on the node directly. See C onfluence Startup Problems Troubleshooting for related information.

Roll back a rolling upgrade

The steps on this page only apply if you have used the rolling upgrade method to upgrade Confluence.

If something goes wrong during a rolling upgrade, you may be able to roll back to the original version.

How you roll back depends on the upgrade stage you have reached, and also how you deploy Confluence. To check the current cluster status go to General Configuration> Rolling upgrades.

On this page:

- Roll back a rolling upgrade - manual and Azure deployments
- Roll back a rolling upgrade - AWS deployments

Upgrade mode	Cluster status	
~)	MIXED Your cluster now has nodes running different Confluence versions.	

Screenshot: Cluster status in the Rolling upgrades screen

Cluster upgrade status	Action to take
READY TO UPGRADE	In the rolling upgrades screen, disable Upgrade mode .
MIXED	Follow the steps below to roll back the application version on each upgraded node:
READY TO RUN UPGRADE TASKS	 Rollback steps - manual and Azure deployments Rollback steps - AWS deployments Rollback steps - Kubernetes deployments Once all nodes are running the original version, the cluster status will change to Ready to upgrade.
RUNNING UPGRADE TASKS UPGRADE TASKS FAILED	You can't roll back once final cluster-wide upgrade tasks have started running. If you stop Confluence on the node running the upgrade task, another node will pick up where the stopped node left off. If upgrade tasks fail, you will need to investigate the problem, then re-run upgrade tasks
	from the rolling upgrades screen.
COMPLETE	You can't roll back to an earlier version, because the upgrade was finalized successfully.

Roll back a rolling upgrade - manual and Azure deployments

These instructions assume your original install directory is still available. If it's not, you may need to restore it from your backup. You don't need to restore the local home directory.

To roll back an upgraded node to its original version:

- 1. Access the node through a command line or SSH.
- 2. Shut down Confluence gracefully on the node.
- 3. Wait for the node to go offline. You can monitor its status on the Node status column of the Rolling upgrade pages Cluster overview section.
- 4. If you run Confluence as a service:

- On Windows, delete the new service then re-install the old serviceby running<old-install-directory>/bin/service.bat.
- On Linux, update the service to point to the old installation directory (or use symbolic links to do this).
- 5. Start Confluence (from the original install directory) on the node. You should not see the setup wizard.

Once all nodes are running the same version, the clusters status will revert back to Ready to upgrade. You can then turn off Upgrade mode.

Roll back a rolling upgrade - AWS deployments

To roll back the upgraded nodes to the original version:

In the AWS console, go to Services > CloudFormation. Select your deployments stack to view its Stack Details.

- 1. In the Stack Details screen, selectUpdate Stack.
- 2. From the Select Templatescreen, select Use current template and select Next.
- 3. Set the Version parameter to your original version.
- 4. SelectNext. Click through the next pages, and then to apply the change using the Update button.

Afterwards, terminate all nodes running the new version of Confluence. AWS will replace each with a node running the original version. Once all nodes are running the same version, the clusters status will revert back to Ready to upgrade. This will also allow you to disable Upgrade mode.

Upgrade task troubleshooting

When introducing a new feature, or making a significant change to your application, we sometimes need to transform existing data in your database or index, or change the way some data is stored.

Here's a simple example. If we stored "Sydney Australia" in a location column in the database, but later decide to store city and country information separately, we might use an upgrade task to take the existing data in the location column, and split it into a city and country column, containing "Sydney" and "Australia" respectively. Actual upgrade tasks are rarely this simplistic, but you get the idea.

On this page:

- When are upgrade tasks run?
- Troubleshooting failed upgrade tasks
- Known issues

We don't include changes that require upgrade tasks in bug fix releases, but they can be quite common in feature and platform releases. You can tell if a version has an upgrade task if the build number is different to your current version.

When are upgrade tasks run?

This depends on the type of upgrade task, and whether you are upgrading with or without downtime.

Rolling upgrade without downtime

If you're performing a rolling upgrade:

- node-specific upgrade tasks happen just prior to the application starting up on that node.
- cluster-wide upgrade tasks happen when all nodes are running the new version, and you select Run
 upgrade tasks and finalize upgrade in the rolling upgrades screen.

During a rolling upgrade, there's a short time where data might exist in both old and new formats. Clusterwide upgrade tasks include tasks that transform data in the database, and may also include changes to the shared home and local home directories on each node. These tasks require all nodes to have been upgraded before they can be run.

Using the example above, a node that has not yet been upgraded would continue to write to the location column, while an upgraded node would write to the new city and country columns, as well as the old loc ation column (to prevent data loss if you need to roll back). Once all nodes are upgraded, it's safe for us to split the existing data in the location field into the new city and country fields. This is often the part of the upgrade task that can take some time, depending on how much data you have.

Upgrade with downtime

When upgrading a non-clustered deployment, upgrade tasks are usually run just prior to the application starting up after the upgrade.

When upgrading a cluster with downtime (not a rolling upgrade), cluster-wide upgrade tasks are run when the first node starts up.

Troubleshooting failed upgrade tasks

If an upgrade task fails, there are a number of things you'll need to do to resolve the issue.

Check the application logs

The first step is to check the application logs. If you're running Confluence in a cluster, you may need to check the logs on more than one node.

Sometimes the cause will be obvious, such as a network timeout, not enough disk space in the local or shared home directory, or the database user / confluence user has inadequate permissions to complete the action.

Check your database configuration

The most common reason upgrade tasks fail include:

- database user does not have adequate permissions to perform the required action
- database configuration is incorrect (for example the character set or encoding)
- database version or edition not supported.

This usually results in a database error being written to the application logs. Check for KB articles about the specific problem, and confirm your setup matches the database setup specified in our documentation.

Re-run upgrade tasks

Once you've resolved any issues, you'll need to re-run the upgrade tasks. How you do this depends on whether you are upgrading with or without downtime.

- If you're performing a rolling upgrade, re-start the application on any failed nodes, then select Re-run upgrade tasks and finalize upgrade.
- If you're upgrading with downtime, re-start the application. Upgrade tasks will run prior to the application starting up.

Don't leave your application in upgrade mode

If you're performing a rolling upgrade, it's important you don't leave your cluster in an upgrade mode longer than is necessary. This is because there may be data that needs to be handled in multiple ways until the final upgrade tasks can be run.

Known issues

Some non-enterprise editions of Microsoft SQL Server don't support online index creation. If an
upgrade task needs to acquire an exclusive table lock, you may experience some performance
degradation or downtime. We'll warn you if we detect that you database edition may be affected.

Supported Platforms

This page describes the additional software and infrastructure you'll need to run Confluence. Please review this info before installing Confluence. The information on this page applies to **Confluence Server and Data Center 7.16**.

- You should only use Confluence with a supported platform. Any platforms and versions not listed on this page are unsupported, which means we don't test, fix bugs or provide assistance.
- SeeEnd of Support Announcements for Confluencefor upcoming changes to supported platforms.
- Go to Seneral Configuration > Troubleshooting and support tools to check your instance health. It looks at things like your license validity, Tomcat version, basic database setup and more.

Definitions:

- Supported you can use Confluence Server and Data Center 7.16 with this platform.
- Limited you can evaluate Confluence on this platform, but you can't use it to run a production Confluence site.

⚠Deprecated - support for this platform will end in an upcoming release. SeeEnd of Support Announcements for Confluence.

On this page:

- Browsers
- Operating systems
- Databases
- Java
- Infrastructure

Related pages:

- Confluence Installation Guide
- Confluence Setup Guide
- Server Hardware Requirements Guide
- Supported Platforms FAQ

Browsers

Desktop browsers

- ✓Microsoft Edge (Chromium)
- Chrome
- Firefox
- Safari (Mac only)

Mobile browsers

- Chrome
- Firefox
- Safari (iOS only)
- Android WebView

Mobile operating system (required for mobile app)

⊘iOS 11 or later

◆Android 4.4 (KitKat) or later

Operating systems

Operating systems

Microsoft Windows

✓Linux (most distributions)

Good to know:

- The Confluence setup wizard requires Javascript to be enabled while installing Confluence.Learn more
- Parts of Confluence won't display correctly if your browser window size is less than 1024x768.
- Although some supported browsers may allow you to enable Adobe Flash in their advanced settings, we recommend leaving Flash disabled, as enabling it may expose you to security vulnerabilities.

Known issues:

- The following operating system variants can't be used with Confluence:
 - Windows Nano

MacOS / OSX (evaluation only)

Alpine Linux (3.5 and earlier)

Good to know:

- You can run Confluence on 32bit or 64bit operating systems, but we only provide installers for 64bit operating systems.
- You can evaluate Confluence on MacOS / OSX, but you can't install and run your production Confluence site on a Mac.

Databases

PostgreSQL

- ♣PostgreSQL 10
- ✓PostgreSQL 11

Amazon Aurora (Data Center only)

- ♣PostgreSQL 10
- PostgreSQL 11

Azure PostgreSQL (Data Center only)

- ♠PostgreSQL 10
- ✓PostgreSQL 11

MySQL

- ✓MySQL 5.7

Oracle

- ✓Oracle 12c Release 2
- ✓Oracle 19c

Microsoft SQL Server

- **▲**SQL Server 2016
- SQL Server 2017
- SQL Server 2019
- Azure SQL

Embedded database

•H2 (Data Center testing installations only)

Java

Oracle JRE / JDK

- Java 1.8
- Java 11

Known issues:

- Confluence will not work on MySQL variants such as:
 - MariaDB see CONFSERVER-29060
 - Percona Server see CONFSERVER-36471
- Confluence can become unresponsive with Oracle's Native Network Encryption - seeCONFSERVER-60152for mitigation options.
- Confluence requires Service Pack 1 or later for Microsoft SQL Server 2016 - seeCONFSERVER-61145.

Good to know:

- The embedded H2 database isonlysupported fortesting and app development purposesonnon-clustered (single node)Confluence Data Center installations.
- You can use Amazon's Relational Database Service (RDS) for the supported databases listed on this page.
- The only supported Amazon Aurora config is a PostgreSQLcompatible clustered database with one writer replicating to zero or more readers. Learn more

Known issues:

• There's a known issue with some Java 11 versions and TLS 1.3. We recommend Java 11.0.8 or later.

AdoptOpenJDK

- Java 8 (HotSpot)
- Java 11 (HotSpot)
- Some Oracle Java versions have bugs that impact Confluence. You can't run Confluence on Java 1.8.0_25 and 1.8.0_31 (seeJDK-8059299) or Java 1.8.0_45 (seeJDK-8068400).
- Some AdoptOpenJDK versions have bugs that impact Confluence.
 We recommend versionjdk8u202 (seeCONFSERVER-58784).
- You can't run Confluence on the OPENJ9 JVM from AdoptOpenJDK. Use the HotSpot JVM.
- We useAdoptOpenJDKto replicate issues raised with OpenJDK. If youre using a different distribution of OpenJDK well still provide support for our products. However, if the bug is caused by a problem in Java distribution, you'll need to contact the Java distributor for help.

Good to know:

- You don't need to install Java if you plan to use the installer to install Confluence, as a Java 11 JRE is bundled with Confluence (provided by AdoptOpenJDK).
- SeeBundled Tomcat and Java versions to see which Java version was bundled with your Confluence version.

Infrastructure

Hardware:

- You can't run Confluence on SPARC based hardware. You'll need to use x86 hardware or 64bit derivatives of x86 hardware.
- You can't use an NFS mount for your installation or home directory due to Lucene requirements. If you're installing Confluence Data Center, an NFS mount is fine for the shared home directory, but not for the local home directories.

Virtualization:

- You can run Confluence and Confluence Data Center in a virtualized environment (including Docker), but our support team can't assist you with problems related to the environment itself. SeeRunning Confluence in a Virtualized Environment
- Our support team can assist you with deploying Confluence Data Center in AWS using the Cloud Formation Template or Quick Start. We won't be able to assist you if you have significantly customised the Cloud Formation Template.

Application server:

We only support the Tomcat version that is bundled with your Confluence version. You can't run
Confluence in your own application server. SeeBundled Tomcat and Java versions to see which
version of Tomcat was bundled with your Confluence version.

Internet protocols:

- You can run Confluence in both IPv4 and IPv6 environments.
- Raw IPv6 addresses are not always recognized. See the Confluence 6.9 Upgrade Notesfor limitations and known issues.

Operating system support:

 You should only install and use Confluence on operating system versions that have active vendor support. For example, you can use Confluence on any Microsoft supported version of Windows, unless specified otherwise above.

For more information see our Server Hardware Requirements Guide and System Requirements.

End of Support Announcements for Confluence

This page is where we announce end of support for various platforms, browsers, and information on features that will be discontinued in Confluence Server.

The table below summarizes the end of support announcements for **upcoming**Confluence releases. If a platform (or version) has already reached its end of support date, it is **not** listed in the table.

Platform	Confluence end of support
PostgreSQL 10	No support in Confluence 7.18 (announcement)
Microsoft SQL Server 2016	No support in Confluence 7.17(announcement)
PostgreSQL 9.6	No support in Confluence 7.14 (announcement)
Microsoft Edge 18	No support in Confluence 7.13 (announcement)
SQL Server 2014 PostgreSQL 9.5	No support in Confluence 7.5 (announcement)
Internet Explorer 11	No support in Confluence 7.5 (announcement)
View File macros	These macros are now fully supported (updated announcement)

Most recent announcements first:

- Deprecated database for Confluence (20 December 2021)
- Deprecated database for Confluence (16 November 2021)
- Deprecated database for Confluence (13 April 2021)
- Deprecated browsers for Confluence (2 February 2021)
- Deprecated databases for Confluence (11 December 2019)
- Deprecated database for Confluence (11 December 2019)
- Deprecated databases for Confluence (14 October 2019)
- Deprecated browsers for Confluence (24 September 2019)
- Deprecated macros for Confluence (12 March 2019)
- Deprecated Gadgets in Confluence (12 March 2019)
- Changes to features in Confluence (12 March 2019)
- Deprecated databases for Confluence (2 October 2018)
- Deprecated databases for Confluence (30 January 2017)
- Deprecated macro for Confluence (31 October 2017)
- Deprecated driver for Microsoft SQL Server
- Deprecated operating system for Confluence (15 May 2017)
- Deprecated mobile browser for Confluence (3 November 2016)
- Changes to Confluence distributions (8 June 2016)
- Deprecated browsers for Confluence (8 June 2016)
- Deprecated databases for Confluence (8 June 2016)
- Deprecated macros for Confluence (13 November 2015)
- Discontinued features for Confluence (10 July 2015)
- Deprecated databases for Confluence (19 May 2015)
- DeprecatedTomcatplatform for Confluence (1 May 2015)
- Deprecated Web Browsers for Confluence (20 April 2015)
- Deprecated Java platform for Confluence (27 January 2015)
- Deprecated distribution for Confluence (2 September 2014)
- Deprecated databases for Confluence (12 June 2014)
- DeprecatedTomcatplatform 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 database for Confluence (20 December 2021)

Atlassian will not support the following database in Confluence 7.18:

PostgreSQL 10

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 7.18 or later with this database.

- Confluence 7.17 is the last version that will support PostgreSQL 10.
- Confluence 7.17 and earlier versions will continue to work with PostgreSQL 10, however we will not fix bugs affecting this database after theend-of-life datefor your version of Confluence.
- Confluence 7.18 and later will not be tested with PostgreSQL 10.

Check out the Supported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated database for Confluence (16 November 2021)

Atlassian will not support the following database in Confluence 7.17:

Microsoft SQL Server 2016

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 7.17 or later with this database.

- Confluence 7.16 is the last version that will support Microsoft SQL Server 2016.
- Confluence 7.16 and earlier versions will continue to work with Microsoft SQL Server 2016, however we
 will not fix bugs affecting this database after theend-of-life datefor your version of Confluence.
- Confluence 7.17 and later will not be tested with Microsoft SQL Server 2016.

Check out the Supported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated database for Confluence (13 April 2021)

Atlassian will not support the following database in Confluence 7.14:

PostgreSQL 9.6

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 7.14 or later with this database.

- Confluence 7.13 is the last version that will supportPostgreSQL 9.6.
- Confluence 7.13 and earlier versions will continue to work with PostgreSQL 9.6, however we will not fix bugs affecting this database after theend-of-life datefor your version of Confluence.
- Confluence 7.14 and later will not be tested with PostgreSQL 9.6.

Check out the Supported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated browsers for Confluence (2 February 2021)

In January 2020 Microsoft released a new Microsoft Edge browser based on Chromium. This new version is compatible with all supported Windows versions, and replaces the previous version, now known as Microsoft Edge Legacy. Read more about the difference between the new Microsoft Edge and Microsoft Edge Legacy on the Microsoft support site.

As Microsoft have announced plans to end support for Microsoft Edge Legacy, we have also decided to end support for Microsoft Edge Legacy.

End of support means we will not fix bugs specific to Microsoft Edge Legacy, and will begin to introduce features that aren't compatible with this browser.

When is this happening?

- Confluence 7.12 is the last version to support Microsoft Edge Legacy.
- Confluence 7.13 and subsequent versions will not support Microsoft Edge Legacy.

What this means for you

We recommend switching to one of oursupported browsers, such as the new Microsoft Edge (Chromium), Google Chrome, or Mozilla Firefox.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated databases for Confluence (11 December 2019)

Atlassian will not support the following databases in Confluence 7.5:

- Microsoft SQL Server 2014
- PostgreSQL 9.5

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 7.5 or later with this database.

- Confluence 7.4 is the last version that will support these databases.
- Confluence 7.4 and earlier versions will continue to work withthese databases, however we will not fix bugs affecting these databases after theend-of-life datefor your version of Confluence.
- Confluence 7.5 and later will not be tested withthese databases.

Check out the Supported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated database for Confluence (11 December 2019)

Atlassian will not support the following database in Confluence 7.4:

Microsoft SQL Server 2012

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 7.4 or later with this database.

- Confluence 7.3 is the last version that will support Microsoft SQL Server 2012.
- Confluence 7.3 and earlier versions will continue to work with Microsoft SQL Server 2012, however we
 will not fix bugs affecting this database after theend-of-life datefor your version of Confluence.
- Confluence 7.4 and later will not be tested with Microsoft SQL Server 2012.

Check out the Supported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated databases for Confluence (14 October 2019)

Atlassian will not support the following databases in Confluence 7.4:

- PostgreSQL 9.4
- MySQL 5.6
- Oracle 12c R1

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 7.4 or later with these databases.

- Confluence 7.3 is the last version that will support these databases.
- Confluence 7.3 and earlier versions will continue to work with these databases, however we will not fix bugs affecting these databases after theend-of-life datefor your version of Confluence.
- Confluence 7.4 and later will not be tested with these databases.

Check out the Supported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated browsers for Confluence (24 September 2019)

In 2015 Microsoft released Edge as the browser to supersede Internet Explorer, and in recent times Microsoft has discouraged the use of Internet Explorer as a default browser. To allow us to continue to take advantage of modern web standards to deliver improved functionality and the best possible user experience across all of our products, we have decided to end support for Internet Explorer 11.

End of support means we will not fix bugs specific to Internet Explorer 11, and will begin to introduce features that aren't compatible with this browser.

When is this happening?

- Confluence 7.4 is the last version to support Internet Explorer. Confluence 7.4 will be an Enterprise release.
- Confluence 7.5 and subsequent versions will not support Internet Explorer 11.

What this means for you

We recommend switching to one of our supported browsers, such as Microsoft Edge, Google Chrome, or Mozilla Firefox.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated macros for Confluence (12 March 2019)

We will end support for the following macros in Confluence 7.0, and hide them from the macro browser. Any existing instances of these macros will still work, but you won't be able to insert these macros into the editor using the macro browser:

- IM Presence macro
- Netwok macro
- Search results macro
- Space details macro

End of support means Atlassian will not fix bugs related to these macros in Confluence 7.0 or later versions. We will remove these macro entirely in a future Confluence release, and will provide more information at that time.

To check whether a macro is used in your site, go to O-Seneral Configuration > Macro Usage. Some macros will be listed under the system app that provides them.

IM Presence macro

TheIM Presence Macroshows when a given user is online in a selected chat service. The macro only supports a small number of chat services, and we feel that most modern chat tools provide better ways to see this information.

If you have questions or concerns, please comment on this issue CONFSERVER 57596 CLOSED



Network macro

The Network Macroallows you display the people a particular user is following, or people who are following that user. Following someone is a useful way to get notifications about their activity, and this network information is also available on each user's profile page.

If you have questions or concerns, please comment on this issue CONFSERVER-57597 CLOSED



Search results macro

The Search Results Macroallows you to display the results of a keyword search on a page. We are making some great changes to Search over the next few releases, and have observed that this macro is rarely used.

If you have questions or concerns, please comment on this issue CONFSERVER-57598 CLOSED



Space details macro

The Space Details Macroallows you to display basic information about the current space on a page. This information is available at all times from Space Tools > Overview.

If you have questions or concerns, please comment on this issue CONFSERVER-57599 CLOSED



Deprecated Gadgets in Confluence (12 March 2019)

We will end support for the following Gadgets in Confluence 7.0, and hide them from the macro browser. Any existing instances of these gadgets will still work, but you won't be able to insert these gadgets into the editor using the macro browser:

- Activity Stream
- Confluence Page Gadget
- Confluence Quick Nav Gadget
- News gadget

End of support means Atlassian will not fix bugs related to these gadgets in Confluence 7.0 or later versions. We will remove these gadgets entirely in a future Confluence release, and will provide more information at that time.

Gadgetswere designed to allow you to display information dynamically from sources like iGoogle or Jira, for example, in Confluence. The first gadgets were introduced in Confluence 3.1, and much of the technology they were based on is now superseded or obsolete. Since then we have also implemented a number of better ways to display dynamic information using macros and other integration points.

Activity Stream gadget

The Activity stream gadget shows a list of recently changed content in your site. We recommend using the Recently Updated macro as an alternative in Confluence.

Confluence Page Gadget

This gadget displays the contents of a Confluence page. We recommend using the Include Page macro as an alternative in Confluence.

Confluence Quick Nav Gadget

This gadget provides a search field that can be used to search for page titles in Confluence. We recommend using the Livesearch macro as an alternative in Confluence.

News gadget

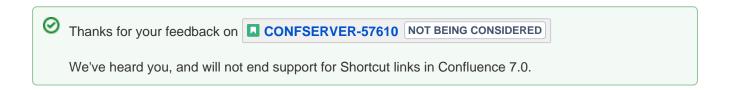
This gadget previously displayed blogs and other news from Atlassian. It has not been displaying content for some time. We will remove this gadget completley in 7.0.

If you have questions or concerns, please comment on this issue CONFSERVER 57614 CLOSED

Changes to features in Confluence (12 March 2019)

Shortcut links

Shortcut Linkswere introduced in Confluence 2.3 and provided a quick way to add links to websites in wiki markup. Shortcut links can only be configured by an administrator, are not easily discoverable, and seldom used by end users.



Trackback and external referrers

We will remove the trackback and referrers features completley in Confluence 7.0.

Trackbackenables Confluence to send and receive trackback pings when pages are linked to.External Referrers appear on the Page Information view of a page, and list clicks from external websites to the page. Trackback is no longer widely used in modern websites, and because it relies on accepting unauthenticated requests to a particular URL, is a spam vector.

If you have questions or concerns, please comment on this issue CONFSERVER-57611 CLOSED

Orphaned pages screen

We will remove the Orphaned pages screen in the default theme in Confluence 7.0.

The Orphaned pages screen provided a list of all pages that Confluence considersorphaned pages(not a child of a space homepage, and not linked to by any other page). Since the introduction of the Confluence 5 default theme, the orphaned pages screen has been less useful because it's always possible to see all pages in a space via**Space Tools>Reorder pages**.

If you have questions or concerns, please comment on this issue CONFSERVER 57601 CLOSED

Hipchat integration

We have discontinued development on all chat products. Hipchat Cloud services were shut down in February 2019, and Hipchat Data Center and Server will both reach end of life within the next year.

We will end support for all bundled Hipchat plugins in Confluence 7.0. These will be disabled by default for new installations. This will have no impact on existing installations, and can be easily enabled if required.

End of support means that Atlassian will not fix bugs related to Hipchat integration in Confluence 7.0 or later versions.

If you have questions or concerns, please comment on this issue CONFSERVER 57602 CLOSED

Deprecated databases for Confluence (2 October 2018)

Atlassian will end support for **PostgreSQL 9.3** in Confluence 6.13. End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 6.13 or later with this database.

- Confluence 6.12 is the last version that will support PostgreSQL 9.3.
- Confluence 6.12 and earlier versions will continue to work with PostgreSQL 9.3, however we will not fix bugs affecting these databases after the end-of-life datefor your version of Confluence.
- Confluence 6.13 and later will not be tested with PostgreSQL 9.3.

Check out the Supported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated databases for Confluence (30 January 2017)

Atlassian will end support for **PostgreSQL 9.2** in Confluence 6.8.End of support means thatAtlassian will not offer support for, or fix bugs related to, running Confluence 6.8 or later with this database.

- Confluence 6.7 is the last version that will support PostgreSQL 9.2.
- Confluence 6.7 and earlier versions will continue to work with PostgreSQL 9.2, however we will not fix bugs affecting these databases after the end-of-life datefor your version of Confluence.
- Confluence 6.8 and later will not be tested with PostgreSQL 9.2.

Check out the Supported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated macro for Confluence (31 October 2017)

We will end support for the **JUnit Report macro** with the release of Confluence 6.6. This macro is used to display the results of JUnit tests on a Confluence page and, based on our research, is rarely used.

End of support means that Atlassian will not fix bugs related to this macro past the support end date for your version of Confluence. We will remove this macro entirely in a future Confluence release, and will provide more information at that time.

To check whether this macro is used in your site, go to General Configuration Macro Usage. The JUnit Report macro will be listed under Advanced Macros if it's used.

If you have questions or concerns, please comment on this issue

CONFSERVER-53942 - Plans to remove JUnit macro CLOSED

Deprecated driver for Microsoft SQL Server

We are replacing the open source jTDS driver for SQL Server with the official Microsoft JDBC Driver for SQL Server. This new driver is bundled with Confluence 6.4 and later.

Atlassian will end support for the jTDS driver in Confluence 6.6. End of support means that Atlassian will not offer support for, or fix bugs related to, installing and running Confluence 6.6 or later with this driver.

- Confluence 6.5.x will be the last major release to bundle the jTDS driver.
- Confluence 6.5.x and earlier versions will continue to be supported with the jTDS driver, until their support end date.
- Confluence 6.6.x will not bundle or support the jTDS driver. We'll provide plenty of information on how to migrate to the new driver at that time.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated operating system for Confluence (15 May 2017)

Atlassian will end support for the Oracle Solaris operating systemin Confluence 6.3.End of support means thatAtlassian will not offer support for, or fix bugs related to, installing and running Confluence 6.3 or later on this operating system.

- Confluence 6.2.x will be the last major release that can be installed on Solaris.
- Confluence 6.2.x and earlier versions will continue to be supported on Solaris, until their support end date.

Check out the Supported Platformspage for the full list of supported operating systems.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated mobile browser for Confluence (3 November 2016)

Atlassian will end support for the default browser provided with Android4.0.3 (Ice Cream Sandwich)in Confluence 6.0.End of support means that Atlassian will not fix bugs related to this browser past the support end date, except for security related issues. This means:

- Confluence 5.10 will be the last major release that supports the default browser provided with Android4.
 0.3 (Ice Cream Sandwich).
- Confluence 5.10.x and earlier versions will continue to work on the default browser provided with Android4.0.3 (Ice Cream Sandwich).

With the release of Confluence 6.0 we have added support for the default browser provided with current Android versions from 4.4 (KitKat) and later. Check out the Supported Platformspage for the full list of supported browsers.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Changes to Confluence distributions (8 June 2016)

To help us bring you new Confluence Server releases faster, we are considering only providing 64-bit installers. Confluence 5.10 would be the last Confluence release to provide a 32-bit installer.

Q: Can I upgrade using the 64-bit installer?

Yes. If you installed Confluence using the 32-bit installer on a 64-bit operating system, you will be able to upgrade using the 64-bit installer.

Q: What if I am not able to use the 64-bit installer?

We'd love to hear from you to better understand how this change would impact you. Comment on this issue **CONFSERVER-42817** - Planned deprecation of 32-bit installers **CLOSED** or contact us directly ateolannouncement at atlassian dot com.

Deprecated browsers for Confluence (8 June 2016)

Atlassian will end support for Internet Explorer 10 in Confluence 6.0.End of support means that Atlassian will not fix bugs related to Internet Explorer 10 past the support end date, except for security related issues.

This changeallows us to use more modern browser technologies to give you the best user experience in Confluence. Check out the Supported Platformspage for the full list of supported browsers.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Internet Explorer 10 (IE10) end of support notes

- Confluence 5.10 will be the last major release that supportsInternet Explorer 10.
- Confluence 5.10.x and earlier versions will continue to work on Internet Explorer 10.
- No Confluence releases after 5.10.x will be tested with Internet Explorer 10.

Deprecated databases for Confluence (8 June 2016)

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 emaileolannouncement at atlassian dot com.

End of life announcement for database support

Database	Support End Date
MySQL 5.5	After Confluence 5.10.x

Notes:

Confluence 5.10 is the last version that will support MySQL 5.5.

- Confluence 5.10 and previously-released versions will continue to work with the database version listed above, however we will not fix bugs affecting these databases after the end-of-life datefor your version of Confluence.
- No Confluence releases after 5.10.x will be tested with the database listed above.

Deprecated macros for Confluence (13 November 2015)



Update 22 January 2019

We know from your feedback that the existing View File macros provide important functionality that the newer file upload and preview experience does not. For this reason, we've decided to reverse the decision to stop supporting these macros.

This means from Confluence 6.14, we will fix bugs relating to these macros, and will not remove the macros from Confluence.

With the release of Confluence 5.9 we will be ending support for the following macros, known collectively as the 'View File' macros:

- Office Excel
- Office Word
- Office PowerPoint
- PDF

End of support means that Atlassian will not fix bugs related to these macros past the support end date for your version of Confluence. We plan to remove these macros in a future Confluence release, and will provide plenty of information to help you make the transition when the time comes.

The View File macros will still be available in future Confluence releases (including Confluence 5.9, 5.10 and later), but we recommend inserting Office and PDF files as a thumbnail or link, and using the preview to view the file in full, as it provides a much better way to display Office and PDF files on your pages. See Display Files and Images for more info.

If you have any questions or concerns, please comment on this issue

CONFSERVER-39829 - Plans to remove the view file macros CLOSED

Discontinued features for Confluence (10 July 2015)

Status updates

As part of our work to make Confluence simpler and easier to use we've decided to remove the Status Updates feature in **Confluence 5.9**. This includes the ability to:

- update your status
- see other people's status via their profile or the User Status List macro.

Our research tells us that this feature isn't widely used, and we believe that HipChatgives your team much better ways to share their status.

We'll provide more information at the time of the Confluence 5.9 release. If you have questions or concerns, please comment on this issue CONFSERVER 38253 - Plans to remove status updates CLOSED

Documentation theme

In order to better focus our development efforts on a single theme, we plan to remove the Documentation theme inConfluence 6.0.

We know that many customers use the Documentation theme because they like to have a page tree in their space sidebar. This has been available in the default theme for some time now, plus other great features like sidebar shortcuts, JIRA links, and sticky table headers.

To help you switch to the more modern default theme, we've added some of your favorite documentation theme features, including the ability to add:

- a header and footer
- · custom content to the sidebar.

These new additions to the default theme are available in Confluence 5.9. As these fields will continue to use wiki markup, you will be able to drop your existing wiki markup straight from the Documentation theme into the default theme.

To help you switch themes we've put together a FAQ and step-by-step guidewhich covers everything from how to turn on the default theme, find out which spaces are using the theme, and what to do if the Documentation theme is the global theme for your whole site.

If you have any questions or concerns please comment on this issue

CONFSERVER-38256 - Plans to remove the documentation theme CLOSED

Deprecated databases for Confluence (19 May 2015)

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.

End of Life Announcement for Database Support

Database	Support End Date
Microsoft SQL 2008	
Oracle 11.1	After Confluence 5.8.x
Oracle 11.2	

Notes:

- Confluence 5.8 is the last version that will support the database versions listed above.
- Confluence 5.8 and previously-released versions will continue to work with the database versions listed above, however we will not fix bugs affecting these databases after the end-of-life date for your version of Confluence.
- No Confluence releases after 5.8.x will be tested with the databases listed above.

DeprecatedTomcatplatform for Confluence (1 May 2015)

This section announces the end of Atlassian support for Tomcat 7.0.x for Confluence. Aspreviously announced, we now only support the version of Tomcat that is bundled with your version of 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. If you have questions or concerns regarding this announcement, please emaileol-announcement at atlassian dot com.

End of Life Announcement for Tomcat 7.0.x Support

Platform	Support End Date
Tomcat 7.0.x	When Confluence 5.8 is released

Tomcat 7.0.x notes:

- Confluence 5.7 is the last major version that will support Tomcat 7.0.x. The Confluence 5.7.x bug-fix releases will also continue to support Tomcat 7.0.x.
- Confluence 5.7.x and previously-released versions will continue to work with Tomcat 7.0.x. However, we
 will not fix bugs affectingTomcat 7.0.x after the end-of-life date for your version of Confluence.
- Confluence 5.8 will not be tested with Tomcat 7.0.x.

Deprecated Web Browsers for Confluence (20 April 2015)

Atlassian will end support for Internet Explorer 9 in the next major release after Confluence 5.8.x.End of support means thatAtlassian will not fix bugs related to Internet Explorer 9 past the support end date, except for security related issues.

This change allows us to use more modern browser technologies to give you the best user experience in Confluence. Check out the Supported Platformspage for the full list of supported browsers.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Internet Explorer 9 (IE9) End of Support Notes

- Confluence 5.8 will be the last major release that supportsInternet Explorer 9
- Confluence 5.8.x and earlier versions will continue to work onInternet Explorer 9
- No Confluence releases after 5.8.x will be tested with Internet Explorer 9

Deprecated Java platform for Confluence (27 January 2015)

This section announces the end of Atlassian support for Java 7 for Confluence. Please note that Oracle is planning to stop providing public updates for JRE 7 in April 2015.

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 at atlassian dot com.

End of Life Announcement for Java 7 Support

Platform	Support End Date
Java 7 (JRE and JDK 1.7)	When Confluence 5.8 is released

Java 7 notes:

- Confluence 5.7 is the last major version that will support Java 7. The Confluence 5.7.x bug-fix releases will also continue to support Java 7.
- Java 7 (JRE and JDK 1.7) will still be supported in Confluence 5.7.
- Confluence 5.7.x and previously-released versions will continue to work with Java 7, but we will not fix bugs affecting Java 7 after the end-of-life date for your version of Confluence.
- Confluence 5.8 will not be tested with Java 7.

Deprecated distribution for Confluence (2 September 2014)

To help us to make Confluence a more robust and scalable application, we have decided to stop providing an EAR/WAR distribution. This means that the only supported application server will be will be the version of Tomcat that is bundled with each release.

Confluence 5.6 will be the last Confluence release to provide an EAR/WAR edition.

Q: Do I need to use the installer?

No, the removal of the EAR/WAR distribution does not force you to use the installer. You can still use the standalone distribution, which doesn't have an install script - it's just a copy of Tomcat with Confluence configured inside it. Essentially it's a directory that you unpack and then run yourself.

Q: What if a security problem is found in the bundled version of Tomcat?

Our security team monitors vulnerabilities in all our dependencies, including Tomcat, and fixes continue to follow ourSecurity Bugfix Policy. If at any time you become aware of a vulnerability we've missed, please report it as described inHow to report a security issue.

If you have more questions or concerns regarding this announcement, please contact us ateolannouncement at atlassian dot com.

Deprecated databases for Confluence (12 June 2014)

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.

End of Life Announcement for Database Support

Database	Support End Date
PostgreSQL 8.4	
PostgreSQL 9.0	
PostgreSQL 9.1	With the release of Confluence 5.7
MySQL 5.1	

Notes:

- Confluence 5.6 is the last version that will support the database versions listed above.
- Confluence 5.6 and previously-released versions will continue to work with the database versions listed above, however we will not fix bugs affecting these databases after the end-of-life date for your version of Confluence.
- Confluence 5.7 has not been tested with the databases listed above.

DeprecatedTomcatplatform 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.

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 affectingTomcat 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.

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 utilize 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 minimize 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 supportInternet Explorer 9 (IE9) and Internet Explorer 10 (IE10) as well asthe latest versions of Chrome, Firefox and Safari. For further information, please refer to the Supported Platformspage. If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot 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 at atlassian dot 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 at atlassian dot 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 affectingTomcat 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 at atlassian dot 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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 forseeable 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot com.

Why is Atlassian doing this?

We have chosen to standardize 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 specialize in providing enterprise grade support contracts for it, ranging from customizations 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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 eol-announcement at atlassian dot 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.

Bundled Tomcat and Java versions

This page lists the specific versions of Apache Tomcat and Adopt OpenJDK that we bundle with Confluence. This information is useful if you want to check whether your Confluence version might be using a Tomcat or Java version that's affected by a specific issue, vulnerability, or bug.

We also list the specific Java versions we use when testing Confluence, which can be handy if you don't run Confluence with the bundled JRE.

Confluence 7.16

Confluence	Tomcat	Bundled JRE	Tested JREs
7.16.0	9.0.54	AdoptiumOpenJDK 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08

Confluence 7.15

Confluence	Tomcat	Bundled JRE	Tested JREs
7.15.0	9.0.54	AdoptiumOpenJDK 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08
7.15.1	9.0.54	AdoptiumOpenJDK 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08

Confluence 7.14

Confluence	Tomcat	Bundled JRE	Tested JREs
7.14.0	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08
7.14.1	9.0.54	AdoptiumOpenJDK 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08
7.14.2	9.0.54	AdoptiumOpenJDK 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08

Confluence	Tomcat	Bundled JRE	Tested JREs
7.13.0	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Adoptium OpenJDK8u292b10
7.13.1	9.0.45	Adopt OpenJDK11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08
7.13.2	9.0.45	AdoptiumOpenJDK 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08

7.13.3	9.0.45	AdoptiumOpenJDK 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08
--------	--------	---------------------------	--

Confluence 7.12

Confluence	Tomcat	Bundled JRE	Tested JREs
7.12.0	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.12.1	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.12.2	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.12.3	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.12.4	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Adopt OpenJDK8u292b10
7.12.5	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Adoptium OpenJDK8u292b10

Confluence 7.11

Confluence	Tomcat	Bundled JRE	Tested JREs
7.11.0	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.11.1	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.11.2	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.11.3	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.11.6	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Adoptium OpenJDK8u292b10

Confluence	Tomcat	Bundled JRE	Tested JREs
7.10.0	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01

7.10.1	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.10.2	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01

Confluence 7.9

Confluence	Tomcat	Bundled JRE	Tested JREs
7.9.0	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.9.1	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.9.2	9.0.37	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01

Confluence 7.8

Confluence	Tomcat	Bundled JRE	Tested JREs
7.8.0	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.8 Adopt OpenJDK 8u252-b09
7.8.1	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.8 Adopt OpenJDK 8u252-b09
7.8.3	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01

Confluence 7.7

Confluence	Tomcat	Bundled JRE	Tested JREs
7.7.2	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.7.3	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.8 Adopt OpenJDK 8u252-b09
7.7.4	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.8 Adopt OpenJDK 8u252-b09

Note: 7.7.0 and 7.7.1 were internal releases.

7.6.0	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.6.1	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.6.2	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09

Confluence 7.5

Confluence	Tomcat	Bundled JRE	Tested JREs
7.5.0	9.0.33	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.5.1	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.5.2	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09

Confluence	Tomcat	Bundled JRE	Tested JREs	
7.4.0	9.0.33	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09	
7.4.1	9.0.33	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09	
7.4.2	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09	
7.4.3	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09	
7.4.4	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.8 Adopt OpenJDK 8u252-b09	
7.4.5	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01	
7.4.6	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01	
7.4.7	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01	

7.4.8	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.4.9	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.4.10	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Adopt OpenJDK8u292b10
7.4.11	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Adoptium OpenJDK8u292b10
7.4.12	9.0.45	Adoptium OpenJDK 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08
7.4.13	9.0.45	Adoptium OpenJDK 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08
7.4.14	9.0.45	Adoptium OpenJDK 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Adoptium OpenJDK8u302b08

Note: Adopt OpenJDK 11.0.3 and 11.0.4 had known issues in Linux and Windows in 7.4.0.

Confluence 7.3

Confluence	Tomcat	Bundled JRE	Tested JREs
7.3.1	9.0.27	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09
7.3.2	9.0.27	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09
7.3.3	9.0.27	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09
7.3.4	9.0.33	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09
7.3.5	9.0.33	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09

Confluence	Tomcat	Bundled JRE	Tested JREs
7.2.0	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13

7.2.1	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.2.2	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13

Note: Java 11 is supported, but not bundled in Confluence 7.2.

Confluence 7.1

Confluence	Tomcat	Bundled JRE	Tested JREs
7.1.0	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.1.1	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.1.2	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13

Note: Java 11 is supported, but not bundled in Confluence 7.1.

Confluence 7.0

Confluence	Tomcat	Bundled JRE	Tested JREs
7.0.1	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.0.2	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.0.3	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.0.4	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.0.5	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13

Confluence	Tomcat	Bundled JRE	Tested JREs
6.15.0	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 8u202
6.15.1	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 8u202
6.15.2	9.0.12	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.3	9.0.17	Adopt OpenJDK 8u202b08	Oracle JDK 8u202

6.15.4	9.0.19	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.5	9.0.19	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.6	9.0.19	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.7	9.0.21	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.8	9.0.22	Adopt OpenJDK 8u222b10	Oracle JDK 8u202
6.15.9	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.10	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202

Note: There was a known issue with Adopt OpenJDK 8u222b10, which was bundled with Confluence 6.15.8 CONFSERVER-58784 CLOSED .

Confluence 6.14

Confluence	Tomcat	Bundled JRE	Tested JREs
6.14.0	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 8u202
6.14.1	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 8u202
6.14.2	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 8u202
6.14.3	9.0.12	Adopt OpenJDK 8u202b08	Oracle JDK 8u202

Confluence	Tomcat	Bundled JRE	Tested JREs
6.13.0	9.0.12	Oracle JDK 1.8.0_192	-
6.13.1	9.0.12	Oracle JDK 1.8.0_192	-
6.13.2	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 1.8.0_192
6.13.3	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 1.8.0_192
6.13.4	9.0.12	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.5	9.0.19	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.6	9.0.19	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.7	9.0.22	Adopt OpenJDK 8u222b10	Oracle JDK 8u202
6.13.8	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.9	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.10	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.11	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.12	9.0.33	Adopt OpenJDK 8u252b09	Oracle JDK 8u251
6.13.13	9.0.33	Adopt OpenJDK 8u252b09	Oracle JDK 8u251
6.13.15	9.0.33	Adopt OpenJDK 8u252b09	Oracle JDK 8u251
6.13.17	9.0.33	Adopt OpenJDK 8u252b09	Oracle JDK 8u261

6.13.18	9.0.33	Adopt OpenJDK 8u252b09	Oracle JDK 8u261
6.13.19	9.0.40	Adopt OpenJDK 8u265b01	Oracle JDK 8u261
6.13.20	9.0.40	Adopt OpenJDK 8u265b01	Oracle JDK 8u261
6.13.21	9.0.40	Adopt OpenJDK 8u265b01	Oracle JDK 8u261
6.13.23	9.0.40	Adopt OpenJDK 8u265b01	Oracle JDK 8u291

Note: There was a known issue with Adopt OpenJDK 8u222b10, which was bundled with Confluence 6.13.7

CONFSERVER-58784 CLOSED

Note: Confluence 6.13.14 and 6.13.16 were internal releases.

Confluence 6.12

Confluence	Tomcat	Bundled JRE
6.12.0	9.0.11	Oracle JDK 1.8.0_181
6.12.1	9.0.11	Oracle JDK 1.8.0_181
6.12.2	9.0.12	Oracle JDK 1.8.0_181
6.12.3	9.0.12	Oracle JDK 1.8.0_181
6.12.4	9.0.12	Oracle JDK 1.8.0_181

Confluence 6.11

Confluence	Tomcat	Bundled JRE
6.11.0	9.0.10	Oracle JDK 1.8.0_162
6.11.1	9.0.10	Oracle JDK 1.8.0_162
6.11.2	9.0.11	Oracle JDK 1.8.0_181

Confluence 6.10

Confluence	Tomcat	Bundled JRE
6.10.0	9.0.8	Oracle JDK 1.8.0_162
6.10.1	9.0.8	Oracle JDK 1.8.0_162
6.10.2	9.0.10	Oracle JDK 1.8.0_162
6.10.3	9.0.19	Oracle JDK 1.8.0_202

Confluence	Tomcat	Bundled JRE
6.9.0	8.0.51	Oracle JDK 1.8.0_162
6.9.1	8.0.51	Oracle JDK 1.8.0_162
6.9.2	8.0.52	Oracle JDK 1.8.0_162

6.9.3 8.0	.52 Oracle JDK 1.8.0_162
-----------	--------------------------

Confluence 6.8

Confluence	Tomcat	Bundled JRE
6.8.0	8.0.50	Oracle JDK 1.8.0_162
6.8.1	8.0.50	Oracle JDK 1.8.0_162
6.8.2	8.0.51	Oracle JDK 1.8.0_162
6.8.3	8.0.51	Oracle JDK 1.8.0_162
6.8.4	8.0.52	Oracle JDK 1.8.0_162
6.8.5	8.0.52	Oracle JDK 1.8.0_162

Confluence 6.7

Confluence	Tomcat	Bundled JRE
6.7.0	8.0.48	Oracle JDK 1.8.0_162
6.7.1	8.0.48	Oracle JDK 1.8.0_162
6.7.2	8.0.48	Oracle JDK 1.8.0_162
6.7.3	8.0.50	Oracle JDK 1.8.0_162

Confluence	Tomcat	Bundled JRE
6.6.0	8.0.47	Oracle JDK 1.8.0_152
6.6.1	8.0.48	Oracle JDK 1.8.0_162
6.6.2	8.0.48	Oracle JDK 1.8.0_162
6.6.3	8.0.50	Oracle JDK 1.8.0_162
6.6.4	8.0.50	Oracle JDK 1.8.0_162
6.6.5	8.0.51	Oracle JDK 1.8.0_162
6.6.6	8.0.51	Oracle JDK 1.8.0_162
6.6.7	8.0.52	Oracle JDK 1.8.0_162
6.6.8	8.0.53	Oracle JDK 1.8.0_162
6.6.9	8.0.53	Oracle JDK 1.8.0_181
6.6.10	8.0.53	Oracle JDK 1.8.0_192
6.6.11	8.0.53	Oracle JDK 1.8.0_192
6.6.12	8.0.53	Oracle JDK 1.8.0_192
6.6.13	8.0.53	Oracle JDK 1.8.0_202

6.6.14	8.0.53	Oracle JDK 1.8.0_202
6.6.15	8.0.53	Oracle JDK 1.8.0_202
6.6.16	8.0.53	Oracle JDK 1.8.0_202
6.6.17	8.0.53	Oracle JDK 1.8.0_202

Confluence 6.5

Confluence	Tomcat	Bundled JRE
6.5.0	8.0.47	Oracle JDK 1.8.0_144
6.5.1	8.0.47	Oracle JDK 1.8.0_144
6.5.2	8.0.47	Oracle JDK 1.8.0_152
6.5.3	8.0.51	Oracle JDK 1.8.0_162

Confluence 6.4

Confluence	Tomcat	Bundled JRE
6.4.0	8.0.43	Oracle JDK 1.8.0_131
6.4.1	8.0.43	Oracle JDK 1.8.0_131
6.4.2	8.0.46	Oracle JDK 1.8.0_144
6.4.3	8.0.47	Oracle JDK 1.8.0_144

Confluence 6.3

Confluence	Tomcat	Bundled JRE
6.3.0	8.0.43	Oracle JDK 1.8.0_131
6.3.1	8.0.43	Oracle JDK 1.8.0_131
6.3.2	8.0.43	Oracle JDK 1.8.0_131
6.3.3	8.0.43	Oracle JDK 1.8.0_131
6.3.4	8.0.43	Oracle JDK 1.8.0_131

Confluence 6.2

Confluence	Tomcat	Bundled JRE
6.2.0	8.0.41	Oracle JDK 1.8.0_112
6.2.1	8.0.43	Oracle JDK 1.8.0_131
6.2.2	8.0.43	Oracle JDK 1.8.0_131
6.2.3	8.0.43	Oracle JDK 1.8.0_131
6.2.4	8.0.43	Oracle JDK 1.8.0_131

Confluence	Tomcat	Bundled JRE
6.1.0	8.0.41	Oracle JDK 1.8.0_112
6.1.1	8.0.41	Oracle JDK 1.8.0_112
6.1.2	8.0.41	Oracle JDK 1.8.0_112
6.1.3	8.0.41	Oracle JDK 1.8.0_112
6.1.4	8.0.43	Oracle JDK 1.8.0_131

Confluence	Tomcat	Bundled JRE
6.0.1	8.0.36	Oracle JDK 1.8.0_102
6.0.2	8.0.36	Oracle JDK 1.8.0_102
6.0.3	8.0.39	Oracle JDK 1.8.0_112
6.0.4	8.0.39	Oracle JDK 1.8.0_112
6.0.5	8.0.39	Oracle JDK 1.8.0_112
6.0.6	8.0.41	Oracle JDK 1.8.0_112
6.0.7	8.0.41	Oracle JDK 1.8.0_112

On this page:

- Confluence 7.16
- Confluence 7.15
- Confluence 7.14
- Confluence 7.13
- Confluence 7.12
- Confluence 7.11
- Confluence 7.10
- Confluence 7.9
- Confluence 7.8
- Confluence 7.7
- Confluence 7.6
- Confluence 7.5
- Confluence 7.4
- Confluence 7.3
- Confluence 7.2
- Confluence 7.1
- Confluence 7.0
- Confluence 6.15
- Confluence 6.14
- Confluence 6.13
- Confluence 6.12
- Confluence 6.11
- Confluence 6.10
- Confluence 6.9
- Confluence 6.8
- Confluence 6.7
- Confluence 6.6
- Confluence 6.5
- Confluence 6.4
- Confluence 6.3
- Confluence 6.2
- Confluence 6.1
- Confluence 6.0

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 customizing 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 organization has standardized 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 standardized operating environment and set up Confluence based on its supported platforms.

Migrate your Confluence site

Whether you're ready to make the move to cloud, or need the deployment and administrative flexibility of Data Center, we have everything you need to migrate successfully.

- Migrate from Server to Data Center
- Migrate from Confluence Cloud to Server
- Migrating Confluence Between Servers
- Move to a non-clustered installation
- From Confluence Evaluation through to Production Installation
- Cloud Migration Assistant for Confluence

Onsidering a move to cloud? Check our the Server to cloud migration guide.

Migrate from Server to Data Center

This page outlines the process for upgrading an existing Confluence Server site to Confluence Data Center.

If you're installing Confluence for the first time (you don't have any existing Confluence data to migrate), seeInstalling Confluence Data Center.

On this page:

Before you begin Migrate to non-clustered Data Center Review and upgrade your apps Upgrade your Confluence license Set up your cluster

Before you begin

Things you should know about when setting up your Data Center:

Your Confluence license determines the type of Confluence you have: Server or Data Center. Confluence will auto-detect the license type when you enter your license key, and automatically unlock any license-specific features.

See our Supported Platformspage for information on the database, Java, and operating systems you'll be able to use. These requirements are the same for Server and Data Center deployments.

Apps extend what your team can do with Atlassian applications, so it's important to make sure that your team can still use their apps after migrating to Data Center. When you switch to Data Center, you'll be required to switch to the Data Center compatible version of your apps, if one is available.

SeeEvaluate apps for Data Center migrationfor more information.

To use Confluence Data Center you must:

- Have a Data Center license (you canpurchase a Data Center licenseor create an evaluation license atmy.atlassian.com)
- Use asupportedexternal database, operating system and Java version
- Use OAuth authentication if you haveapplication linksto other Atlassian products (such as Jira)

If you plan to run Confluence Data Center in a cluster there are some additional infrastructure requirements. SeeClustering with Confluence Data Centerfor more information.



Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Migrate to non-clustered Data Center

Review and upgrade your apps

If you have any apps installed on your site, youll need to upgrade to the Data Center app version, if one is available. To avoid any impact to your apps, we recommend you do this before you enter your Confluence Data Center license key. Learn more about Upgrading server apps when you migrate to Data Center.

Upgrade your Confluence license

To move from Confluence Server to Confluence Data Center:

- 1. Go to Seneral Configuration
- 2. ChooseLicense Details from the sidebar under the Administration heading

3. Enter your Confluence Data Center license key.

There's no need to restart Confluence. Data Center features such as read-only mode, SAML single sign-on, and CDN will now be available.

Set up your cluster

If your organisation requirescontinuous uptime, scalability, and performance under heavy load, you'll want to run Confluence Data Center in a cluster.

To find out more about clustering, including infrastructure requirements see Clustering with Confluence Data Center.

If you're ready to set up your cluster now, head to Set up a Confluence Data Center cluster.



① Looking to migrate all your Atlassian applications to Data Center? Weve got you covered:

- Migrate to Bitbucket Data Center
- Migrate to Crowd Data Center
- Migrate to Confluence Data Center
- Migrate to Jira Data Center

Considering moving to cloud? Plan your cloud migration.

Migrate from Confluence Cloud to Server

①

Important changes to our server and Data Center products

Weve ended sales for new server licenses, and will end support for server on February 2, 2024. Were continuing our investment in Data Center with several key improvements. Learn what this means for you

This page is for people who are currently using Confluence Cloud, and wish to move to Confluence Server (a self-managed Confluence site).

Not moving from Cloud to Server?

These resources will help you plan your migration from:

- Confluence Server to Cloud
- Confluence Server to Data Center
- Confluence Server to Server

On this page:

- · Before you begin
 - Minimum Confluence version
 - Features and app availability
 - Templates
 - Team Calendars and Questions data
 - Migration approach
 - Infrastructure and database choice
 - Licenses
 - Account visibility
- Migration steps
 - Step 1: Check your apps
 - Step 2: Install Confluence Server or Data Center
 - Step 3: Export your Confluence Cloud Site
 - Step 4: Import your Confluence Cloud site export file
 - Step 5: Recover sys admin permissions
 - Step 6: Install any apps
 - Step 7: Check your application links
- Troubleshooting

Before you begin

There's a few things to understand before you begin this process. Ready to migrate? Skip to the migration steps

Minimum Confluence version

You can migrate from Confluence Cloud to **Confluence Server /Data Center 6.0 or later** only. You can't import Cloud data (either the whole site or individual spaces) into any earlier versions of Confluence.

We recommend installing either latest version of Confluence, or the latest Enterprise Release. The Confluence e Upgrade Matrix will help you choose the right version for your organisation.

Features and app availability

Some Cloud features won't be available in Confluence Server or Data Center. The navigation and user experience will also be different in some places. The core functionality of Confluence is the same however.

Marketplace apps are not automatically migrated. When you set up your Confluence Server or Data Center site, you'll need to reinstall each of your apps.

Not all apps are available for both Cloud and Server / Data Center. When planning your migration, we recommend you check that your essential apps are available for Server / Data Center in the Atlassian Marketplace and make a list of the ones you'll need to reinstall.

Templates

All pages that were created from a template will be migrated.

However, any custom templates you may have created in your Confluence Cloud site will not be migrated. You'll need to re-create your templates once your migration is complete.

You should also be aware that the range of built-in templates (known as blueprints) is much smaller in Confluence Server and Data Center, so some of the default templates you have previously used may not be available. See the full list of blueprints

Team Calendars and Questions data

Confluence Questions and Team Calendars data can't be migrated, as there is currently no way to export this data from Confluence Cloud.

Migration approach

You can choose to migrate your entire site in one go, or to import your team's content, space by space.

A **full site migration** involves a full site export (backup), and importing this file into Confluence Server or Data Center. Users and groups are included in this export. All spaces will be migrated, including archived spaces and personal spaces.

SeeMigration steps below to find out how to do this.

A **space by space migration** involves exporting each space individually, and importing these files into Confluence Server or Data Center one at a time. This means you can choose which spaces you want to migrate, or migrate in stages over time. Users and groups are not automatically migrated. If you've connected Confluence Server or Data Center to an external user directory, or have already populated your new site with user accounts, we'll attempt to attribute content to the right people on import.

SeeImport a space from Confluence Cloud if you plan to migrate your spaces one by one.

Infrastructure and database choice

Where you choose to host Confluence Server or Data Center is up to you. SeeSupported Platforms to find out which operating systems and databases are supported.

You can use any database listed on the Supported Platformspage, but you don't already have a database server, we recommend PostgreSQL, which is what Confluence Cloud is run on.

Licenses

You will need a new license to migrate to Confluence Server or Data Center. Your existing Confluence Cloud license can't be used. You can get a new license athttps://my.atlassian.com. You'll also need new licenses for any paid Marketplace apps.

Account visibility

In Confluence Cloud, people can choose not to make their profile information visible. This means when a Cloud site is imported into Server, user account information such as their full name, may not be included.

As long as you are logged in as a Site Admin when you complete the site export, email addresses will always be included, and used as the username when the user accounts are created. Users can then log in, and update their profile to provide the missing information.

Migration steps

This page will guide you through a **full site migration**. Seelmport a space from Confluence Cloudif you plan to migrate your spaces one by one.

Step 1: Check your apps

To check your apps are compatible:

- 1. In Confluence Cloud, go to Settings>Manage Apps.
- 2. Make a note of all User-installed apps.

3. Go tohttps://marketplace.atlassian.comand look up each app to see if a Server or Data Center edition is available.

Step 2: Install Confluence Server or Data Center

The way you do this depends on whether you're migrating to Server or Data Center, and how you plan to host the application.

SeeConfluence Installation Guide for links to all the installation options.

Step 3: Export your Confluence Cloud Site

To export your Confluence Cloud site:

- 1. Log in to Confluence Cloud as a Site Admin
- 2. In Confluence Cloud, go to Settings > Backup Manager
- 3. Follow the prompts to back up the site, and download the XML file.

The file will include all spaces and pages (including attachments), and all your users and groups.

Step 4: Import your Confluence Cloud site export file

Unless your site export file is quite small (less than 25mb) we recommend importing via the home directory method.



The import will overwrite all spaces, pages, and user accounts in your site - including your administrator account. You'll recover that account in the next step.

You should back up your database, home directory, and installation directory before you begin, in case you need to roll back.

To import a site from the home directory:

- 1. Copy your export file to < confluence home > / restore. (If you're not sure where this directory is located, the path is listed in the **Backup and Restore**screen)
- 2. Go to Co Seneral Configuration Backup and Restore.
- 3. Select your site export file under Import from home directory
- 4. Make sure **Build Index** is checked so that your index is created automatically.
- 5. Choose**Import**.

SeeRestoring a Site for more information about the import process.

Step 5: Recover sys admin permissions

When you import a site export file, all user accounts are overwritten, including the system administrator account that was created when you installed Confluence. Your existing Cloud Site Admin account will not automatically have system administrator permissions for Confluence Server or Data Center.

To recover system administrator permissions:

- 1. Stop Confluence.
- 2. Edit<installation-directory>/bin/setenv.shorsetenv.batand add the following system property, replacingyour-password>with a unique, temporary password.

-Datlassian.recovery.password=<your-password>

SeeConfiguring System Properties for more information on using system properties.

- 3. Start Confluencemanually (don't start Confluence as a service).
- 4. Log in to Confluence with the username recovery adminand the temporary password you specified in the system property.

- 5. Go to User Management > Add Users.
- 6. Enter the details for your new system administrator account and hit **Save**. Make sure to use a strong password.
- 7. Choose **Edit Groups** and select the confluence-administrators group. This is a super-group that has system administrator permissions.
- 8. Log out, and confirm that you can successfully log in with your new account.
- 9. Stop Confluence.
- 10. Edit<installation-directory>/bin/setenv.shorsetenv.batand remove the system property.
- 11. Restart Confluence using your usual method (manually or by starting the service).

SeeRestore Passwords To Recover Admin User Rights for more information on this process.

Step 6: Install any apps

To re-install your apps:

- 1. Log in to Confluence Server or Data Center as an administrator.
- 2. Go to 2 > Manage apps.
- 3. Follow the prompts to search for or upload the apps you identified in step 1. You'll need to purchase new licenses for these apps.

Remember that Team Calendars and Questions data is not included in your export, and cannot be migrated from Cloud at this time.

Step 7: Check your application links

If you had multiple Cloud products, such JIRA Software, you may need to make some changes to the application links.

To remove or update application links:

- 1. Go to General Configuration Application Links.
- 2. Follow the prompts to check and update any application links that are now pointing to the wrong place.

If you're unable to remove the Jira Cloud application link from your Confluence after the import, you'll need to remove those references directly from the Confluence database. SeeAlternative Methods of Deleting Application Links in Confluence.

Troubleshooting

There are a few known issues that you might encounter when importing your Cloud site.

Can't load pages in your new site

If you experience problems loading pages after the import, head to General Configuration to check your base URL as the port may have changed.

User management admin screens are missing

This is a fairly uncommon problem caused by a dark feature flag that is included in your Cloud site export file. See

CONFSERVER-35177 - User and Group Links Missing from Admin Console After Migrating From Cloud to Server GATHERING IMPACT

for a workaround.

Jira issues macros are broken

If your Confluence Cloud site has macros that depend on the Application Links back to a Jira Cloud instance, and you are migrating Jira as well, these references will need to be updated to work properly. See

▲ APL-1144 - Allow relocation of application links even if the target application is still accessible.
UNDER TRIAGE

for a workaround.

You can also edit the XML file prior to importing it into Confluence Server or Data Center, or by bulk editing those references in Confluence database. SeeHow to bulk update JIRA Issue Macro to point to a different JIRA instance.

Users' favourites (starred pages, or saved for later) are missing

If you find that some of your users' favorites (pages saved for later) are missing due to

CONFSERVER-36348 - Favourites missing after importing GATHERING IMPACT. . See How to restore missing favorites after import from XML for more information.

Some user accounts are missing or created without user details

Users in Confluence Cloud have the ability to change their profile visibility settings. To ensure all user data is included in the export, ask a site admin to perform the export.

Migrating Confluence Between Servers

This page describes how to move Confluence between physical servers using the same or a different operating system.

It doesn't coverdatabase migrationorupgradingyour Confluence version. We suggest you do each of these steps separately.

Transferring Confluence to another server

On this page:

Transferring Confluence to another server

To transfer Confluence to another server you will copy the home and install folders straight into an identical external database and user management setup. If your new server is using a different operating system there may be some additional changes at step 4.

- 1. Run the Confluence installer on your new server
- 2. Shut down Confluence on both your old and new servers
- 3. If you're using Oracle or MySQL, copythe driversfrom your old server to the new one
- 4. Delete the contents of the home directory on your new Confluence server, then copy in the contents of the home directory from your old Confluence server.
- 5. Make any additional changes required for your environment.

If the path to your home directory is different on the new server open the <code>Confluence_install_directory/confluence/WEB-INF/classes</code> directory and edit <code>confluence-init.</code> properties by changing the line starting with 'confluence.home='.

If you have also moved your database from one server to another you can change the JDBC URL in<confluence.home>/confluence.cfg.xmlif you are using a direct JDBC connection or in the definition of your datasource (if you are connecting via a datasource).

If you're migrating from **Windows to Linux**, you'll need to replace the backslashes with forward slashes in the following lines inconfluence.cfg.xml:

If you're migrating from **Linux to Windows**, you'll need to replace the forward slashes with backslashes in the following lines inconfluence.cfg.xml:

```
<property name="attachments.dir">${confluenceHome}\attachments/
<property name="lucene.index.dir">${localHome}\index/
cproperty name="webwork.multipart.saveDir">${localHome}\temp/
property>
```

- 6. Copy the <confluence-install>/conf/server.xmlfile from your old server to the same location on your new server
- 7. If you use a data source, ensure the data source points to the new database. See Configuring a datasource connection.
- 8. Start Confluence, then head to General configuration > License Details to add your license key

Westrongly recommend you perform are build of your content indices after performing a migration, to ensure Confluence search works as expected.

Move to a non-clustered installation

This page outlines how to switch from a clustered Confluence deployment to a non-clustered deployment. In these instructions we'll assume that you'll use one of your existing cluster nodes as your new, non-clustered installation.

If you have a valid Server license and want to move from Data Center to Server, see Moving from Data Center to Server.

Run Confluence in a cluster with one node

If you no longer need clustering for high availability or managing load, you can simply reduce the number of application nodes in your cluster to one. There are some advantages to this setup, as it is very easy to add more nodes if you require them in future, but there is a small performance overhead as Confluence will still operate as a cluster.

Move to a non-clustered installation

If you no longer need clustering, and want to avoid the overhead that comes from running a cluster with just one node, you can go back to a non-clustered (sometimes known as standalone) Data Center installation.

In these instructions we'll assume that you'll use one of your existing cluster nodes as your new, non-clustered installation. You'll also need to make some infrastructure changes as part of the switch. We recommend completing this process in a staging environment, and running a set of functional tests, integration tests, and performance tests, before making these changes in production.

Terminology

In this guide we'll use the following terminology:

- Installation directory The directory where you installed Confluence.
- Local home directory The home or data directory stored locally on each cluster node (if Confluence is not running in a cluster, this is simply known as the home directory).
- Shared home directory The directory you created that is accessible to all nodes in the cluster via the same path.

1. Shut down Confluence

Make sure read only mode is turned off, then stop Confluence on all cluster nodes before you proceed.

2. Configure your load balancer

Configure your load balancer to redirect traffic away from all Confluence nodes, except the node you plan to keep.

If you no longer need your load balancer, you can remove it at this step.

3. Move items in the cluster shared home back to local home

To move everything back to your local home:

- 1. Create a directory called/shared-homein the<local home> directory on the node you plan to keep (if you removed this directory when you set up clustering).
- 2. Move the following directories and files from your<shared home>directory to the<local home> /shared-homedirectory
 - config
 - · confluence.cfg.xml
 - dcl-document
 - dcl-document_hd
 - dcl-thumbnail
- 3. Move the remaining contents of your <shared home> directory to the root of your <local home> directory. Make sure your attachments directory is moved as part of this step.

Your cluster's shared home directory should now be empty.

▲Make sure you don't accidentally overwrite theconfluence.cfg.xmlin your local home directory. Theconfl uence.cfg.xmlfile from your shared home directory doesn't contain the same parameters as the one in your local home directory.

From Confluence 7.12, you can choose to skip this step and keep your existing shared home directory. For example, this may be beneficial if you're using elastic storage for the <shared home>/attachments directory and want to keep that setup.

4. Modify cluster properties

- 1. Take a backup of the existing < local home > / confluence.cfg.xml
- 2. Edit<local home>/confluence.cfg.xml
- 3. Change thesetupType parameter from cluster tocustom:

```
<setupType>custom</setupType>
```

4. Remove all cluster properties that begin withconfluence.cluster.

Here are some example cluster properties that should be removed. These will vary depending on how you configured your cluster.

```
confluence.cluster
confluence.cluster.address
confluence.cluster.home
confluence.cluster.interface
confluence.cluster.join.type
confluence.cluster.name
```

A If you chose to keep your shared home directory at the previous step, do not remove the confluen ce.cluster.home property, or Confluence will not know where to find your shared home, or attachments directory.

5. Save the file.

5. Start Confluence

Restart Confluence.



To confirm you're now running a standalone installation, go to General Configuration > Clustering.

The active cluster should no longer appear. Instead, you'll see information about getting started with clustering, and the option to enable cluster mode.

Additional steps if you have a Synchrony cluster

If you also have a Synchrony cluster, but would prefer to let Confluence manage Synchrony for you, you'll need to make some additional changes.

SeeMigrate from a standalone Synchrony cluster to managed Synchrony. This guide assumes you're running Confluence in a cluster, but the steps are similar for a non-clustered installation.

From Confluence Evaluation through to Production Installation

i Important changes to our server and Data Center products

Weve ended sales for new server licenses, and will end support for server on February 2, 2024. Were continuing our investment in Data Center with several key improvements. Learn what this means for you

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 Cloud site.
- You will evaluate Confluence on an installed version too, not a Confluence Cloud site.

If you are using Confluence Cloud to evaluate Confluence, please refer to the following guide when you want to move to an installed version: Migrate from Confluence Cloud to Server.

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:

- Download the 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 **Trial**or**Production** installation type. If you are not sure, choose**Trial Installation**.
 - The **Trial**option will install Confluence with default settings, including the embedded database which is automatically set up for you. You'll need to migrate to an external database before running Confluence in a production environment (more info below).

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 Marketplace apps as part of your evaluation
- Step 4. Set up your production Confluence site

Related pages:

- Supported Platforms
- Add and Invite Users
- Getting Started as Confluence Administrator
- Confluence installation and upgrade guide

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 'Trial 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.
- ChooseCreate to add pages and blog posts.

To add users: Go to User management.

Step 3. Look for interesting Marketplace apps as part of your evaluation

If you have finished evaluating Confluence, you can skip this step.

Apps, also called plugins or add-ons, provide additional features that you can install into your Confluence site. Some of them are provided free of charge. Many of the commercial apps are available free for an evaluation period.

You can browse and download app on the Atlassian Marketplace. You can also find apps via the Confluence user interface, which interacts with the Atlassian Marketplace for you.

To find useful apps via the Confluence user interface:

- 1. Go to O > Manage apps.
- 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:

- 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. SeeRunning 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. SeeSetting Up Public Access.
- Set up your permission scheme. See Permissions and restrictions.
- Connect Confluence to Jira applications such as Jira Software or Jira Service Management or 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.

Cloud Migration Assistant for Confluence

robots	noindex
descr iption	Step-by-step instructions on how to use the Confluence Cloud Migration Assistant to move your data from Confluence Server or Data Center to Confluence Cloud.

robots	noindex	
robots	noindex	

Before you migrate, check your cloud organization

Were currently rolling out changes that may affect your migration experience. From your organization at admin. atlassian.com, if the Users list and Groups list are under the Directory tab, you have the improved user management experience. This means that the users and groups across sites will be merged under the organization. Read more about how groups and permissions are migrated. If you have any concerns, contact support.





For a test migration or UAT, we recommend that your test cloud site is not part of the organization that also hosts your prod site. The prod site should be hosted in a different organization. This is to ensure smooth migration of the relevant users and groups.

The Confluence Cloud Migration Assistant is an app that helps you easily move content, users, and groups from Confluence Server or Data Center to Confluence Cloud. Built and maintained by Atlassian, the app is free to install and use.

With the app, you can choose what you want to move to the cloud, start migrating at your convenience, and monitor the progress of everything throughout the migration process.



(i) Important changes to our server and Data Center products

Weve ended sales for new server licenses, and will end support for server on February 2, 2024. Were continuing our investment in Data Center with several key improvements. Learn what this means for you

When to use the Confluence Cloud Migration Assistant

- When you want to move users or data from Confluence Server or Data Center to Confluence Cloud.
- When you want to assess your apps before moving from Confluence Server or Data Center to Confluence Cloud.
- When you want to run a test or trial migration from Confluence Server or Data Centerto Confluence
- When the Atlassian Support team has recommended using the app.

The Confluence Cloud Migration Assistant will not work for Jira products. You can download the Jira Cloud Migration Assistant for Jira migrations to cloud.



Before you begin

Make sure you have reviewed the server to cloud migration guide. This guide will walk you through the migration process step-by-step and help you identify what to look out for.

Before attempting a test or production migration, ensure you've completed all of the steps for the Confluence Cloud Migration Assistant in thepre-migration checklist. The checklist will help you prepare yourself and your data for migration, and ensure you avoid common sources of migration failure.

Install the Confluence Cloud Migration Assistant app

If your Confluence Server site is version 6.13 or above you won't need to install anything because it comes preinstalled, although you may be asked to update the app.

To install the app on versions 5.10 to 6.12:

- 1. In Confluence Server go to > Manage apps.
- 2. ChooseFind new add-ons.
- 3. Search for the Confluence Cloud Migration Assistantapp.
- 4. Choose**Install**and you're all set.

Alternatively, you can install it from the Atlassian Marketplace.

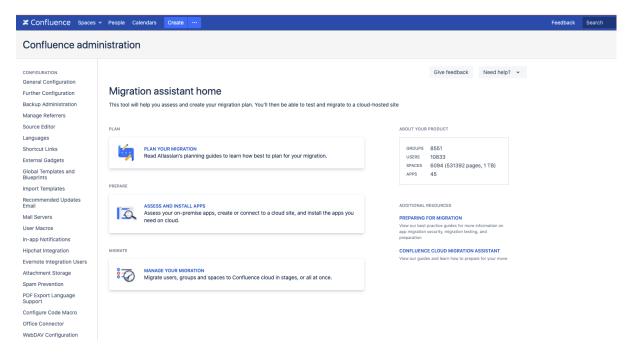
Once installed, you can access the migration assistant by going to Confluence Administration>look for the Atla ssian Cloudcategory > select Migration Assistant.



If your Confluence Server site is behind a firewall, you'll need to allow access to the domain: atlassian. com

Use the migration assistant to assess your apps

Carrying out an assessment of your apps helps you to establish which apps are needed for a migration.



You can find step-by-step instructions for this process in Assessing and migrating apps with the Confluence Cloud Migration Assistant.

Check for possible data conflicts in your cloud site

You can reduce the risk of running into issues, or the migration failing, if you conduct somemanual checks in your server and cloud sites.

1. Check for groupconflicts

Make sure that there are no groups already in your cloud site with the same name as groups from your server site, unless you are intentionally trying to merge them.

If we find a group in your server site that has the samename as a groupin your cloud site (either Jira or Confluence), we will merge the users from the server group into the cloud group. The server group users will inherit the permissions of the cloud group. This also applies to groups with Jira product access that have the same name as a Confluence group you are migrating. This is because all users and groups are managed in a central location in your cloud site.

If you dont want this to happen, youll need to make sure all groups across server and cloudhave unique namesb efore running your migration.



The following groups manage admin access and are blacklisted. They will not be migrated at all: "siteadmins", "system-administrators", "atlassian-addons", "atlassian-addons-admin". Users in these groups will still be migrated; if you want them to be in one of the blacklisted groups youll need to manually add them after migration.

2. Check for space keyconflicts

Before migrating check that there are no spaces with the samespace keybetween your server and cloud sites.

If a space from your server site has the samespace keyas a space in your cloud site your migration will fail. This is because every space in Confluence Cloud must have a unique space key. If you find aconflict youcan:

- delete duplicate spaces from yourcloudorserversites
- reset your cloud site
- choose not to migrate these spaces

If the migration assistant finds a conflict, the space will not migrate.

If a space key conflict is caused by a previous test migration you canreset your cloud sitebefore migrating.

Use the app to set up and run your migration

Once you have the app installed, there are five key steps to set up and run your migration from server or Data Center to cloud:

- 1. Connect to cloud
- 2. Choose what to migrate
- 3. Check for errors
- 4. Review your migration
- 5. Migrate



The sections below describe each step in detail and explain some common errors that you may come across. If you have technical questions or issues while using the migration assistant, get in touch with our support team.



Running a test migration

We strongly recommend doing a trial run of your migration to a test or staging site before running your final migration. Check out our guidance on testing your migration.

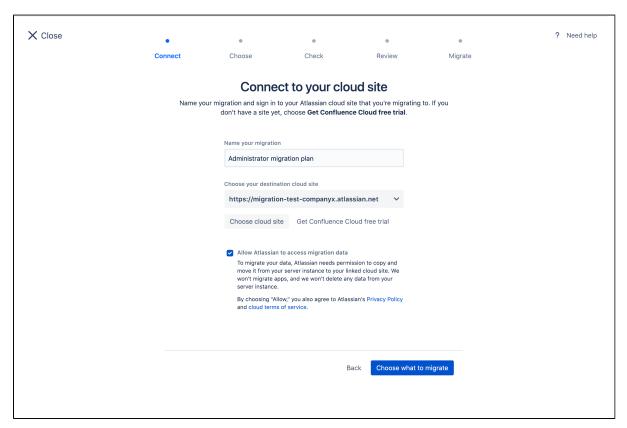
1. Connect to your destination Confluence Cloud site

Youll be asked to add a name for your migration and choose which cloud site you would like to migrate to. You need to be an admin in both your server and the destination cloud sites.

If you have already connected a cloud site, you should see it in the dropdown. If there is nothing there, you will need to either connect a new cloud site or sign up for a new cloud license.

When youre ready to go, check the box to allow Atlassian to move your data from your server site to your cloud site. If youre unable to grant Atlassian this access, you wont be able to migrate with the migration assistant and will need to do a space import instead.

If your Confluence Server site is behind a firewall, you'll need to allow access to the domain:atlassian.com. You also might need to allow access toother Atlassian domains.

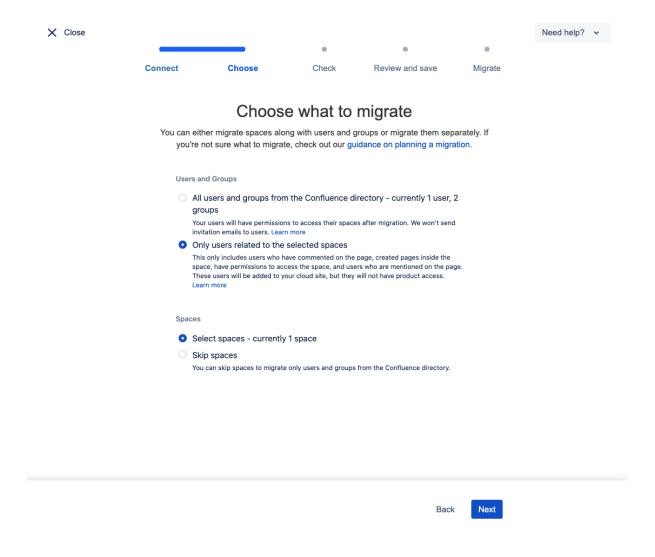


2. Choose what to migrate

You can migrate everything together or break it up into different stages.

You can choose:

- all or some of your users and groups
- which individual spaces (and their attachments) you'd like to migrate



Users and groups

You can choose to either migrate all or some of your users.

If you choose the migrate your users, the first time you do so all your users will be added to your cloud site. Every migration, after the first, we will just link your data to the users that already exist in cloud. If you have a large userbase we suggest following our recommendations.

When you migrate your users, they will be added to their groups when they get to cloud. You will need to review and approve group permissions after you migrate. When you approve group permissions, your users will be given Confluence access and will be added to your bill.

We wont send an invitation to your users. To invite your users you can choose to send an invitation from the **Ad ministration** space after you have migrated, or send a link for them to log in themselves.

When you select Only users related to the selected spacesunder users and groups, we will still migrate some user data connected to the spaces you are migrating. This is to make sure that mentions, comments, and page history stay active.

User data that will be migrated every time includes:

- full name
- username (discarded after migration)
- email address

We will only migrate this information for users directly connected to the spaces you are migrating. We will not give these users product access or add them to any groups. They will appear in your cloud site user list.

If you choose to migrate users later, their product and group access will be updated.

Also, if you choose not to migrate users and groups and you have a space permission granted by a group that don't exist in cloud, the Confluence Cloud Migration Assistant will not migrate the respective space permission. To avoid this scenario, we recommend you to create the specific group in the cloud site before migration.

Other things to be aware of when migrating users and groups:

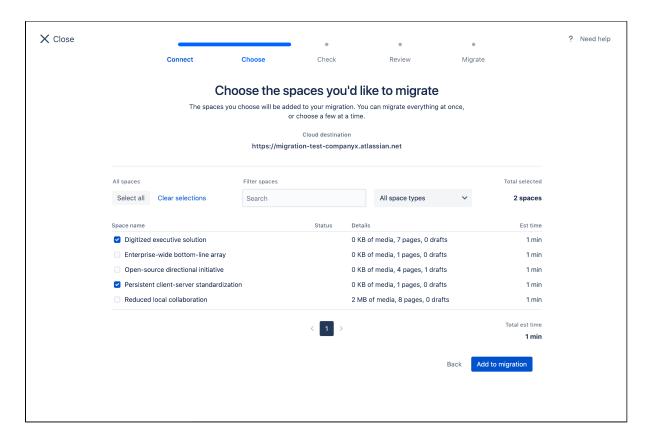
- Users are migrated using email address as the source of truth. On subsequent migrations, the migration assistant will link users by email address rather than re-migrating them. Check out ourtips for migrating a large number of users.
- You must validate all your user accounts (email addresses) before migrating to cloud. Migrating unknown user accounts can potentially allow unauthorized access to your cloud sites. For example, if you had users in your server instance with emails that you dont own, say email@example.com, you might be inviting someone who owns @example.com to your site in cloud.
- Confluence Cloud is subscription-based and billed on a per-user basis. If you plan to migrate your users, make sure you check the licensing options available.
- If you use an external user management system, we recommend synchronizing it with your local directory before migrating. This is to make sure that your users and groups are up to date before you transfer any data.
- Users with disabled status in your server site will be migrated as active but without any product access. This means they will not be counted as active Confluence users for billing purposes.
- If we find a group in your server site that has the same name as a group in your cloud site, we will merge the users from the server group into the cloud group.
- Global settings and global site permissions are not migrated with this tool. Youll need to set these manually after migration.
- If you have users that already exist in your destination cloud site and you choose to migrate users with this app, the following will occur:
 - If a user has product accessin cloud, but has disabled status in your server site, they will continue to have product access in cloud after migration.
 - If a user does not have product access in cloud, but is enabled in your server site, they will be granted product access through the migration process.



 If you use Confluence as a knowledge base for Jira Service Management (formerly Jira Service Desk), your Jira Service Management users may also be migrated along with your Confluence users. This will happen if you can see your Jira Service Management users in the cwd_user table in Confluence.

Spaces

If you want to migrate all or some of your spaces choose Select spaces from the options. You will then be able to select what spaces you want to migrate. If you arent migrating any spaces you will be taken straight to check for errors.



Select the spaces you want to add to your migration. You can filter the list or search for particular spaces, or click **Select all** if you want to migrate everything at once. You wont be able to migrate spaces with space keys that already exist in your Confluence Cloud destination site.

If a space has a MIGRATED status, we have detected that you have already migrated this space to the same cloud site.

If a space has a QUEUED status, it has already been added to a migration that is waiting to be run.



If you have lots of spaces and attachments or you are on Data Center, you might want to break the migration up into a few smaller migrations. The migration assistant can be slow to load and process tasks when there is a lot to manage.

When you've chosen all your spaces, selectAdd to migration.

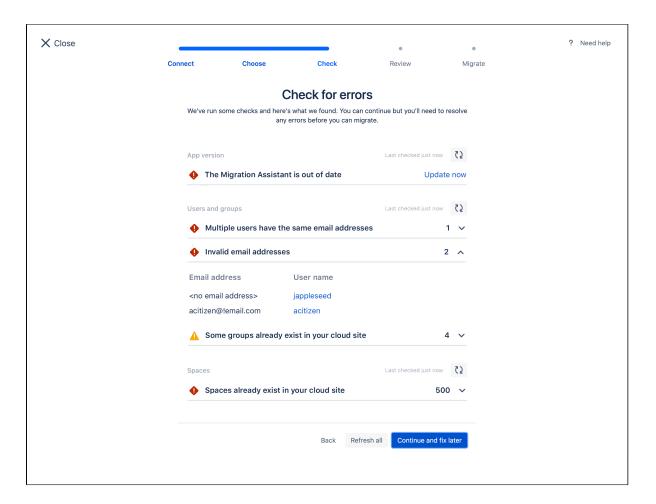
3. Check for errors

In this step, the Confluence Cloud Migration Assistant will review your migration and check for common errors. It will check if your:

- migration assistant app is up to date
- users have valid and unique email addresses
- groups will merge through the migration process
- spaces already exist in your cloud site
- spaces are publicly available and searchable online



You may also encounter other issues during the migration process; this step only checks for the issues mentioned here.



If there is a green tick othen the check has passed. If you get a warning sign ▲then you can continue, but you need to be aware of a potential issue.

If a check comes back with a red error othen you will need to resolve the error before you can run your migration.

If you decide to **Continue and fix later**, you can come back to view the errors once you have saved your migration.

Updating the app

The migration assistant may be out of date. If you get this error, youll need to update it before running any migrations.

Users and groups errors

All users will need to have a valid and unique email address. If we detect invalid emails or multiple users with the same email, you will get an error. You will need to fix these email addresses before you can run your migration.

If you have chosen to migrate all users, we will check to see if you have any groups with the same name already in your cloud site. If we find groups with the same name, we will merge the users from the server group into the cloud group with the same name. You can continue with your migration without fixing this issue, but its important to check that this wont cause permission escalation.



The following groups manage admin access and are **blacklisted.** They will not be migrated at all: "site-admins", "system-administrators", "atlassian-addons", "atlassian-addons-admin". Users in these groups will still be migrated; if you want them to be in one of the blacklisted groups youll need to manually add them after migration.

Space errors

If youre migrating spaces we will check to see if there will be any space key conflicts. If you get an error you can:

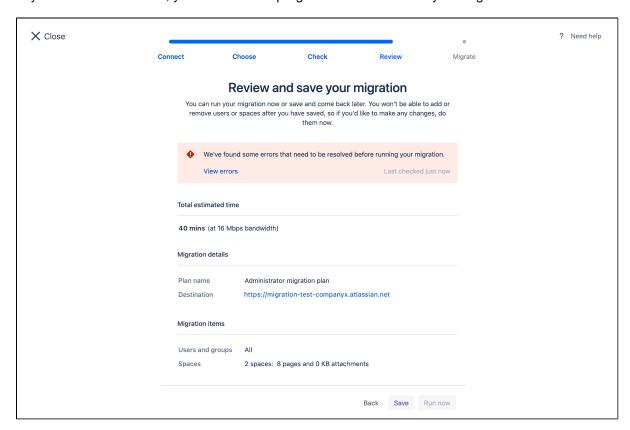
- delete duplicate spaces from your cloud or server sites
- reset your cloud site
- choose not to migrate these spaces by removing them from your migration

You will need to resolve any space key conflicts before you can run your migration.

4. Review your migration

This is the final step in setting up your migration.

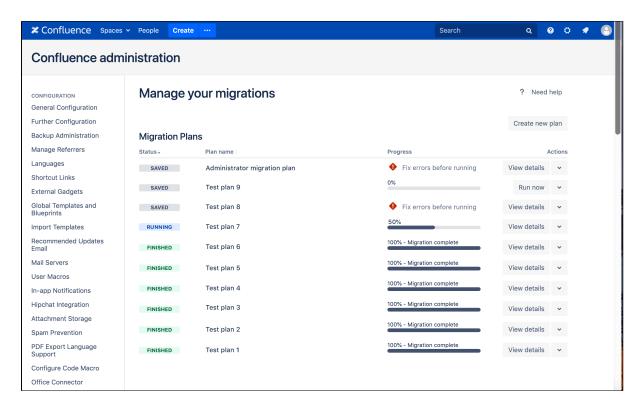
If everything looks correct and you want to start your migration, click **Run**. If you would like to start your migration later or you still have errors to fix, click **Save**. If you choose to run your migration, it will still be saved to your dashboard. There, you can view the progress and details of all your migrations.



5. Manage your migrations

Your saved migration will be listed on the migration dashboard, where you can **view details** or **run** it. You can also check the status of a migration, monitor the progress, stop a migration that's currently running, or create a new one.

You can create as many migrations as you need. At this time, migrations can't be edited or deleted, so if you create a migration that can't be used, just create a new one.



Status definitions

Your migration is saved and ready to run.

RUNNING

Your migration is currently in progress.

FINISHED

All tasks in your migration have been completed.

STOPPED

Your migration has been manually stopped. Once stopped, it can't be resumed. Any step already in progress will first need to finish before the migration is shown as fully stopped. Some users, groups, and spaces may already have been migrated to your Confluence Cloud site.

FAILED We were unable to complete the migration. This might be because a space key already exists in the destination site, or the migration hit an unexpected error. Some users, groups, and spaces may already have been migrated to your Confluence Cloud site.

After migrating

After migrating spaces, it may take a while for them to appear in the space directory. However, you can still access them via a direct link.

Depending on the type of migration, there may be some things you need to do once your migration is finished.

Users and groups

To make sure your users and groups are set up correctly:

- Review members of groups and approve their permissions by going to Review imported groups. (If you
 have the Free plan, permissions cant be modified; users and groups retain the same permissions that
 they had on your original site.)
- Add users to the generic groups if necessary. The generic groups are: "site-admins", "system-administrators", "atlassian-addons", "atlassian-addons-admin".
- If you use an external user management system, check that your users have synced correctly.
- When you are ready, invite your users. Go to **Administration > Users > Show details** and then **Resend invite**. When they first log in they may be prompted to set a new password and add personal details.



If you have the improved user management experience, go toAdministration > Directory > Users > Show details and then Resend invite.

We recommend providing some training or sending an onboarding email to your users tohelp them get familiar with their new cloud workspace.

Spaces

To check that your spaces have migrated successfully:

- Review content and spaces, or ask your users to review their own content.
- Check for any instances of Former User. This means that we were unable to match content to a user.
- Link your other Atlassian products by going to Settings > Application links.
- Use the Jira macro repair to update any links to Jira. On your cloud site go to Settings > Jira macro repair and follow the steps.



Confluence short links like https://confluence.example.com/x/PywS may not work after migrating. Replacing them with internal links (or full URLs if theyre not in your Confluence site)\before migrating should solve this issue.

You can then install any apps you wish to use and onboard your users.

For the full overview of post-migration actions check out the server to cloud migration guide.

More information and support

We have a number of channels available to help you with your migration.

- For more migration planning information and FAQs, visit the Atlassian Cloud Migration Center.
- Have a technical issue or need more support with strategy and best practices? Get in touch.
- Looking for peer advice? Ask the Atlassian Community.
- Want expert guidance? Work with anAtlassian Partner.

Confluence Data Center documentation

Data Center is our self-managed edition of Confluence built for enterprises. It provides the deployment flexibility and administrative control you need to manage mission-critical Confluence sites.Learn more aboutC onfluence Data Centeron our website.

Data Center architecture

You can deploy Confluence Data Center in two ways.



Non-clustered (single node)

Run the Confluence Data Center application on a single server. (Available for Confluence 7.2 and later).

This allows you to take advantage of Data Centeronly features without adding to your infrastructure.

Learn more about Data Center features.

Clustered

Run Confluence Data Center in a cluster with multiple application nodes, and a load balancer to direct traffic.

Clustering is designed for large, or mission-critical, Confluence sites, allowing you to provide high availability, and maintain performance as you scale.

Learn more about clustering with Data Center.

Get started

Install or upgrade Confluence Data Center

- Install Confluence Data Center from scratch
- Migrate from Server to Data Center

Clustering with Confluence Data Center

- Learn about clustering architecture and requirements
- Set up a Data Center cluster
- Add or remove application nodes
- Turn off clustering (revert to a non-clustered Data Center installation)
- Troubleshoot a clustering issue

Want to see what's included with Data Center? Head to the Confluence Server and Data Center feature comparison.



You can purchase a Data Center license or create an evaluation license at my.atlassian.com

Getting Started with Confluence Data Center

Data Center is our self-managed edition of Confluence built for enterprises. It provides the deployment flexibility and administrative control you need to manage mission-critical Confluence sites.

You can run Confluence Data Center documentation in a cluster, or as standalone (non-clustered) installation.

This guide covers**clustered** Data Center deployments.

On this page:

- 1. Define your requirements
- 2. Provision your infrastructure
- 3. Plan your deployment
- 4. Install and configure Confluence Data Center
- 5. Maintain and scale your Confluence cluster

1. Define your requirements

Before you get started, its a good idea to define your organizations goals and requirements. If you need high availability, scalability, and performance under heavy load, you'll want to run Confluence Data Center in a cluster.

To prepare, we recommend assessing:

- the number of users you have
- the amount of data you have
- your expected usage patterns
- the resources your organization has allocated to maintain your Confluence site.

For more information about disaster recovery for Confluence, head to Confluence Data Center disaster recovery.

Our sizing and performance benchmarks can help you assess your expected load, and predict performance:

- Confluence Data Center load profiles
- Confluence Data Center performance
- Infrastructure recommendations for enterprise Confluence instances on AWS

2. Provision your infrastructure

Once you've identified your organization's needs, you can prepare your clustered environment.Read ourClustering with Confluence Data Centerfor importanthardware and infrastructure considerations.

To help you get started with clustering, we've provided a Confluence Data Center sample deployment and monitoring strategy.

We've also provided some general advice about node sizing and load balancers, to help you find your feet if this is your first clustered environment:

- Node sizing overview for Atlassian Data Center
- Load balancer configuration options
- Traffic distribution with Atlassian Data Center

3. Plan your deployment

If you're new to Confluence, you can try out Confluence Data Center bydownloading a free trial. This can help you identify dependencies and plan your path to production.

Migrating from Confluence Server to Confluence Data Center? Read throughthese guides to help minimize disruption during the switch:

- Moving to Confluence Data Center
- Atlassian Data Center migration plan
- Atlassian Data Center migration checklist

It's also important to take aninventory of your third-party apps (also known as add-ons) to make sure they're compatible with Data Center. Using a large number of add-ons can degrade performance, soit's a good idea to remove any add-ons that aren't crucial to functionality.

Find out how to evaluate add-ons for Data Center migration.

4. Install and configure Confluence Data Center

Once your environment is ready, it's time to install and configure Confluence Data Center in a cluster.

How you install depends on your environment:

- Your own hardware seeInstalling Confluence Data Center
- Kubernetes seeRunning Data Center products on a Kubernetes cluster
- Azure seeGetting started with Confluence Data Center on Azure
- AWS (Amazon Web Services) seeRunning Confluence Data Center in AWS

If you're migrating from Confluence Server to Confluence Data Center, follow the instructions outlined in Migrate from Server to Data Center.

Before deploying Confluence Data Center to production, we recommend thoroughly testing the installation. Head to our Data Center migration plan for detailed advice about testing and launching to production.

5. Maintain and scale your Confluence cluster

Once you've deployed your Confluence Data Center cluster in production, here are some resources for monitoring the health of the cluster, and scaling it up to accommodate moreusers:

Tools for monitoring your Data Center application

Ready to grow? Read up on scaling and adding nodes to your new Confluence Data Centercluster:

- Scaling with Atlassian Data Center
- Adding or removing Confluence Data Center nodes

Confluence Server and Data Center feature comparison

Important changes to our server and Data Center products

Weve ended sales for new server licenses, and will end support for server on February 2, 2024. Were continuing our investment in Data Center with several key improvements. Learn what this means for you

If you manage your own Confluence site (it's not hosted by Atlassian), you'll have either a Confluence Server orConfluence Data Centerlicense. If we manage Confluence for you, you'll have a Confluence Cloud license.

Your Confluence license determines which features and infrastructure choices are available.

We want all teams to get the most out of Confluence, so the core features are available for everyone including creating pages, working together, and organizing your work.

Feature comparison

Heres a summary of available features for Confluence Server and Confluence Data Center. If youre interested in having Atlassian host and manage your products, see how a cloud plan compares on our Conflu ence features page.

Core features	Server license	Data Center license
Create spaces Create spaces to store your team or project work.	•	•
Create pages Create pages and blog posts, and work on them with your team.	•	•
Collaborative editing Up to 12 people can work on the same page at the same time. Learn more	•	•
Browser and mobile Use your browser, or use the iOS or Android app. Learn more	•	•
Team calendars Create and view calendars from your organization. Learn more	8	7.11+
Analytics Track engagement with all the content in your site. Learn more	8	7.11+
User management		
External user directories Store users in Active Directory, Crowd, Jira or another LDAP directory. Lear n more		•
Single sign-on Use a SAML or OpenID Connect identity provider for authentication and single-sign on. Learn more	via Crowd	6.1 +
Advanced permissions management Inspect user and group permissions for auditing and troubleshooting purposes. Learn more	8	7.3 +
High availability and performance at scale		

Clustering Run Confluence on multiple nodes high availability. Learn more	8	5 .6 +
Content Delivery Network (CDN) support Improve geo-performance for distributed teams. Learn more		7.0 +
Infrastructure and Control		
Read-only mode Limit what users can do in your site while you perform maintenance. Learn more	8	6.10 +
Sandboxed processes Run resource intensive tasks in external sandboxes for greater stability. Learn more	8	6.12 +
Rate limiting Control how many external REST API requests users and automations can make. Learn more	8	7 .3 +
Advanced auditing Access a wider range of audit events, and integrate with third-party logging systems. Learn more	8	? 7.5 +
Retention rules Set rules to automatically delete historical page versions and purge deleted items from the trash. Learn more	8	7 .16 +
Rolling upgrades	×	7 .9 +
Upgrade to the latest bug fix update of the same feature release with no downtime. Learn more		
Deployment options		
Your own hardware Run Confluence on your own physical servers, virtualized servers, or in the data center of your choice.	•	•
AWS Quick Start Use our Cloud Formation Templates to deploy Confluence on AWS. Learn more	8	6.1 +
Azure template Use our template to deploy Confluence on Azure. Learn more	8	6.6+
Kubernetes Helm charts Use our Helm charts to deploy Confluence on Kubernetes.Learn more	8	7 .13 +

Clustering with Confluence Data Center

Confluence Data Center allows you to run a cluster of multiple Confluence nodes, providing high availability, scalable capacity, and performance at scale.

This guide describes the benefits of clustering, and provides you an overview of what youll need to run Confluence in a clustered environment, including infrastructure and hardware requirements.

Ready to get started? See Set up a Confluence Data Center cluster

Is clustering right for my organization?

On this page

- Is clustering right for my organization?
- Clustering architecture
- Infrastructure and hardware requirements
- App compatibility
- Ready to get started?

Clustering is designed for enterprises with large or mission-critical Data Center deployments that require continuous uptime, instant scalability, and performance under high load.

There are a number of benefits to running Confluence in a cluster:

- **High availability and failover:** If one node in your cluster goes down, the others take on the load, ensuring your users have uninterrupted access to Confluence.
- Performance at scale: each node added to your cluster increases concurrent user capacity, and improves response time as user activity grows.
- **Instant scalability:** add new nodes to your cluster without downtime or additional licensing fees. Indexes and apps are automatically synced.
- Disaster recovery: deploy an offsite Disaster Recovery system for business continuity, even in the
 event of a complete system outage. Shared application indexes get you back up and running quickly.
- Rolling upgrade: upgrade to the latest bug fix update of your feature release without any downtime.
 Apply critical bug fixes and security updates to your site while providing users with uninterrupted access to Confluence.

Clustering is only available with a Data Center license. Learn more about upgrading to Data Center.

Clustering architecture

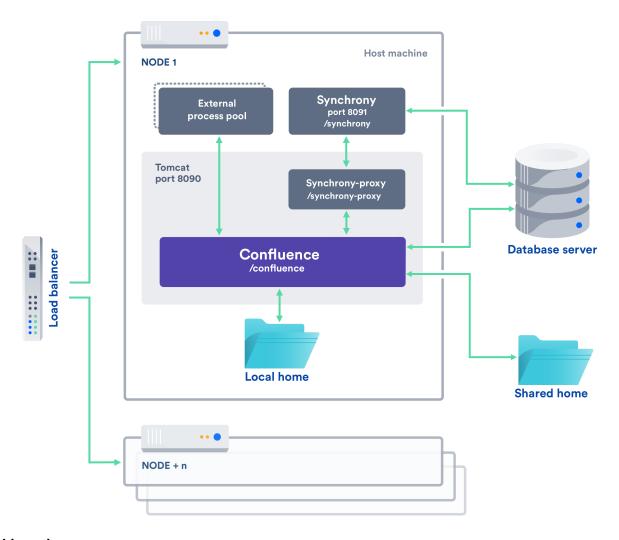
The basics

A Confluence Data Center cluster consists of:

- Multiple identical application nodes running Confluence Data Center.
- A load balancer to distribute traffic to all of your application nodes.
- A sharedfilesystem that stores attachments, and other shared files.
- A database that all nodes read and write to.

All application nodes are active and process requests. A user will access the same Confluence node for all requests until their session times out, they log out, or a node is removed from the cluster.

The image below shows a typical configuration:



Licensing

Your Data Center license is based on the number of users in your cluster, rather than the number of nodes. This means you can scale your environment without additional licensing fees for new servers or CPU.

You can monitor the available license seats in the License Details page in the admin console.

If you wanted to automate this process (for example to send alerts when you are nearing full allocation) you can use the REST API.

The following GET requests require an authenticated user with system administrator permissions. The requests return JSON.

<pre><confluenceurl>/rest/license/1.0/license /userCount</confluenceurl></pre>	Number of active users
<pre><confluenceurl>/rest/license/1.0/license /remainingSeats</confluenceurl></pre>	Number of users you can add before reaching your license limit
<pre><confluenceurl>/rest/license/1.0/license /maxUsers</confluenceurl></pre>	Maximum number of users allowed by your license

Your Confluence license determines which features and infrastructure choices are available. Head to Confluence Server and Data Center feature comparison for a full run down of the differences between a Server license and a Data Center license.

Home directories

To run Confluence in a cluster, you'll need an additional home directory, known as the shared home.

Each Confluence node has a local home that contains logs, caches, Lucene indexes and configuration files. Everything else is stored in the shared home, which is accessible to each Confluence node in the cluster. Marketplace apps can choose whether to store data in the local or shared home, depending on the needs of the app.

Here's a summary of what is found in the local home and shared home:

Local home	Shared home
 logs caches Lucene indexes configuration files plugins 	 attachments avatars / profile pictures icons export files import files plugins

If you are currently storing attachments in your database you can continue to do so, but this is not available for new installations.

Caching

When clustered, Confluence uses a combination of local caches, distributed caches, and hybrid caches that are managed using Hazelcast. This allows for better horizontal scalability, and requires less storage and processing power than using only fully replicated caches. SeeCache Statistics for more information.

Because of this caching solution, to minimize latency, your nodes should be located in the same physical location, or region (for AWS and Azure).

Indexes

Each individual Confluence application node stores its own full copy of the index. A journal service keeps each index in sync.

When you first set up your cluster, you will copy the local home directory, including the indexes, from the first node to each new node.

When adding a new Confluence node to an existing cluster, you will copy the local home directory of an existing node to the new node. When you start the new node, Confluence will check if the index is current, and if not, request a recovery snapshot of the index from either the shared home directory, or a running node (with a matching build number) and extract it into the index directory before continuing the start up process. If the snapshot can't be generated or is not received by the new node in time, existing index files will be removed, and Confluence will perform a full re-index.

If a Confluence node is disconnected from the cluster for a short amount of time (hours), it will be able to use the journal service to bring its copy of the index up-to-date when it rejoins the cluster. If a node is down for a significant amount of time (days) its Lucene index will have become stale, and it will request a recovery snapshot from an existing node as part of the node startup process.

If you suspect there is a problem with the index, you can rebuild the index on one node, and Confluence will propagate the new index files to each node in the cluster.

SeeContent Index Administration for more information on reindexing and index recovery.

Cluster safety mechanism

The ClusterSafetyJob scheduled task runs every 30 seconds in Confluence. In a cluster, this job is run on one Confluence node only. The scheduled task operates on asafety number a randomly generated number that is stored both in the database and in the distributed cache used across the cluster. The ClusterSafetyJob compares the value in the database with the one in the cache, and if the value differs, Confluence will shut the node down - this is known as cluster split-brain. This safety mechanism is used to ensure your cluster nodes cannot get into an inconsistent state.

If cluster split-brain does occur, you need to ensure proper network connectivity between the clustered nodes. Most likely multicast traffic is being blocked or not routed correctly.

Balancing uptime and data integrity

By changing how often the cluster safety scheduled job runs and the duration of the Hazelcast heartbeat (which controls how long a node can be out of communication before it's removed from the cluster) you can fine tune the balance between uptime and data integrity in your cluster. In most cases the default values will be appropriate, but there are some circumstances where you may decide to trade off data integrity for increased uptime for example.

Uptime over data integrity

Cluster safety job	Hazelcast heartbeat	Effect
1 minute	1 minute	You could have network interruptions or garbage collection pauses of up to 1 minute without triggering a cluster panic. However, if two nodes are no longer communicating, conflicting data could be being written to the database for up to 1 minute, affecting your data integrity.
10 minutes	30 seconds	You could have network interruptions or garbage collection pauses of up to 30 seconds without nodes being evicted from the cluster. Evicted nodes then have up to 10 minutes to rejoin the cluster before the Cluster Safety Job kicks in and shuts down the problem node. Although this may result in higher uptime for your site, conflicting data could be being written to the database for up to 10 minutes, affecting your data integrity.

Data integrity over uptime

Cluster safety job	Hazelcast heartbeat	Effect
15 seconds	15 seconds	Network interruptions or garbage collection pauses longer than 15 seconds will trigger a cluster panic. Although this may result in higher downtime for your site, nodes can only write to the database while out of communication with each other for a maximum of 15 seconds, ensuring greater data integrity.
15 seconds	1 minute	You could have network interruption or garbage collection pauses up to 1 minute without nodes being evicted from the cluster. Once a node is evicted, it can only write to the database for a maximum of 15 seconds, minimizing the impact on your data integrity.

To find out how to change the cluster safety scheduled job, see Scheduled Jobs.

You can change the Hazelcast heartbeat default via the confluence.cluster.hazelcast.max.no. heartbeat.seconds system property. See Configuring System Properties.

Cluster locks and event handling

Where an action must only run on one node, for example a scheduled job or sending daily email notifications, Confluence uses a cluster lock to ensure the action is only performed on one node.

Similarly, some actions need to be performed on one node, and then published to others. Event handling ensures that Confluence only publishes cluster events when the current transaction is committed and complete. This is to ensure that any data stored in the database will be available to other instances in the cluster when the event is received and processed. Event broadcasting is done only for certain events, like enabling or disabling an app.

Cluster node discovery

When configuring your cluster nodes you can either supply theIP address of each cluster node, or a multicast address.

If you're using multicast:

Confluence will broadcast a join request on the multicast network address. Confluence must be able to open a UDP port on this multicast address, or it won't 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.

A multicast address can be auto-generated from the cluster name, or you can enter your own, during the setup of the first node.

Infrastructure and hardware requirements

The choice of hardware and infrastructure is up to you. Below are some areas to think about when planning your hardware and infrastructure requirements.

AWS Quick Start deployment option

If you plan to run Confluence Data Center on AWS, a Quick Start is available to help you deploy Confluence Data Center in a new or existing Virtual Private Cloud (VPC). You'll get your Confluence and Synchrony nodes, Amazon RDS PostgreSQL database and application load balancer all configured and ready to use in minutes. If you're new to AWS, the step-by-step Quick Start Guide will assist you through the whole process.

Confluence can only be deployed in a region that supports Amazon Elastic File System (EFS). SeeRunning Confluence Data Center in AWS for more information.

It is worth noting that if you deploy Confluence using the Quick Start, it will use the Java Runtime Engine (JRE) that is bundled with Confluence (/opt/atlassian/confluence/jre/), and not the JRE that is installed on the EC2 instances (/usr/lib/jvm/jre/).

Server requirements

You should not run additional applications (other than core operating system services) on the same servers as Confluence. Running Confluence, Jira and Bamboo on a dedicated Atlassian software server works well for small installations but is discouraged when running at scale.

Confluence Data Center can be run successfully on virtual machines. If you plan to use multicast, you can't run Confluence Data Center in Amazon Web Services (AWS) environments as AWS doesn't support multicast traffic.

Cluster nodes

Each node does not need to be identical, but for consistent performance we recommend they are as close as possible. All cluster nodes must:

- be located in the same data center, or region (for AWS and Azure)
- run the same Confluence version on each Confluence node (exceptduring a rolling upgrade)
- run the same Synchrony version on each Synchrony node (if not using managed Synchrony)
- have the same OS, Java and application server version

- have the same memory configuration (both the JVM and the physical memory) (recommended)
- be configured with the same time zone (and keep the current time synchronized). Using ntpd or a similar service is a good way to ensure this.

⚠You must ensure the clocks on your nodes don't diverge, as it can result in a range ofproblems with yourcluster.

How many nodes?

Your Data Center license does not restrict the number of nodes in your cluster. The right number of nodes depends on the size and shape of your Confluence site, and the size of your nodes. See our Confluence Data Center load profiles guide for help sizing your instance. In general, we recommend starting small and growing as you need.

Memory requirements

Confluence nodes

We recommend that each Confluence node has a minimum of 10GB of RAM. A high number of concurrent users means that a lot of RAM will be consumed.

Here's some examples of how memory may be allocated on different sized machines:

RAM	Breakdown for each Confluence node
10GB	 2GB for operating system and utilities 4GB for Confluence JVM (-Xmx 3GB) 2GB for external process pool (2 sandboxes with -Xmx 512MB each) 2GB for Synchrony
16GB	 2GB for operating system and utilities 10GB for Confluence JVM (-Xmx 8GB) 2GB for external process pool (2 sandboxes with -Xmx 512MB each) 2GB for Synchrony

The maximum heap (-Xmx) for the Confluence application is set in the setenv.shor setenv.batfile. The default should be increased for Data Center. We recommend keeping the minimum (Xms) and maximum (Xmx) heap the same value.

The external process pool is used to externalise memory intensive tasks,to minimise the impact on individual Confluence nodes. The processes are managed by Confluence. The maximum heap for each process (sandbox) (-Xmx), and number of processes in the pool, is set using system properties. In most cases the default settings will be adequate, and you don't need to do anything.

Standalone Synchrony cluster nodes

Synchrony is required for collaborative editing. By default, it is managed by Confluence, but you can choose to run Synchrony in its own cluster. SeePossible Confluence and Synchrony Configurations for more information on the choices available.

If you do choose to run your own Synchrony cluster, we recommend allowing 2GB memory for standalone Synchrony. Here's an example of how memory could be allocated on a dedicated Synchrony node.

Physical RAM	Breakdown for each Synchrony node
4GB	 2GB for operating system and utilities 2GB for Synchrony JVM (-Xmx 1GB)

Database

The most important requirement for the cluster database is that it have sufficient connections available to support the number of nodes.

For example, if:

- each Confluence node has a maximum pool size of 20 connections
- each Synchrony node has a maximum pool size of 15 connections (the default)
- you plan to run 3 Confluence nodes and 3 Synchrony nodes

your database server must allow at least 105 connections to the Confluence database. In practice, you may require more than the minimum for debugging or administrative purposes.

You should also ensure your intended database is listed in the currentSupported Platforms. The load on an average cluster solution is higher than on a standalone installation, so it is crucial to use the a supported database.

You must also use a supported database driver. Collaborative editing will fail with an error if you're using an unsupported or custom JDBC driver (ordriverClassNamein the case of a JNDI datasource connection). SeeDatabase JDBC Driversfor the list of drivers we support.

Additional requirements for database high availability

Running Confluence Data Center in a cluster removes the application server as a single point of failure. You can also do this for the database through the following supported configurations:

- Amazon RDS Multi-AZ: this database setup features a primary database that replicates to a standby in a different availability zone. If the primary goes down, the standby takes its place.
- Amazon PostgreSQL-Compatible Aurora: this is a cluster featuring a database node replicating to one
 or more readers (preferably in a different availability zone). If the writer goes down, Aurora will
 promote one of the writers to take its place.

The AWS Quick Start deployment optionallows you to deploy Confluence Data Center with either one, from scratch. If you want to set up an Amazon Aurora cluster with an existing Confluence Data Center instance, refer to Configuring Confluence Data Center to work with Amazon Aurora.

Shared home directory and storage requirements

All Confluence cluster nodes must have access to a shared directory in the same path.NFS and SMB/CIFS shares are supported as the locations of the shared directory. As this directory will contain large amount of data (including attachments and backups) it should be generously sized, and you should have a plan for how to increase the available disk space when required.

Remember me and session timeout

The 'remember me' option is enforced by default in a cluster. Users won't see the 'remember me' checkbox on the login page, and their session will be shared between nodes. See the following knowledge base articles if you need to change this, or change the session timeout.

- How to configure the 'Remember Me' feature in Confluence
- How to adjust the session timeout for Confluence

Load balancers

We suggest using the load balancer you are most familiar with. The loadbalancer needs to support session affinity and WebSockets. This is required for both Confluence and Synchrony. If you're deploying on AWS you'll need to use an Application Load Balancer (ALB).

Here are some recommendations when configuring your load balancer:

- Queue requests at the load balancer. By making sure the maximum number requests served to a
 node does not exceed the total number of http threads that Tomcat can accept, you can avoid
 overwhelming a node with more requests than it can handle. You can check the maxThreads in<inst
 all-directory>/conf/server.xml.
- Don't replay failed idempotent requests on other nodes, as this can propagate problems across all your nodes very quickly.
- Using *least connections* as the load balancing method, rather than *round robin*, can better balance the load when a node joins the cluster or rejoins after being removed.

Many load balancers require a URL to constantly check the health of their backends in order to automatically remove them from the pool. It's important to use a stable and fast URL for this, but lightweight enough to not consume unnecessary resources. The following URL returns Confluence's status and can be used for this purpose.

URL	Expected content	Expected HTTP Status
http:// <confluenceurl>/status</confluenceurl>	{"state":"RUNNING"}	200 OK

HTTP Status Code	Response entity	Description
200	{"state":" RUNNING"}	Running normally
500	{"state":" ERROR"}	An error state
503	{"state":" STARTING"}	Application is starting
503	{"state":" STOPPING"}	Application is stopping
200	{"state":" FIRST_RUN"}	Application is running for the first time and has not yet been configured
404		Application failed to start up in an unexpected way (the web application failed to deploy)

Here are some recommendations, when setting up monitoring, that can help a node survive small problems, such as a long GC pause:

- Wait for two consecutive failures before removing a node.
- Allow existing connections to the node to finish, for say 30 seconds, before the node is removed from the pool.

Network adapters

Use separate network adapters for communication between servers. Cluster nodes should have a separate physical network (i.e. separate NICs) for inter-server communication. This is the best way to get 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.

Additional requirements for collaborative editing

Collaborative editing in Confluence 6.0 and later is powered by Synchrony, which runs as a seperate process.

If you have a Confluence Data Center license, two methods are available for running Synchrony:

- managed by Confluence(recommended)
 Confluence will automatically launch a Synchrony process on the same node, and manage it for you.
 No manual setup is required.
- Standalone Synchrony cluster (managed by you)
 You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need.
 Significant setup is required. During a rolling upgrade, you'll need to upgrade the Synchrony separately from the Confluence cluster.

If you want simple setup and maintenance, we recommend allowing Confluence to manage Synchrony for you. If you want full control, or if making sure the editor is highly available is essential, then managing Synchrony in its own cluster may be the right solution for your organisation.

App compatibility

The process for installing Marketplace apps (also known as add-ons or plugins) in a Confluence cluster is the same as for a standalone installation. You will not need to stop the cluster, or bring down any nodes to install or update an app.

The Atlassian Marketplace indicates apps that are compatible with Confluence Data Center.

If you have developed your own plugins (apps) for Confluence you should refer to our developer documentation on How do I ensure my app works properly in a cluster? to find out how you can confirm your app is cluster compatible.

Ready to get started?

Head to Set up a Confluence Data Center cluster for a step-by-step guide to enabling and configuring your cluster.

External Process Pool for Confluence Data Center

In Confluence Data Center we minimize the impact of particularly memory or CPU intensive actions byhandling them in an external process pool, which is a seperate pool of processes, managed by Confluence. These processes (also known as sandboxes) can crash or be terminated, and will be restarted automatically by Confluence, without affecting the Confluence application itself.

On this page:

- Memory requirements
- Configure the external process pool
- Monitor failed actions

The external process pool currently handles the following actions:

- Document conversion(thumbnail generation for file previews)
- Exporting a space to PDF



The external process pool is only available for Confluence Data Center.

In Confluence Server, these actions are handled by Confluence, so the information on this page does not apply.

Memory requirements

You will need to make sure that Confluence has enough memory for the external process pool. In a clustered Data Center installation, you'll need to do this for each cluster node. The pool contains two processes (sandboxes) by default, sowe recommend allowing an additional 2 GB on top of what is already required for Confluence (1 GB per sandbox).

If you increase the size of the external process pool, make sure each node has enough free memory to cater for the extra processes.

Configure the external process pool

In most cases the default values will be adequate, however system administrators can configure the external process pool usingsystem properties. For example you may want to increase the size of the pool (the number of processes available), or increase the amount of memory a process can consume. Here are the main properties you may need to change:

- conversion.sandbox.pool.size
 - Use this property to increase the number of processes (sandboxes) in the pool. You'll need to allow additional memory on each node for each additional process.
- conversion.sandbox.memory.limit.megabytes Use this property to limit the amount of memory each process (sandbox) in the pool can consume.

SeeRecognized System Properties for a full description of these properties, including additional properties that can be used to fine-tune, or disable sandboxes for particular actions.

Monitor failed actions

When an external process (sandbox) is terminated, we'll write the following to the application log on that node:

```
2018-04-09 17:35:35 WARN [sandbox-terminator]
[impl.util.sandbox.DefaultSandbox] lambda$startTerminator$0 Request
has taken 33384ms exceeds limit 30000ms terminating sandbox
```

This will be followed by anAttempting to restart the sandboxmessage, the next time someone performs an action that uses the external process pool.

Note that the process is not immediatley restarted after termination, as we don't re-attempt failed actions. We wait for the next request to spin up a new sandbox process.

Document conversion for Confluence Data Center

When you insert a file into a page (for example a Word document, or Excel spreadsheet), Confluence will convert the contents to a format that can be viewed inline in the page, in the preview, or in some macros. This can be quite memory and CPU intensive, and has been known to cause out of memory errors when processing very complex files.

In Confluence Data Center we minimize the impact byhandling the conversion in an external process pool, which is a seperate pool of processes, managed by Confluence. These processes (also known as sandboxes) can crash or be terminated, and will be restarted automatically by Confluence, without affecting the Confluence application itself.

For example, If you insert a very complex file, and the process crashes or is terminated, thumbnail generation will fail. When this happens, a placeholder thumbnail will be used on the page, and a download option will be provided in the file preview. Confluence Data Center doesn't re-attempt to generate thumbnails for failed files. A good example of a complex file, is a PowerPoint presentation that contains 50 embedded Excel charts. Most files will be processed without any problems.

The external process pool is used for the following conversions:

- thumbnail generation for images and documents inserted into a page, or viewed in the preview.
- HTML conversion for Word and Office documents viewed using the Office Word and Office Excel macros.



The external process pool is only available for Confluence Data Center.

In Confluence Server, thumbnail generation and HTML conversion is handled by Confluence, so the information on this page does not apply.

Configure the external process pool

In most cases the default values will be adequate, however system administrators can change the behaviour using system properties. For example you may want to increase the size of the pool (the number of processes available), or increase the time limit before a process is terminated. Here are the main properties you may need to change:

- conversion.sandbox.pool.size
 - Use this property to increase the number of processes (sandboxes) in the pool. You'll need to allow additional memory on each node for each additional process.
- conversion.sandbox.memory.limit.megabytes Use this property to limit the amount of memory each thumbnail generation process in the pool can consume.
- document.conversion.sandbox.memory.requirement.megabytes Use this property to limit the amount of memory each HTML conversion process in the pool can consume.
- document.conversion.sandbox.request.time.limit.secs Use this property to change the amount of time (in seconds) that the sandbox will wait for the conversion process to complete, before terminating the process.

SeeRecognized System Properties for a full description of these properties, plus a few additional properties that can be used to fine-tune, or disable the sandboxes completley.

Re-attempt thumbnail generation for failed files

Confluence does not re-attempt to generate thumbnails for a failed attachment, and re-inserting the attached file into the editor will not trigger the process.

If you do want to re-attempt thumbnail generation, for example after increasing the request time limit, you will need to re-upload the file, and then re-insert it into the page.

Other system properties that affect document conversion

The system properties listed on this page apply specifically to the external process pool. However there are some additional properties that apply in both Confluence Server and Confluence Data Center:

- confluence.document.conversion.imaging.enabled.tif
 Use this property to enable document conversion for TIFF files. This is disabled by default.
- confluence.document.conversion.imaging.enabled.psd
 Use this property to enable document conversion for Photoshop PSD files. This is disabled by default.
- confluence.document.conversion.imaging.convert.timeout
 Use this property to change the default 30 second time limit which applies when performing document conversion on complex image files (such as ICO, EMF, WMF).
- confluence.document.conversion.slides.convert.timeout
 Use this property to change the default 30 second time limit which applies when performing document conversion on presentation files (such as PPT, PPTX).

To override the default value of these properties, you'll need to use the conversion.sandbox.java. optionssystem property to pass the property to the JVMs that make up the external process pool.

In this example we'll enable thumbnail generation for TIFF and PSD files.

- 1. Edit the <install-directory>/bin/setenv.bat file.
- 2. Add the following lines

```
set CATALINA_OPTS=-Dconversion.sandbox.java.options=-Dconfluence.document.conversion.imaging.enabled.tif=true -Dconfluence.document.conversion.imaging.enabled.psd=true %CATALINA_OPTS%
```

You can pass multiple properties to the external process pool JVMs this way.

If you're running Confluence as a Windows Service or on AWS, see Configuring System Properties for how to add this property.

In this example we'll enable thumbnail generation for TIFF and PSD files.

- 1. Edit the <install-directory>/bin/setenv.sh file.
- 2. Add the following lines. In this example we're enabling document conversion for TIFF and PSD files.

```
CATALINA_OPTS="-Dconversion.sandbox.java.options=-Dconfluence.document.conversion.imaging.enabled.tif=true -Dconfluence.document.conversion.imaging.enabled.psd=true ${CATALINA_OPTS}"
```

You can pass multiple properties to the external process pool JVMs this way.

If you're running Confluence on AWS, see Configuring System Properties for how to add this property.

If you decide to increase the timeout for generating thumbnails in the external process pool using thedocument .conversion.sandbox.request.time.limit.secssystem property,you may also want to change the timeout for complex image files or presentations using the system properties above. Alternatively you could keep the default, and allow these types of files to fail sooner.

PDF export in Confluence Data Center

When you export a space to PDF, Confluence exports the content of each page to HTML, converts that HTML to PDF, and then finally merges all the pages together into a single PDF file. This can be guite memory and CPU intensive, andhas been known to cause out of memory errors when processing spaces with very long or complex pages.

In Confluence Data Center we minimize the impact byhandling the export in an external process pool, which is a seperate pool of processes, managed by Confluence. These processes (also known as sandboxes) can crash or be terminated, and will be restarted automatically by Confluence, without affecting the Confluence application itself.



The external process pool is only available for Confluence Data Center.

In Confluence Server, PDF export is handled by Confluence, so the information on this page does not apply.

Troubleshooting failed exports

Exporting an entire space to PDF can sometimes fail, especially if the space is very large, or has very long or complex pages. If PDF export fails you'll see one of the following errors in your browser.

Page took too long to convert

This error occurs when the time it takes to convert the HTML of a page to PDF exceeds the set time limit. The page title will be included in the error message.

You should take a look at the page, and see if it can be simplified. It might have a lot of complex macros, or a lot of web images (images that are not attached to the page). If this error happens a lot, you can ask your admin toin crease the time limit.

Error converting page to HTML

This error occurs when Confluence runs out of memory, or hits another error while trying to convert the HTML of a page to PDF. The page title will be included in the error message.

As with the 'page took too long to convert' error above, you should take a look at the page, and see if it can be simplified.

Confluence admins can get more information about the cause of these errors from the Confluence application logs. If the failures are being caused by out of memory errors, your admin may be able to increase the amount of memory available to each sandbox in the external process pool. SeeExternal Process Pool for Confluence Data Center for more information.

Final PDF file wasn't merged in time

This error occurs at the last stage of the process, when the time it took to stitch together all the individual page PDFs into one PDF file, exceeds the set time limit.

If you hit this error you could try exporting the space again, or perhaps export the space in two sections (using the custom option on the PDF export screen). If this error happens a lot, you can ask your admin to increase the time limit.

Error merging the final PDF file

This error occurs when Confluence runs out of memory, or hits another error, when attempting to stitch together all the individual page PDFs into one file.

If you hit this error you could try exporting the space again, or perhaps export the space in two sections (using the custom option on the PDF export screen).

Confluence admins can get more information about the cause of these errors from the Confluenceapplication logs. If the failures are being caused by out of memory errors, they may be able to increase the amount of memory available to each sandbox in the external process pool. SeeExternal Process Pool for Confluence Data Centerfor more information.

Too many concurrent exports

This error occurs when multiple people are exporting to PDF at the same time. Confluence limits the number of PDF exports that can be processed concurrently.

If you hit this error, try exporting the space again later, after the other PDF exports have completed.

If this error happens a lot, your admin can increase the maximum number of concurrent PDF exports, or increase the time Confluence should wait when the maximum number of concurrent PDF exports has been reached using the following system properties:

confluence.pdfexport.permits.size

Use this property to set the maximum number of concurrent PDF exports that can be performed. This property applies per node, not per sandbox process.

confluence.pdfexport.timeout.seconds

Use this property to set the amount of time a new PDF export request should wait before failing, if the maximum number of concurrent PDF exportshas already been reached.

Change the time limit

Processes are automatically terminated once a time limit is exceeded. You can increase the time limit for PDF export using the following system property:

pdf.export.sandbox.request.time.limit.secs

Use this property to set the amount of time (in seconds) that a process should wait to complete, before being terminated. This time limit applies both to the time to convert the content from HTML to PDF, and the time to merge the final PDF file.

SeeRecognized System Propertiesfor a full list of properties, including a few additional properties that can be used to fine-tune, or disable the sandboxes for a particular action.

Don't use the external process pool for PDF export

If you don't want to use the external process pool for PDF export, you can revert back to the Confluence Server method of generating PDF exports using the following system property:

pdf.export.sandbox.disable

Set this property to true if you don't want to handle PDF exports in the external process pool.

Restricted Functions in Confluence Data Center

There are some features that are disabled or limited in clustered Confluence Data Center installations. This is to ensure the integrity and performance of your cluster.

The current restricted functions are:

Restricted function	Data Center Status	Explanation
Workbox plugins	Available from 5.7	The workbox provides notifications collected from Confluence page watches, shares, and mentions. This is disabled in Confluence Data Center 5.6 to ensure notifications are correctly handled across the cluster. Disabled plugins included Workbox common plugin, Workbox Jira provider plugin, Workbox confluence provider plugin, Workbox host plugin. You will not be able to enable these plugins in the universal plugin manager.
Confluence Quick Reload Plugin	Available from 5.6.3	The quick reload function notifies users when a new comment has been added to a page they are currently viewing. New comment from Rachel Admin Show Ignore New edit by Rachel Admin Reload Ignore This is disabled in Confluence Data Center 5.6 and 5.6.1 for performance reasons. You will not be able to enable the Confluence Quick Reload Plugin in the universal plugin manager. See CONFSERVER-34680 - Make quick reload plugin available in Confluence Data Center CLOSED for more info.
Application links authenticati on: Basic access (http) Trusted Applicati ons	RESTRICTED	When creating Application links to other applications (for example Jira) Basic HTTP and Trusted Applications authentication is not supported for Confluence Data Center. All application links must use OAuth authentication in a cluster.
Confluence Usage Stats plugin	DISABLED	The Confluence Usage Stats plugin provides space activity information for a space (statistics). This is disabled by default in Confluence Server and should not be enabled in Confluence Data Center.
Scheduled jobs history and status	LIMITED	On the Scheduled Jobs page in the Confluence Data Center administration console you will not be able to access the last execution time or history for each job. The page will also only show the configured status (scheduled or disabled) of each job, and will not indicate when a job is in progress.

Remember me on by default



Remember me on the log in page is enabled by default (and does not appear) to allow users to move seamlessly between nodes. You can use the cluster.login.rememberme.enabled system property to override the default and show the checkbox - users will be prompted to log in to another node if their current node is unavailable.

Set up a Confluence Data Center cluster

Confluence Data Center allows you to run a cluster of multiple Confluence nodes, providing high availability, scalable capacity, and performance at scale.

This guides walks you through the process of configuring a Data Center cluster on your own infrastructure.

Youll need to be logged in as a System Administrator to do this.

Not sure if clustering is right for you? Check out Clustering with Confluence Data Center for a detailed overview.

On this page

- Clustering with AWS and Azure
- Before you begin
- Set up and configure your cluster
- Add more Confluence nodes
- Troubleshooting
- Were here to help

Clustering with AWS and Azure

You can also choose to deploy a Data Center cluster on public cloud providers, like AWS (Amazon Web Services) and Azure. We have specific guides and deployment templates to help you easily configure a cluster in A WS or Azure. Check them out to find out what's required.

Before you begin

Clustering requirements

To use Confluence Data Center you must:

- Have a Data Center license (you canpurchase a Data Center licenseo r create an evaluation license atmy.atlassian.com)
- Use asupportedexternal database, operating system and Java version
- Use OAuth authentication if you haveapplication linksto other Atlassian products (such as Jira)

To run Confluence in a cluster you must also:

- Use a load balancer with session affinity in front of the Confluence cluster. WebSockets support is also recommended for collaborative editing.
- Have a shared directory accessible to all cluster nodes in the same path (this will be your shared home directory). This must be a separate directory, and not located within the local home or install directory.

See Clustering with Confluence Data Center for a complete overview of hardware and infrastructure considerations.

Security

Ensure that only permittedcluster nodes are allowed to connect to the following portsthrough the use of a firewall and / or network segregation:

- 5801 -Hazelcastport for Confluence
- 5701 -Hazelcastport for Synchrony
- 25500 Cluster base port for Synchrony

If you use multicast for cluster discovery:

 54327- Multicast port for Synchrony (only required if running Synchrony standalone cluster)

Terminology

In this guide we'll use the following terminology:

- Installation directory The directory where you installed Confluence.
- Local home directory The home or data directory stored locally on each cluster node (if Confluence is not running in a cluster, this is simply known as the home directory).
- Shared home directory The directory you created that is accessible to all nodes in the cluster via the same path.

Set up and configure your cluster

We recommend completing this process in a staging environment, and testing your clustered installation, before moving to production.

1. Back up

We strongly recommend that you backup your existing Confluence local home and install directories and your database before proceeding.

You can find the location of your home directory in the <installation -directory>/confluence/WEB-INF/classes/confluence-init.properties file.

This is where your search indexes and attachments are stored. If you store attachments outside the Confluence Home directory, you should also backup your attachments directory.

2. Create a shared home directory

- 1. Create a directory that's accessible to all cluster nodes via the same path. This will be your **shared home** directory.
- 2. In your existing Confluence home directory, move the contents of <lo cal home directory>/shared-hometo the new shared home directory you just created. To prevent confusion, we recommend deleting the empty <local home directory>/shared-homedire ctory once you've moved its contents.
- 3. Move your <local home>/attachments>directory to the new <sh ared home>/attachmentsdirectory.

4. Enable cluster mode

Before you enable cluster mode, you should be ready to restart Confluence and configure your cluster. This will require some downtime.

- 1. Start Confluence.
- 2. Go to Configuration.
- 3. Choose **Clustering** from the sidebar.
- 4. Select Enable cluster mode.
- 5. Select **Enable** to confirm youre ready to proceed.

5. Restart Confluence

Restart Confluence to configure your cluster. Once you restart, Confluence will be unavailable until youve completed the set up process.

6. Configure your cluster

The setup wizard will prompt you to configure the cluster, by entering:

- A name for your cluster
- The path to the shared home directory you created earlier
- The network interface Confluence will use to communicate between nodes
- How you want Confluence to discover cluster nodes:
 - Multicast enter your own multicast address or automatically generate one.
 - TCP/IP enter the IP address of each cluster node
 - AWS enter your IAM Role or secret key, and region.

We recommend using our Quick Start or Cloud Formation Template to deploy Confluence Data Center in AWS, as it will automatically provision, configure and connect everything you need.

If you do decide to do your own custom deployment, you can provide the following information to allow Confluence to auto-discover cluster nodes:

Field	Description
IAM Role or Secre t Key	This is your authentication method. You can choose to authenticate by IAM Role or Secret Key.
Region	This is the region your cluster nodes (EC2 instances) will be running in.
Host head er	Optional. This is the AWS endpoint for Confluence to use (the address where the EC2 API can be found, for example 'ec2.amazonaws. com'). Leave blank to use the default endpoint.
Secur ity grou p name	Optional. Use to narrow the members of your cluster to only resources in a particular security group (specified in the EC2 console).
Tag key and T ag value	Optional. Use to narrow the members of your cluster to only resources with particular tags (specified in the EC2 console).

If Synchrony is managed by Confluence, the same network settings will be applied to Synchrony.

Follow the prompts to create the cluster.

When you restart, Confluence will start setting up the cluster. This can take a few minutes. Some core components of Confluence will also change to become cluster compatible. For example, Confluence will switch to a distributed caching layer, managed by Hazelcast.

⚠Do not restart Confluence until your cluster is set up, and Confluence is back up and running.

Add more Confluence nodes

Your Data Center license doesnt restrict the number of nodes in your cluster. To achieve the benefits of clustering, such as high availability, youll need to add at least one additional cluster node.

Weve found that typically between 2 and 4 nodes is sufficient for most organizations. In general we recommend starting small and growing as needed.

7. Copy Confluence to the second node

To copy Confluence to the second node:

- 1. Shut down Confluence on node 1.
- 2. Copy the installation directory from node 1 to node 2.
- 3. Copy the local home directory from node 1 to node 2.

Copying the local home directory ensures the Confluence search index, the database and cluster configuration, and any other settings are copied to node 2.

Copying the local home directory ensures the Confluence search index, the database and cluster configuration, and any other settings are copied to node 2.

Make sure your database has sufficient connections available to support the number of nodes.

8. Configure your load balancer

Configure your load balancer for Confluence. You can use the load balancer of your choice, but it needs to support session affinity and WebSockets.

You can verify that your load balancer is sending requests correctly to your existing Confluence server by accessing Confluence through the load balancer and creating a page, then checking that this page can be viewed /edited by another machine through the load balancer.

See Clustering with Confluence Data Center for further load balancer guidance.

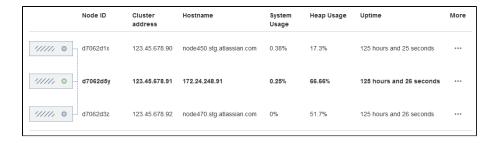
9. Start Confluence one node at a time

You must only start Confluence **one node at a time**. The first node must be up and available before starting the next one.

- 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.

TheCluster monitoring console (Seneral Configuration > Clustering) shows information about the active cluster.

When the cluster is running properly, this page displays the details of each node, including system usage and uptime. Use the *** menu to see more information about each node in the cluster.



10. Test your Confluence cluster

To test creating content you'll need to access Confluence via your load balancer URL. You can't create or edit pages when accessing a node directly.

A simple process to ensure your cluster is working correctly is:

- Access a node via your load balancer URL, and create a new document on this node.
- Ensure the new document is visible by accessing it directly on a different node.
- 3. Search for the new document on the original node, and ensure it appears.
- Search for the new document on another node, and ensure it appears.
- i 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.

11. Set up a Synchrony cluster (optional)

Synchrony is required for collaborative editing. You have two options for running Synchrony with a Data Center license:

- managed by Confluence(recommended)
 This is the default setup. Confluence will automatically launch a Synchrony process on the same node, and manage it for you. No manual steps are required.
- Standalone Synchrony cluster (managed by you)
 You deploy and manage Synchrony standalone in its own cluster
 with as many nodes as you need. Significant setup is required. See S
 et up a Synchrony cluster for Confluence Data Center for a step-by step guide.

Head to Administering Collaborative Editing to find out more about collaborative editing.

Troubleshooting

If you have problems with the above process, check our cluster troubleshooting guide.

Were here to help

Need help setting up your cluster? There are a range of support services available to help you plan and implement a clustered Data Center installation.

- An Atlassian Technical Account Manager can provide strategic guidance. They work with you to develop best practices for configuring, deploying and managing Confluence in a cluster.
- The Atlassian Premier Support team can provide technical support.
 Premier Support also offers health check analyses to validate the readiness of your environment.
- Atlassian Enterprise Partners offers a wide array of services to help you get the most out of your Atlassian tools.
- You can also ask questions in the Atlassian Community.

Confluence Data Center Performance

This document describes the performance tests we conducted on clustered Confluence Data Center within Atlassian, and the results of those tests. You can compare these data points to your own implementation to predict the type of results you might expect from implementing Confluence Data Center in a cluster in your own organization.

We started our performance tests by taking a fixed load profile (read/write ratio), then tested different cluster set ups against multiples of that load profile.

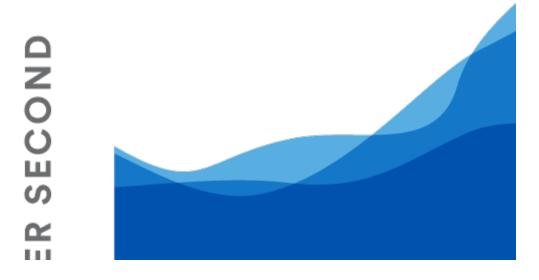
On this page

- Testing results summary
- Testing methodology and specifications
 - How we tested
 - What we tested
 - Hardware
- Comparison to Confluence Server response times

Testing results summary

Performance gains - Under a high load, clustered Confluence has improved performance overall.

Request responses don't diminish under increased load -Adding more nodes increases throughput, handles higher load and decreases response times.





CONCURRENT BROV

You might observe a different trend/behavior based on your configuration and usage. For details, please see the **What we tested**se ction below.

Testing methodology and specifications

The following sections detail the testing environment and methodology we used in our performance tests.

How we tested

Our performance tests were all run on the same controlled isolated lab at Atlassian. For each test, the entire environment was reset and rebuilt. The testing environment included the following components and configuration:

- Apache proxy_balancer
- Postgres database and the required data
- G1GC garbage collector
- 8GB Xmx settings per node
- 6 CPUs per node
- Confluence Server on one machine or Confluence Data Center on two, or four machines as required for the specific test.

To run the test, we used a number of machines in the lab to generate load using scripted browsers and measuring the time taken to perform an action. An action here, means a complete user operation like creating a page or adding comment. Each browser was scripted to perform an action from a predefined list of actions and immediately move on the to next action (i.e. zero think time). Please note that this resulted in each browser performing more tasks than would be possible by a real user and you should not interpret the number of browsersto be equal to the number of real world users. Each test was run for 20 minutes, after which statistics were collected.

What we tested

All tests used the same Postgres database containing the same number of spaces and pages.

The mix of actions we included in the tests represented a sample of the most common user actions*
representing six typical types of users (personas). The table below showthe ratio of actionsperformed
by each of these personas. These user-based actions were repeated until the test was completed.

Persona	Ratio of actions
PageReader	7
Searcher	1
Editor	1
Creator	1
Commenter	1
Liker	1

Tests were performed with differing load sizes, from 4 up to 96 browsers. For larger load sets, profiles were scaled up, that is, doubling each amount for the 24 browser load, tripled for the 36 browser load.

Hardware

All performance tests were all run on the same controlled, isolated lab at Atlassian using the hardware listed below.

Hardware	Description	How many?
Rackform iServ R304.v3	CPU: 2 x Intel Xeon E5-2430L, 2.0GHz (6-Core, HT, 15MB Cache, 60W) 32nm	20
	RAM: 48GB (6 x 8GB DDR3-1600 ECC Registered 2R DIMMs) Operating at 1600 MT/s Max	
	NIC: Dual Intel 82574L Gigabit Ethernet Controllers - Integrated	
	Controller: 8 Ports 3Gb/s SAS, 2 Ports 6Gb/s SATA, and 4 Ports 3Gb/s SATA via Intel C606 Chipset	
	PCIe 3.0 x16: Intel X540-T2 10GbE Dual-Port Server Adapter (X540) 10GBASE-T Cat 6A - RJ45	
	Fixed Drive: 240GB Intel 520 Series MLC (6Gb/s) 2.5" SATA SSD	
	Power Supply: 600W Power Supply with PFC - 80 PLUS Gold Certified	
Arista DCS- 7050T-36-R	4PORT SFP+ REAR-TO-FRONT AIR 2XAC	1
HP ProCurve Switch	1810-48G 48 Port 10/100/1000 ports Web Managed Switch	1

Hardware testing notes:

• In order to quickly put more stress on the Confluence nodes with less load, cluster nodes were set to use only 4 cores out of 6 from each CPU, thereby reducing its processing power.

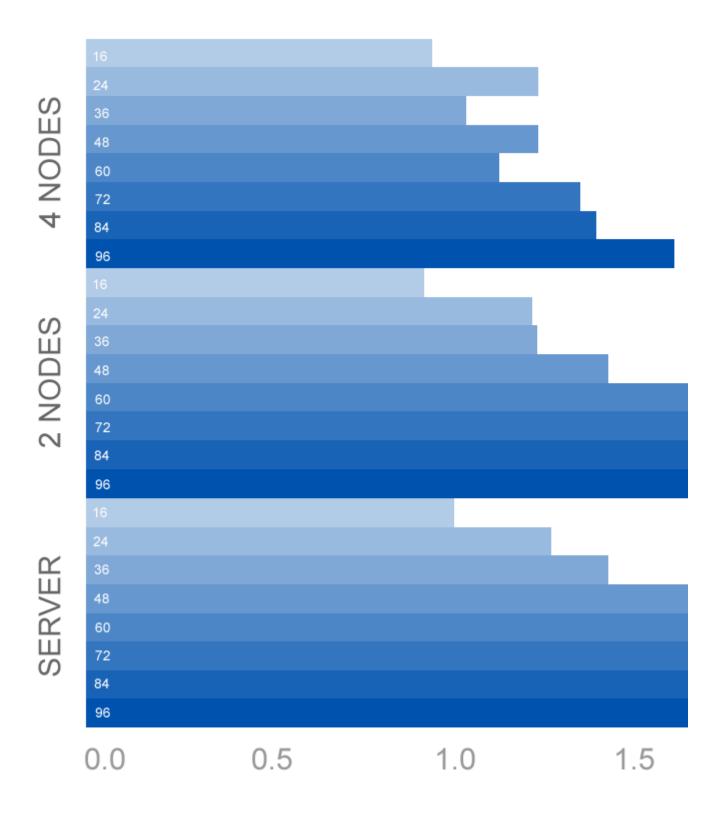
^{*}The tests did not include admin actions as these are assumed to be relatively infrequent.

- For instances being tested, 6 GB of memory was allocated to the JVM consistently across all tests.
 This may not be optimized for all cases but allowed for consistency and comparability between the tests
- During the tests we did not observe high CPU or IO load on either the database or load balancer servers.
- During the tests we did not observe running out of HTTP connections in the load balancer or connections to database.
- The browser and servers are in the same location so there was very low latency between client and server.

Comparison to Confluence Server response times

The following table shows the relative performance as the load increases for each Confluence instance configuration: Confluence Server, two node Confluence Data Center, and four node Confluence Data Center. The table shows the response time relative to the baseline response time which we determined to be Confluence Server with sixteen browsers.

Browsers	16	24	36	48	60	72	84	96
Server	100.00%	125.28%	142.95%	222.76%	276.54%	334.79%	393.03%	451.28%
2 Node	93.79%	122.61%	123.50%	141.98%	168.47%	201.97%	235.47%	268.97%
4 Node	94.24%	122.22%	103.94%	123.47%	114.76%	134.61%	138.90%	160.95%



Ready to get started?

Contact usto speak with an Atlassian orget going with Data Centerstraight away.

For a detailed overview of Confluence's clustering solution see Clustering with Confluence Data Center. For help with installation, take a look at Installing Confluence Data Center.

Confluence Data Center disaster recovery

A disaster recovery strategy is a key part of any business continuity plan. It outlines the processes to follow in the event of a disaster, to ensure that the business can recover and keep operating. For Confluence, this means ensuringConfluence'savailability in the event that your primary site becomes unavailable.

Confluence Data Center documentationis the only Atlassian-supported high-availability solution for Confluence.

Not sure if you should upgrade from Confluence Server to Data Center? Learn more about thebenefits of Confluence Data Center.

This page demonstrates how you can useConfluence Data Center5.9 or laterin implementing and managing a disaster recovery strategy for Confluence. It doesn't, however, cover the broader business practices, like setting the key objectives (RTO, RPO & RCO¹), and standard operating procedures.

What's the difference between high availability and disaster recovery?

The terms "high availability", "disaster recovery" and "failover" can often be confused. For the purposes of this page, we've defined them as follows:

- High availability A strategy to provide a specific level of availability. In Confluence's case, access to the application and an acceptable response time. Automated correction and failover (within the same location) are usually part of high-availability planning.
- Disaster recovery A strategy to resume operations in an alternate data center (usually in another geographic location), if the main data center becomes unavailable (i.e. a disaster). Failover (to another location) is a fundamental part of disaster recovery.
- **Failover** is when one machine takes over from another machine, when the aforementioned machines fails. This could be within the same data center or from one data center to another. Failover is usually part of both high availability and disaster recovery planning.

Overview

Before you start, you needConfluence Data Center 5.9 or laterto implement the strategy described in this guide. We'll also assume you've already set up and configured your cluster. See Set up a Confluence Data Center cluster.

This page describes what is generally referred to as a 'cold standby'strategy, which means the standby Confluence instance isn't continuously running and that you need to take some administrative steps to start the standby instance and ensure it's in a suitable state to service the business needs of your organization.



Maintaining a runbook

The detailed steps will vary from organization to organization and, as such, we recommend you keep a full runbook of steps on file, away from the production system it references. Make your runbook detailed enough such that anyone in the relevant team should be able to complete the steps and recover your service, regardless of prior knowledge or experience. We expect any runbook to contain steps that cover the following parts of the disaster recovery process:

- 1. Detection of the problem
- 2. Isolation of the current production environment and bringing it down gracefully
- 3. Synchronization of data between failed production and intended recovery point
- 4. Warm up instructions for the recovery instance
- 5. Documentation, communication, and escalation guidelines

The major components you need to consider in your disaster recovery plan are:

Confluence installation

Your standby site should have exactly the same version of Confluence installed as your production site.

Database	This is the primary source of truth for Confluence and contains most of the Confluence data (except for attachments, avatars, etc). You need to replicate your database and continuously keep it up to date to satisfy your RPO1
Attachments	All attachments are stored in the Confluence Data Center shared home directory, and you need to ensure it's replicated to the standby instance.
Search Index	The search index isn't a primary source of truth, and can always be recreated from the database. For large installations, though, this can be quite time consuming and the functionality of Confluence will be greatly reduced until the index is fully recovered. Confluence Data Center stores search index backups in the shared home directory, which are covered by the shared home directory replication.
Plugins	User installed plugins are stored in the database and are covered by the database replication.
Other data	A few other non-critical items are stored in the Confluence Data Center shared home. Ensure they're also replicated to your standby instance.

Set up a standby system

Step 1. InstallConfluence Data Center5.9or higher

Install the same version of Confluence on your standby system. Configure the system to attach to the standby database.



DO NOT start the standby Confluence system

Starting Confluence would write data to the database and shared home, which you do not want to do.

You may want to test the installation, in which case you should temporarily connect it to a different database and different shared home directory and start Confluence to make sure it works as expected. Don't forget to update the database configuration to point to the standby database and the shared home directory configuration to point to the standby shared home directory after your testing.

Step 2. Implement a data replication strategy

Replicating data to your standby location is crucial to a cold standby failover strategy. You don't want to fail over to your standby Confluence ins tance and find that it's out of date or that it takes many hours to re-index.

Database	All of the following Confluence supported database suppliers provide their own database replication solutions:
	 Oracle:http://www.oracle.com/technetwork/database/features/data-integration/index.html PostgreSQL:https://wiki.postgresql.org/wiki/Binary_Replication_Tutorial MySQL:http://dev.mysql.com/doc/refman/5.7/en/replication.html Microsoft SQL Server:http://msdn.microsoft.com/en-us/library/ms151198.aspx
	You need to implement a database replication strategythat meets your RTO, RPO and RCO ¹ .
Files	You also need to implement a file server replication strategy for the Confluence shared home directory that meets your RTO, RPO and RCO ¹ .

Clustering considerations

For your clustered environment you need to be aware of the following, in addition to the information above:

Standby cluster	There's no need for the configuration of the standby cluster to reflect that of the live cluster. It may contain more or fewer nodes, depending on your requirements and budget. Fewer nodes may result in lower throughput, but that may be acceptable depending on your circumstances.
File locations	Where we mention <confluencesharedhome> as the location of files that need to be synchronized, we're referring to the shared home for the cluster. <confluencelocalhome> refers to the local home of the node in the cluster.</confluencelocalhome></confluencesharedhome>
Starting the standby cluster	It's important to initiallystart only one node of the cluster, allow it to recover the search index, and check it's working correctly before starting additional nodes.

Disaster recovery testing

You should exercise extreme care when testing any disaster recovery plan. Simple mistakes may cause your live instance to be corrupted, for example, if testing updates are inserted into your production database. You may detrimentally impact your ability to recover from a real disaster, while testing your disaster recovery plan.



The key is to keep the main data center as isolated as possible from the disaster recovery testing.

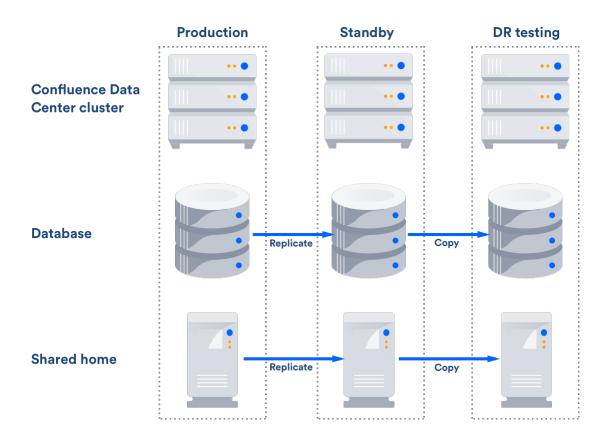


⚠ This procedure will ensure that the standby environment will have all the right data, but as the testing environment is completely separate from the standby environment, possible configuration problems on the standby instance are not covered.

Prerequisites

Before you perform any testing, you need to isolate your production data.

Database	 Temporarily pause all replication to the standby database Replicate the data from the standby database to another database that's isolated
	and with no communication with the main database
Attachments, plugins and indexes	You need to ensure that no plugin updates or index backups occur during the test: 1. Disable index backups 2. Instruct sysadmins to not perform any updates in Confluence 3. Temporarily pause all replication to the standby shared home directory 4. Replicate the data from the standby shared home directory to another directory that's isolated and with no communication with the main shared home directory
Installation folders	 Clone your standby installation separate from both the live and standby instances Change the connection to the database in the <confluencelocalhome> /confluence.cfg.xml file to avoid any conflict</confluencelocalhome> Change the location of the shared home directory in the <confluencelocalhome>/confluence.cfg.xml file to avoid any conflict</confluencelocalhome> If using TCP/IP for cluster setup, change the IP addresses to that of your testing instances in <confluencelocalhome>/confluence.cfg.xml</confluencelocalhome>



After this you can resume all replication to the standby instance, including the database.

Perform disaster recovery testing

Once you have isolated your production data, follow the steps below to test your disaster recovery plan:

- 1. Ensure that the new database is ready, with the latestsnapshot andno replication
- 2. Ensure that the new shared home directory is ready, with the latestsnapshot andno replication
- 3. Ensureyou have a copy of Confluence on a clean server with the right database and shared home directory settings in<confluencelocalhome>/confluence.cfg.xml
- 4. Ensure you have confluence. homemapped, as it was in the standby instance, in the test server
- 5. Disable email (Seeatlassian.mail.senddisabledinConfiguring System Properties)
- 6. StartConfluence

Handling a failover

In the event your primary site is unavailable, you'll need to fail over to your standby system. The steps are as follows:

- 1. Ensure your live system is shutdown and no longer updating the database
- 2. Ensure the contents of <confluencesharedhome > is synced to your standby instance
- 3. Perform whatever steps are required to activate your standby database
- 4. Start Confluenceon one node in the standby instance
- 5. Wait for Confluence to start and check it is operating as expected
- 6. Start up other Confluence nodes
- 7. Update your DNS, HTTP Proxy, or other front end devices to route traffic to your standby server

Returning to the primary instance

In most cases, you'll want to return to using your primary instance after you've resolved the problems that caused the disaster. This is easiest to achieve if you can schedule a reasonably-sized outage window.

You need to:

- Synchronize your primary database with the state of the secondary
- Synchronize the primary shared home directory with the state of the secondary

Perform the cut over

- 1. Shutdown Confluence on the standby instance
- 2. Ensure the database is synchronized correctly and configured to as required
- 3. Use rsync or a similar uililty to synchronize the shared home directory to the primary server
- 4. Start Confluence
- 5. Check that Confluence is operating as expected
- 6. Update your DNS, HTTP Proxy, or other front end devices to route traffic to your primary server

Other resources

Troubleshooting

If you encounter problems after failing over to your standby instance, check these FAQs for guidance:

If your database doesn't have the data available that it should, then you'll need to restore the database from a backup.

Once you've restored your database, the search index will no longer by in sync with the database. You can ei ther do a full re-index, background or foreground, or recover from the latest index snapshot if you have one. This includes the journal id file for each index snapshot. The index snapshot can be older than your database backup; it'll synchronize itself as part of the recovery process.

If the search index is corrupt, you can either do a full re-index, background or foreground, or recover from an earlier index snapshot from the shared home directory if you have one.

You may be able to recover them from backups if you have them, or recover from the primary site if you have access to the hard drives. Tools such as rsync may be useful in these circumstances. Missing attachments won't stop Confluence performing normally; the missing attachments won't be available, but users may be able to upload them again.

Application links are stored in the database. If the database replica is up to date, then the application links will be preserved.

You do, however, also need to consider how each end of the link knows the address of the other:

- If you use host names to address the partners in the link and the backup Confluence server has the same hostname, via updates to the DNS or similar, then the links should remain intact and working.
- If the application links were built using IP addresses and these aren't the same, then the application links will need to be re-established.
- If you use IP addresses that are valid on the internal company network but your backup system is remote and outside the original firewall, you'll need to re-establish your application links.

Definitions

RPO	Recovery Point Objective	How up-to-date you require your Confluence instance to be after a failure.
RTO	Recovery Time Objective	How quickly you require your standby system to be available after a failure.
RCO	Recovery Cost Objective	How much you are willing to spend on your disaster recovery solution.

Data Center Troubleshooting

This page covers troubleshooting for a Data Center installation of Confluence.

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.

⚠You must ensure the clocks on your cluster nodes don't diverge, as it can result in a range ofproblems with yourcluster.

Symptoms

Below is a list of potential problems with Confluence Data Center, and their likely solutions.

On this page:

- Symptoms
- Didn't find a solution?

Related pages:

 Troubleshooting a Data Center cluster outage

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	Check firewall, Check intermediate routers, Increase multicast TTL
App is unlicensed on some nodes after updating the license on one node.	Disable and re-enable the app in the Universal Plugin Manager.
After an app update, strings appear in the UI instead of buttons and icons on some nodes.	Restart the affected node.
Hazelcast CANNOT start on this node. No matching network interface found.	See Hazelcast CANNOT start on this node. No matching network interface found KB article
Any issue not covered here	Contact support

Multicast

Which multicast address?

The multicast address and port used by Confluence can be found on the Cluster Configuration page, or in confluence.cfg.xml in the Confluence home directory.

Multicast address generation.

Confluence uses a hashing algorithm to take the inputted name during setup and it is then turned into a multicast address stored in the config file. Thus, once the initial setup is completed, Confluence will use the address this is the reason why user can change the address if needed, without actually changing the name. Consequently the additional nodes using the same multicast address specified in the config file are able to join the cluster.

Each node has a multicast address configured in the confluence.cfg.xml file

```
name="confluence.cluster.address">xxx.xxx.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
netstat -gn	Lists multicast groups. Does not work on Mac OS X.
netstat -rn	Lists system routing table.
tcpdump -i inte rface	Captures network traffic on the given interface. Most useful on an interface that only receives cluster traffic.

Add multicast route

Multicast networking requirements vary across operating systems. Some operating systems require little configuration, while some require the multicast address to be explicitly added to a network interface before Confluence can use it. If multicast traffic can't be sent or received correctly, adding a route for multicast traffic on the correct interface will often fix the problem. The example below is for a Ubuntu Linux system:

```
route add -net 224.0.0.0 netmask 240.0.0.0 dev eth0
```

To support multiple applications using multicast on different interfaces, you may need to specify a route specific to the Confluence multicast address.

Check firewall

Ensure your firewall allows UDP traffic on the multicast address and port used by Confluence.

Prefer IPv4

There are known issues relating to IPv6. You should configure your JVM to try binding to an IPv4 address first

Change multicast interface

Confluence might have selected the incorrect interface for multicast traffic, which means it cannot connect to other nodes in the cluster. To override the interface used for multicast traffic after initial setup, edit the confluence.cluster.interface property in<local-home>/confluence.cfg.xmland specify the network interface. For example to tell Confluence to use eth1:

cproperty name="confluence.cluster.interface">eth1

Overriding Hazelcast Configuration

If the solution to your problem involves changes to the Hazelcast configuration, these changes should**not**be made to the Confluence configuration files. Instead, to ensure your configuration survives upgrades, make your changes by creating a Hazelcast override file.

Increase multicast TTL

The multicast time-to-live (TTL) specifies how many *hops* a multicast packet should be allowed to travel before it is discarded by a router. It should be set to the number of routers in between your clustered nodes: 0 if both are on the same machine, 1 if on two different machines linked by a switch or cable, 2 if on two different machines with one intermediate router, and so on.

To increase the multicast TTL by edit the confluence.cluster.ttl property in the <local home > /confluence.cfg.xml file on each node. For example to set the TTL to 3:

cproperty name="confluence.cluster.ttl">3</property>

Check intermediate routers

Advanced switches and routers have the ability to understand multicast traffic, and route it appropriately. Unfortunately sometimes this functionality doesn't work correctly with the multicast management information (IGMP) published by the operating system running Confluence.

If multicast traffic is problematic, try disabling advanced multicast features on switches and routers in between the clustered nodes. These features can prevent multicast traffic being transmitted by certain operating systems.

Didn't find a solution?

Check Related Articles from the Confluence Knowledge Base

- Hazelcast CANNOT start on this node. No matching network interface found.
- List of REST APIs available to configure SSO on Confluence Data Center
- Starting Confluence node fails with 'Port [5801] is already in use and auto-increment is disabled.
 Hazelcast cannot start' error
- "Exception bootstrapping cluster:Shared home directory is not configured correctly" Error during Confluence Data Center startup
- Recovering from a Data Center cluster split-brain
- How to suppress cluster warning messages in the Confluence log files
- Cannot find "external_id" column when trying to upgrade to a Confluence CDC license after upgrading from a pre-5.5 Confluence Clustered installation
- Multicast communication works only one-way

- Cluster Panic due to Multicast Traffic Communication Problem
- Configuration of Confluence Cluster Fails with 'Cannot assign requested address'

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.

Troubleshooting a Data Center cluster outage

Confluence Data Center cluster outages can be difficult to troubleshoot as the environments are complex and logging can be very verbose.

This page provides a starting point for investigating outages in your cluster.

Establish the originating node

The most common outage scenario is when something, such as database connectivity issue, network outage or a long garbage collection (GC) process, causes a node to fail to communicate with the cluster for 30 seconds or more and is removed by Hazelcast. The affected node then continues to write to the database, causing a cluster panic.

On this page:

- Establish the originating node
- Investigate common root causes
 - Garbage collection
 - Database connections
 - Network connectivity
- Still having trouble?

To establish the originating node:

- 1. Gather the application log filefrom each node as soon as possible after the outage. Time is critical as the logs will roll over and you may lose the relevant time period.
- 2. Record identifying information about each node to help you interpret the log messages (IP address, node ID and name of each node).
- 3. Make a chronological timeline of the events:
 - a. Record the time that users or monitoring systems started reporting problems.
 - b. View the logs for each node side by side (Hint: we find opening three tabs in node number order helps you always know which logs you are viewing).
 - c. Search the logs for 'removing member' and 'panic'. This will give you a good idea of which nodes caused the issue and when.
 - d. Make a chronological timeline of events from errors to node removal to panics. You can essentially disregard all logging that happens post-panic because once a node panics it needs to be restarted to function effectively. There will be a lot of noise in the logs, but it won't be very useful. The time period we're most interested in will be the minute or so leading up to the first removal or panic event in the logs.

For example:

```
2:50:15 (approx) Node 3 stopped heartbeating to the cluster for 30s (we can estimate this from the time of node removal)
02:50:45 Node 3 was removed by Node 2
02:53:15 Node 4 panics
02:54:15 Node 1, Node 3 and Node 4 receive the panic event and stop processing Node 2 remains serving requests
```

e. When you've established when the first affected node was removed, or when the first cluster panic occurred, look back in time in the logs on that node, to look for root causes.

Investigate common root causes

Once you know when the first affected node was removed you can start investigating root causes. From this point on, you're only looking at events on the affected node around the time of removal (in our example above, this is Node 3 at around 2:50). The subsequent removals and panics are usually flow-on effects of the original node removal event, and aren't likely to provide useful root cause information.

Garbage collection

Check the GC logs for the node that was removed (Node 3 in our example). Were there any GC pauses longer than the Hazelcast heartbeat interval (30 seconds by default)? Nodes can't heartbeat during Garbage Collection, so they will be removed from the cluster by one of the other nodes.

If there was a cluster panic, but the node was not removed from the cluster first, check the GC logs for pauses around the time of the panic - pauses that are relatively short (less than 30 seconds) can sometimes still cause panics (due to a race condition) in Confluence 5.10.1 and earlier.

Database connections

Check any database monitoring tools you may have. How many connections to the database were there at the time of the outage? Heartbeats can fail to send if a node can get a connection from its connection pool but not from the database itself, which can lead to nodes being removed from the cluster.

You won't be able to diagnose this from the Confluence logs and will need to look at any external monitoring tools you have for your database. If the outage happens again, check the current number of connections at the db level during the outage.

Network connectivity

Check your network monitoring tools. If a node drops off the network for a short time and cannot communicate with the cluster, it can be removed by the other nodes. Your load balancer logs may be useful here.

Still having trouble?

ContactSupportfor help troubleshooting these outages. Provide them with as much of the information above as possible, to help their investigation.

Use a CDN with Atlassian Data Center applications

On this page:

· Get started with CDN

How it works

How to determine whether a CDN will help your users

What is cached?

Planning your CDN implementation

- Infrastructure requirements
- Considerations for private instances
- · Marketplace apps and third party customizations

If your users are distributed across the world and experience poor performancewhen using Data Center products, you may be able to improve their experience by using a Content Delivery Network (CDN). Common CDNs include AWS CloudFront, Cloudflare, Akamai, and others.

CDN support is available in Data Centereditions of:

- Jira Software 8.3
- Jira Service Management (formerly Jira Service Desk)4.3
- Confluence 7.0
- Bitbucket 6.8.

Get started with CDN

Here's a quick summary of what's involved to enable your CDN in Confluence Data Center:

- Useour templateto spin up an AWS CloudFront distribution, or create an account with the CDN vendor of your choice.
- 2. Update yourload balancer and firewallto allow the CDN to reach your site.
- 3. In Confluence Data Center, provide the CDN URL, and enable CDN support.

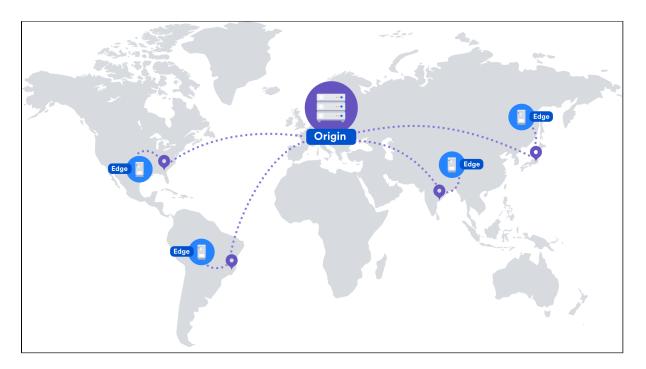
As end users access your site, static assets will be cached on the edge server closest to them, and served from there until they expire. This means it might take some time before you can start measuring the impact of the CDN, depending on when your users are online and accessing the site in each location. We don't provide the ability to preload the cache, so assets will be cached as they are served for the first time.

SeeConfigure your CDN for Confluence Data Centerfor the full step-by-step guide.

As always, we recommend testing this on your staging environment, before making any changes to your production site.

How it works

Static assets (such as JavaScript, CSS, and fonts) are cached on edge servers provided by a CDN vendor that are geographically closer to the user. This means when someone views a page, some of the assets needed to display the page are delivered by a server in their region, rather than from yourserver, known as the origin server. This can speed up page load times.



For example, if your server (known as the origin) is in Germany, a CDN can improve page load time by as much as 50% for users located in Rio de Janeiro, as static assets can be served from an edge server in Brazil.If you're new to CDNs and would like to learn more about how they work, CloudFlare provides a great introduction, seehttps://www.cloudflare.com/learning/cdn/performance/.

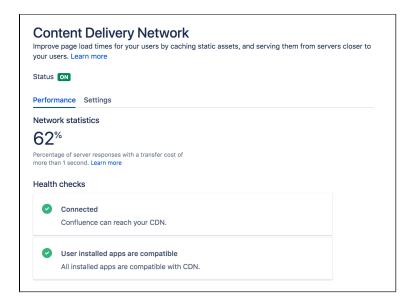
It's important to note that using a CDN will not make your application inherently faster, what it will do is reduce the load on your cluster, and reduce the latency experienced by some users, which should result in faster page load times for users.

Tests on our internal dogfooding instances located in Gdask, Poland have shown the response time for the View Issue action in Jira Data Center is ~50% faster for people accessing from US East, when CDN is enabled.

How to determine whether a CDN will help your users

A good starting point when assessing whether a CDN will help your users, is to take a look at the network overhead experienced in your site.

Go to**Content Delivery Network**in the admin console of your Data Center application. On the **Performance** tab you'll see the percentage of requests that had a transfer cost of more than one second. Put simply, the higher the percentage, the more likely it is that your users requests are being affected by network conditions, such as latency and connection quality.



This network statistic is a useful indicator of the network conditions your users experience when using the product. If the percentage is high, it's likely that using a CDN will benefit your users in these conditions.

As users access pages in your site (for example a Confluence page, Jira issue, or Bitbucket pull request page), we measure the amount of time the browser has to wait to get the content of that page. We then subtract the time required to render the page on the server. This leaves us with the time it took to send the request and retrieve the response.

This time is dependent mostly on the latency between the server and the browser, but also includes things like SSL connection setup time.

This metric is collected on requests that don't use CDN, so it will continue to provide consistent statistics on your network, even after you enable CDN.

You should also consider where your users are geographically located. For example, if your servers are located in Frankfurt, and the majority of your teams are located in Germany and Austria, your team based in Malaysia may be suffering from high latency, resulting in slow page load times.

Network diagnostic tools such as traceroute, ping, and mtrcan be helpful to determine the amount of latency being experienced.

In these examples we'll use traceroute todisplay some basic network statistics, including latency information. Remember to replace yoursite.comwith your base URL.

In Windows, open Command Prompt and enter the following:

```
> tracert yoursite.com
```

In Linux or Mac OS, open Terminal and enter the following:

```
$ traceroute yoursite.com
```

This will display the number of hops, and three latency times, in milliseconds, for each server. Average the three figures to get the latency for that server.

The mtr command (my traceroute) is a useful combination of ping and traceroute. You will need to install mtr to be able to use it in MacOS or Windows.

What is cached?

We only cache static assets served by a Data Center application or Marketplace app. These are things that are only going to change when you upgrade your Data Center application or app.Dynamic content is not cached.

Here's a summary of what will be cached when you enable CDN:

Cached	Not cached
CSSJavaScriptFonts	 attached files pages or issues personal information, including avatars assets that are part of a theme

You shouldn't need to ever manually invalidate the cache, as we handle this when you upgrade your Data Center product, or an app.

Planning your CDN implementation

Infrastructure requirements

You can use any origin pull CDN. You're responsible for any costs associated with your CDN.

We've prepared a CloudFormation template that you can use to configure Amazon CloudFront with minimal effort. You can find all our AWS deployment resources in this repositoryhttps://bitbucket.org/atlassian/atlassian-aws-deployment/src/master/templates/cdn/.

There are some other infrastructure requirements that you need to be aware of before you start:

HTTP/2 is highly recommended

Your load balancer, firewall, or proxy should allow HTTP/2 traffic. Using HTTP/2 will provide the best performance for your end users. Check the documentation for your particular provider to find out how to do this.

Firewall considerations

Your CDN must be able to access and cache static assets. If your instance is not publicly accessible will you need to make some changes to your firewall to allow requests from the CDN to pass through. We recommend using application firewalls instead of standard IP range filtering, as CDN IP ranges can change without notice.

Considerations for private instances

If your site is publicly accessible on the internet, you should be able to enable CDN without any problems.

If your site is not publicly accessible you can:

- configure your firewall to allow requests from your CDN to pass through. More information on how to do this is provided in our step-by-step guides below.
- set up your own caching servers closer to your users which will not require opening any traffic to the
 internet, instead of using a CDN vendor. SeeHow to configure Apache for caching and HTTP/2to learn
 more about this workaround.

Marketplace apps and third party customizations

Some marketplace apps or customizations may not be compatible with the CDN feature. A health check, on the Content Delivery Network admin screen will let you know if any of your apps are not compatible.

SeeUser-installed apps health check fails in Data Center when configuring CDNto find out what to do if any of your apps are incompatible.

If you've developed your own plugin, seePreparing for Confluence 7.0for information about the APIs you can use to confirm your plugin is compatible.

Configure your CDN for Confluence Data Center

On this page:

- Configure an internet facing load balancer (optional)
 - · Add an internet-facing load balancer
 - Update your firewall rules for the internet-facing load balancer
- Configure your CDN to cache assets
- Enable CDN inConfluence
 - Configure CDN in Confluence via REST API
- Troubleshooting
 - Frequently asked questions

If your users are distributed across the world and experience high latency when using Confluence Data Center, you may be able to improve their experience by using a Content Delivery Network (CDN). Common CDNs include AWS CloudFront, Cloudflare, Azure CDN, Akamai, and others.

Head to Use a CDN with Atlassian Data Center applications to learn about our CDN capabilities, and how to assess whether it will improve your users' experience.

Once you're ready to start using a CDN, there are three main steps:

- 1. Configure an internet-facing load balancer (optional)
- 2. Configure your CDN.
- 3. Enable the CDN feature in Confluence.

Configure an internet facing load balancer (optional)

If your site is not publicly accessible, you'll need to make sure that your CDN can reach it, but only to access and cache static assets. The way you do this depends on your particular load balancer and web application firewall. Refer to the documentation for your load balancer and firewall for detailed guidance.

Add an internet-facing load balancer

Add an internet-facing load balancer to your setup. This is in addition to your primary load balancer. Your CDN is the only entity that will interact with this load balancer. We recommend you:

- Enable HTTPS the traffic from this load balancer will be sent over the public internet and should be encrypted.
- Enable HTTP/1.1 currently, the caching proxies and CDNs do not handle HTTP/2 well (or at all) on the way to the origin.
- For AWS deployments, you would set up an internet-facing application load balancer.

Update your firewall rules for the internet-facing load balancer

Unlike your primary load balancer, this internet-facing load balancer must be locked down to ensure that your CDN can only pull data it is allowed to cache. When configuring your firewall rules we recommend:

- 1. The configuration should only allow requests for paths that start with "/s/". If your application is deployed with a context path (for example yoursite.com/wiki or yoursite.com/jira) you will need to include it in the path. All other requests must be blocked.
- 2. You can also choose to limit the allowedHTTP methods to GET, HEAD, OPTIONS.

For AWS deployments, you will configure a Web Access Control List (WebACL) in the Web Application Firewall attached to your application load balancer. The condition to use is a "string match condition" applied to "URI".

To check that your setup is secure, perform the following manual tests:

- 1. A GET onhttps://internet-facing-proxy/should return "403 FORBIDDEN".
- 2. A GET onhttps://internet-facing-proxy/s should return "403 FORBIDDEN".

- 3. A GET onhttps://internet-facing-proxy/s/should return "404 NOT FOUND".
- 4. A GET onhttps://internet-facing-proxy/s/.should return "403 FORBIDDEN".
- 5. A GET onhttps://internet-facing-proxy/s/../s/should return "404 NOT FOUND".

Configure your CDN to cache assets

You'll need an account with a CDN provider. You're responsible for all costs associated with your CDN.We only support serving static assets from a CDN at this time. This means page content, attached files, and personally identifiable information, including things like user avatars, won't be cached by your CDN.

We've prepared a CloudFormation template that you can use to configure Amazon CloudFront with minimal effort. You can find all our AWS deployment resources in this repositoryhttps://bitbucket.org/atlassian/atlassian-aws-deployment/src/master/templates/cdn/.

If you choose not to use our template, define the following in your CDN configuration. This example is based on AWS CloudFront.

Origin domain name	This is your Atlassian application base URL, including the context path if you've configured one. For example:mycompany.com/confluence
Origin path	Leave blank. There is no need to specify a path.
Allowed HTTP methods	Optionally limit to: GET, HEAD, OPTIONS
Viewer protocol policy	redirect HTTP to HTTPS
Object caching	Use origin cache headers
Forward cookies	None This is important to make sure static assets are cached without the user context.
Query String Forwarding and Caching	Forward all, cache based on all
HTTP protocols	Must include HTTP/2
Error pages/Error Caching Minimum TTL (seconds)	The default error page caching time for CloudFront is 5 minutes. Consider lowering it to a value in the range of 10-30 seconds to decrease the time required to recover from an outage.
Compress Objects Automatically	Yes

Using the default should be fine for most of the other settings.

You will need to adapt this information for your particularCDN provider. You should refer to the documentation for your CDN for details, as we've found that terminology differs between CDNs.

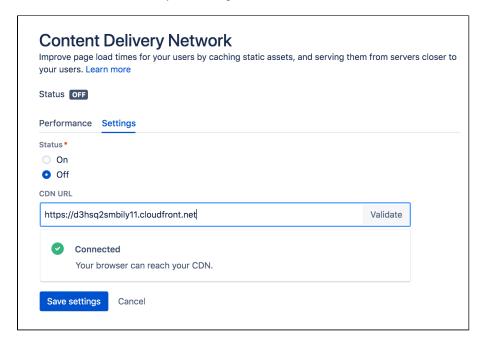
Enable CDN inConfluence

Once you've configured your CDN, you can enable the CDN option in Confluence.

To turn on CDN:

- 1. Go to Content Delivery Network.
- 2. Navigate to the Settingstab.
- 3. Set the status to On
- 4. Paste the URL generated by your CDNinto the URL field and hitValidate.

5. If successful, save your changes.



As end users access Confluence, static assets will be cached on the edge server closest to them, and served from there until they expire. This means it might take some time before you can start measuring the impact of the CDN, depending on when your users are online and accessing the site in each location.

Configure CDN in Confluence via REST API

You can also interact with the CDN feature using the following REST endpoint:

caching/configuration

- **GET** -returns the current CDN status, and URL.
- **DELETE** -deletes the existing configuration and reverts to the default state (CDN disabled, no URL). This is useful if you can't access the UI because of a caching problem.
- PUT -sets the CDN URL and status to the values passed in the body of the request as follows:

```
{
  enabled: true,
  url: https://yourcdnurl.com
}
```

Troubleshooting

Here are some common problems that you may encounter.

We only accept HTTPS CDN URLs

This is particularly important if you're using Azure CDN, as Azure CDN will mirror the same protocol as the originating request, which means your Data Center application will need to be provisioned with HTTPS.

Data Center application UI is inaccessible or not functional

Although unlikely, a misconfiguration of your CDNor a CDN service outage may mean your application's UI is not accessible. If this happens, you will need to disable the CDN feature using the REST API, as follows.

```
\verb| curl -v -u < admin username>: < admin password> -X | DELETE | | http://<your-base-url>/rest/static-asset-aching/configuration| | http://cyour-base-url>/rest/static-asset-aching/configuration| | http://cyour-base-url>/rest/static-asset-aching/configuration| | http://cyour-base-url>/rest/static-asset-aching/configuration| | http://cyour-base-url>/rest/static-asset-aching/configuration| | http://cyour-base-url>/rest/static-asset-aching/configuration| | http://cyour-base-url>/rest/static-asset-aching/configuration| | https://cyour-base-url>/rest/static-asset-aching/configuration| | https://cyour-base-aching/configuration| | https://cyour-base-url>/rest/static-asset-aching/configuration| | https://cyour-base-url>/rest/static-asset-aching/configuration| | https://cyour-base-url>/rest/static-asset-aching/configuration| | https://cyour-base-aching/configuration| | https://cyour-base-aching/configuration| | https://cyour-base-aching/configuration| | https://cyour-base-aching/configuration| | https://cyour-base-aching/configuration| | https://cyour-base-aching/configuration| | https://cyour-base-aching/configu
```

This example uses Curl, but you can use any language. Don't forget to replace the username, password, and base URL placeholders with your own details.

HTTP/2 disabled

Your load balancer, firewall, or reverse proxy shouldallow HTTP/2 traffic.Using HTTP/2 will provide the best performance for your end users. See HTTP/2 health check fails in Data Center when configuring CDNfor more information.

User-installed apps may not be compatible

This warning is displayed when we detect that a Marketplace or other user-installed app is using a deprecated method, which may result in assets being cached incorrectly. SeeUser-installed apps health check fails in Data Center when configuring CDNfor more information on what to do if you see this warning.

Frequently asked questions

Can I control what static assets are cached?

No, the application controls this. All requests for static assets are routed to the CDN. Requests for non-static assets are routed directly to your product.

Is personally identifiable information cached?

User created content, usernames, mentions, avatars etc are not static assets, so are not cached. Your CDN should also be configured to pull content from your product with cookies stripped to make sure it operates without user context.

Is dynamic content such as batch. jscached?

Although dynamically generated, batch. jsis considered static content, so is cached.

Improving instance stability with rate limiting

When automated integrations or scripts send requests to Confluence in huge bursts, it can affect Confluences stability, leading to drops in performance or even downtime. With rate limiting, you can control how many external REST API requests automations and users can make and how often they can make them, making sure that your Confluence instance remains stable.



Rate limiting is available for **Confluence Data Center**.

On this page:

- How rate limiting works
- How to turn on rate limiting
- Limiting requests what its all about
- Adding exemptions
- Identifying users who have been rate limited
- Viewing limited requests in the Confluence log file
- Getting rate limited users perspective
- Other tasks

How rate limiting works

Heres some details about how rate limiting works in Confluence.

Rate limiting targets only external REST API requests, which means that requests made within Confluence arent limited in any way. When users move around Confluence, creating pages, commenting, and completing other actions, they wont be affected by rate limiting, as were seeing this as a regular user experience that shouldnt be limited.

Lets use an example to better illustrate this:

- When a user visits a space in Confluence, a number of requests are sent in the background these
 requests ask Confluence for the pages, blog posts, etc. Since this traffic is internal to Confluence,
 it wont be limited.
- When the same user opens up the terminal on their laptop and sends a request (like the one below) to get the contents of a space, it will be rate limited because its made outside of Confluence.

 $\verb|curl -u user:password http://localhost:8090/rest/api/space/SPACEKEY/content| \\$

Authentication mechanisms

To give you more details on how we recognize which requests should be limited, were targeting external HTTP requests with these authentication mechanisms:

- Basic auth
- OAuth
- JSESSIONID cookie

Out of the many available techniques for enforcing rate limits, weve chosen to use token bucket, which gives users a balance of tokens that can be exchanged for requests. Heres a summary of how it works:

Users are given tokens that are exchanged for requests. One token equals one request.

Users get new tokens at a constant rate so they can keep making new requests. This is their Requests allowed, and can be, for example, 10 every 1 minute.

Tokens are added to a users personal bucket until its full. This is their Max requests and allows them to adjust the usage of tokens to their own frequency, for example 20 every 2 minutes instead of 10 every 1 minute, as specified in their usual rate.

When a user tries to send more requests than the number of tokens they have, only requests that can draw tokens from the bucket will be successful. The remaining ones will end in a 429 error message (too many requests). The user can retry those requests once they get new tokens.

Confluence tastes best when used with our other products like Jira. Technically, products like these are external to Confluence, so they should be limited. In this case, however, were treating them as belonging to the same user experience and dont want to enforce any limits for requests coming from or to these products.

The way it is now:

- Server: Not limited in any way.
- Cloud: Theres a known issue that applies rate limits to requests coming from/to cloud products.
 Were working hard to disable rate limits for cloud products and should make that happen soon.
 For now, if youre integrating Confluence with Jira cloud, you should make rate limits higher than usual.

The general assumption is that Marketplace apps are installed on a Confluence instance, make internal requests from within Confluence, and shouldnt be limited. But, as always, **it depends on how an app works**.

- Internal: If an app in fact works internally, enhancing the user experience, it wont be limited. An example of such app would be a special banner thats displayed in a Confluence space. Lets say this banner checks all pages that were created and shows this spaces winner a user whos created the most pages in the last month. Traffic like that would be internal, not limited.
- External: Apps whose requests are external to Confluence are limited. Lets say we have an app
 that displays a wallboard on TV. It asks Jira for details about boards, issues, assignees, etc. and
 then reshuffles and displays them in its own way as the earlier mentioned wallboard. An app like
 that sends external requests and behaves just like a user sending requests over a terminal.

It really depends on the app, but were assuming most of them shouldnt be limited.

Rate limiting is available for Data Center, so you most likely have a cluster of nodes behind a load balancer. You should know that each of your users will have a separate limit on each node (rate limits are applied per node, not per cluster).

In other words, if they have used their Requests allowed on one node and were rate limited, they could theoretically send requests again if they started a new session on a different node. Switching between the nodes isnt something users can do, but keep in mind that this can happen.

Whatever limit youve chosen (e.g. 100 requests every 1 hour), the same limit will apply to each node, you dont have to set it separately. This means that each users ability to send requests will still be limited, and Confluence will remain stable regardless of which node their requests are routed to.

Setting the right limit depends on many factors, so we cant give you a simple answer. We have some suggestions, though.

Finding the right limit

The first step is to understand the size of traffic that your instance receives. You can do this by parsing the access log and finding a user than made the most REST requests over a day. Since UI traffic is not rate limited, this number will be higher than what you need as your rate limit. Now, thats a base number you need to modify it further based on the following questions:

- Can you afford to interrupt your users work? If your users integrations are mission-critical, consider upgrading your hardware instead. The more critical the integrations, the higher the limit should be consider multiplying the number you found by two or three.
- 2. Is your instance already experiencing problems due to the amount of REST traffic? If yes, then choose a limit thats close to the base number you found on a day when the instance didnt struggle. And if youre not experiencing significant problems, consider adding an extra 50% to the base number this shouldnt interrupt your users and you still keep some capacity.

In general, the limit you choose should keep your instance safe, not control individual users. Rate limiting is more about protecting Confluence from integrations and scripts going haywire, rather than stoping users from getting their work done.

How to turn on rate limiting

You need the System Administrator global permission to turn on rate limiting.

To turn on rate limiting:

- 1. In Confluence, go to Configuration Rate limiting.
- 2. Change the status to Enabled.
- 3. Select one of the options: Allow unlimited requests, Block all requests, or Limit requests. The first and second are all about allowlisting and blocklisting. For the last option, youll need to enter actual limits. You can read more about them below.
- 4. Save your changes.

Make sure to add exemptions for users who really need those extra requests, especially if youve chosen allowlisting or blocklisting. See Adding exemptions.

Limiting requests what its all about

As much as allowlisting and blocklisting shouldnt require additional explanation, youll probably be using the **L imit requests** option quite often, either as a global setting or in exemptions.



Lets have a closer look at this option and how it works:

- 1. **Requests allowed**: Every user is allowed a certain amount of requests in a chosen time interval. It can be 10 requests every second, 100 requests every hour, or any other configuration you choose.
- 2. Max requests (advanced): Allowed requests, if not sent frequently, can be accumulated up to a set maximum per user. This option allows users to make requests at a different frequency than their usual rate (for example, 20 every 2 minutes instead of 10 every 1 minute, as specified in their rate), or accumulate more requests over time and send them in a single burst, if thats what they need. Too advanced? Just make it equal to Requests allowed, and forget about this field nothing more will be accumulated.

Examples

Requests allowed: 10/hour | Max requests: 100

One of the developers is sending requests on a regular basis, 10 per hour, throughout the day. If they try sending 20 requests in a single burst, only 10 of them will be successful. They could retry the remaining 10 in the next hour when theyre allowed new requests.

Another developer hasnt sent any requests for the past 10 hours, so their allowed requests kept accumulating until they reached 100, which is the max requests they can have. They can now send a burst of 100 requests and all of them will be successful. Once they used up all available requests, they have to wait for another hour, and theyll only get the allowed 10 requests.

If this same developer sent only 50 out of their 100 requests, they could send another 50 right away, or start accumulating again in the next hour.

Requests allowed: 1/second | Max requests: 60

A developer can choose to send 1 request every second or 60 requests every minute (at any frequency).

Since they can use the available 60 requests at any frequency, they can also send all of them at once or in very short intervals. In such a case, they would be exceeding their usual rate of 1 request per second.

Finding the right limit

Setting the right limit depends on many factors, so we cant give you a simple answer. We have some suggestions, though.

Finding the right limit

The first step is to understand the size of traffic that your instance receives. You can do this by parsing the access log and finding a user that made the most REST requests over a day. Since UI traffic is not rate limited, this number will be higher than what you need as your rate limit. Now, thats a base number you need to modify it further based on the following questions:

- 1. Can you afford to interrupt your users work? If your users integrations are mission-critical, consider upgrading your hardware instead. The more critical the integrations, the higher the limit should be consider multiplying the number you found by two or three.
- 2. Is your instance already experiencing problems due to the amount of REST traffic? If yes, then choose a limit thats close to the base number you found on a day when the instance didnt struggle. And if youre not experiencing significant problems, consider adding an extra 50% to the base number this shouldnt interrupt your users and you still keep some capacity.

In general, the limit you choose should aim at keeping your instance safe, not to control individual users. Rate limiting is more about protecting Jira from integrations and scripts going haywire, rather than stoping users from getting their work done.

Adding exemptions

Exemptions are, well, special limits for users who really need to make more requests than others. Any exemptions you choose will take precedence over global settings.

0

After adding or editing an exemption, youll see the changes right away, but it takes up to 1 minute to apply the new settings to a user.



To add an exemption:

- 1. Go to the **Exemptions** tab.
- 2. Click Add exemption.
- Find the user and choose their new settings.You cant choose groups, but you can select multiple users.
- 4. The options available here are just the same as in global settings: Allow unlimited requests, Block all requests, or Assign custom limit.
- 5. Save your changes.

If you want to edit an exemption later, just click **Edit** next to a users name in the **Exemptions** tab.

Recommended: Add an exemption for anonymous access

Confluence sees all anonymous traffic as made by one user: **Anonymous**. If your site is public, and your rate limits are not too high, a single person may drain the limit assigned to anonymous. Its a good idea to add an exemption for this account with a higher limit, and then observe whether you need to increase it further.

Identifying users who have been rate limited

When a user is rate limited, theyll know immediately as theyll receive an HTTP 429 error message (too many requests). You can identify users that have been rate limited by opening the **List of limited accounts** tab on the rate limiting settings page. The list shows all users from the whole cluster.



When a user is rate limited, it takes up to 5 minutes to show it in the table.



Unusual accounts

Youll recognize the users shown on the list by their name. It might happen, though, that the list will show some unusual accounts, so heres what they mean:

- Unknown: Thats a user that has been deleted in Confluence. They shouldnt appear on the list for
 more than 24 hours (as they cant be rate limited anymore), but you might see them in the list of
 exemptions. Just delete any settings for them, they dont need rate limiting anymore.
- Anonymous: This entry gathers all requests that werent made from an authenticated account. Since
 one user can easily use the limit for anonymous access, it might be a good idea to add an exemption
 for anonymous traffic and give it a higher limit.

Viewing limited requests in the Confluence log file

You can also view information about rate limited users and requests in the Confluence log file. This is useful if you want to get more details about the URLs that requests targeted or originated from.

When a request has been rate limited youll see a log entry similar to this one:

```
2019-12-24 10:18:23,265 WARN [http-nio-8090-exec-7] [ratelimiting.internal.filter.RateLimitFilter] lambda$userHasBeenRateLimited$0 User [2c9d88986ee7cdaa016ee7d40bd20002] has been rate limited -- url: /rest/api/space/DS/content | traceId: 30c0edcb94620c83 | userName: exampleuser
```

Getting rate limited users perspective

When users make authenticated requests, theyll see rate limiting headers in the response. These headers are added to every response, not just when youre rate limited.

Header	Description
X-RateLimit- Limit	The max number of requests (tokens) you can have. New tokens wont be added to your bucket after reaching this limit. Your admin configures this as Max requests.
X-RateLimit- Remaining	The remaining number of tokens. This value is as accurate as it can be at the time of making a request, but it might not always be correct.
X-RateLimit- Interval- Seconds	The time interval in seconds. You get a batch of new tokens every time interval.
X-RateLimit- FillRate	The number of tokens you get every time interval. Your admin configures this as Requests allowed.
retry-after	How long you need to wait until you get new tokens. If you still have tokens left, it shows 0; this means you can make more requests right away.

When youre rate limited and your request doesnt go through, youll see the HTTP 429 error message (too many requests). You can use these headers to adjust scripts and automations to your limits, making them send requests at a reasonable frequency.

Other tasks

Allowlisting URLs and resources

Weve also added a way to allow whole URLs and resources on your Confluence instance using a system property. This should be used as quick fix for something that gets rate limited, but shouldnt.

For example, a Marketplace app added some new API to Confluence. The app itself is used from the UI, so it shouldnt be limited, but it might happen that Confluence sees this traffic as external and applies the rate limit. In this case, you could disable the app or increase the rate limit, but this brings additional complications.

To work around issues like this, you can allowlist the whole resource added by the app so it works without any limits.

To allow specific URLs to be excluded from rate limiting:

- 1. Stop Confluence.
- 2. Add the com. atlassian.ratelimiting.whitelisted-url-patterns system property, and set the value to a comma-separated list of URLs, for example:

```
- D com. at lassian. rate limiting. white listed-url-patterns=/**/rest/applinks/**,/**/rest/capabilities,/**/rest/someapi
```

The way you add system properties depends on how you run Confluence. See Configuring System Properties for more information.

3. Restart Confluence.

For more info on how to create URL patterns, see AntPathMatcher: URL patterns.

Allowlisting external applications

You can also allowlist consumer keys, which lets you remove rate limits for external applications integrated through AppLinks.



If you're integrating Confluence with other Atlassian products, you don't have to allowlist them as this traffic isn't limited.

- 1. Find the consumer key of your application.
 - a. Go to O-> General Configuration Application Links.
 - b. Find your application, and click Edit.
 - c. Copy the Consumer Key from Incoming Authentication.
- 2. Allowlist the consumer key.
 - a. Stop Confluence.
 - b. Add the com. at lassian.ratelimiting.whitelisted-oauth-consumers system property, and set the value to a comma-separated list of consumer keys, for example:

 $-\texttt{Dcom.atlassian.ratelimiting.whitelisted-oauth-consumers=app-connector-for-confluence-confluenc$ server

The way you add system properties depends on how you run Confluence. SeeConfiguring System Properties for more information.

c. Restart Confluence.

After entering the consumer key, the traffic coming from the related application will no longer be limited.

Adjusting your code for rate limiting

Weve created a set of strategies you can apply in your code (scripts, integrations, apps) so it works with rate limits, whatever they are.

For more info, see Adjusting your code for rate limiting.

Adjusting your code for rate limiting

Whether its a script, integration, or app youre using if its making external REST API requests, it will be affected by rate limiting. Until now, you could send an unlimited number of REST API requests to retrieve data from Confluence, so were guessing you havent put any restrictions on your code. When admins enable rate limiting in Confluence, theres a chance your requests will get limited eventually, so we want to help you prepare for that.

Before you begin

To better understand the strategies weve described here, its good to have some some basic knowledge about rate limiting in Confluence. When in doubt, head to Improving instance stability with rate limiting and have a look at the first paragraph.

Quick reference

Success: When your request is successful, youll get a 2xx code.

Error: When your request fails, youll get a 4xx code. If youre rate limited, it will be 429 (too many requests).

The following HTTP headers are added to every authenticated request affected by rate limiting:

Header	Description
X-RateLimit- Limit	The max number of requests (tokens) you can have. New tokens wont be added to your bucket after reaching this limit. Your admin configures this as Max requests.
X-RateLimit- Remaining	The remaining number of tokens. This value is as accurate as it can be at the time of making a request, but it might not always be correct.
X-RateLimit- Interval- Seconds	The time interval in seconds. You get a batch of new tokens every time interval.
X-RateLimit- FillRate	The number of tokens you get every time interval. Your admin configures this as Requests allowed.
retry-after	How long you need to wait until you get new tokens. If you still have tokens left, it shows 0; this means you can make more requests right away.

Strategies

Weve created a set of strategies you can apply in your code so it works with rate limits. From very specific to more universal, these reference strategies will give you a base, which you can further refine to make an implementation that works best for you.

1. Exponential backoff

This strategy is the most universal and the least complex to implement. Its not expecting HTTP headers or any information specific to a rate limiting system, so the same code will work for the whole Atlassian suite, and most likely non-Atlassian products, too. The essence of using it is observing whether youre already limited (wait and retry, until requests go through again) or not (just keep sending requests until youre limited).

◆Universal, works with any rate limiting system.

Doesnt require too much knowledge about limits or a rate limiting system.

⊗High impact on a Confluence instance because of concurrency. Were assuming most active users will send requests whenever theyre available. This window will be similar for all users, making spikes in Confluence performance. The same applies to threads most will either be busy at the same time or idle.

☼Unpredictable. If you need to make a few critical requests, you cant be sure all of them will be successful.

Summary of this strategy

Heres the high-level overview of how to adjust your code:

- 1. **Active**: Make requests until you encounter a 429. Keep concurrency to a minimum to know exactly when you reached your rate limit.
- 2. **Timeout**: After you receive a 429, start the timeout. Set it to 1 second for starters. Its a good idea to wait longer than your chosen timeout up to 50%.
- 3. Retry: After the timeout has passed, make requests again:
 - a. Success: If you get a 2xx message, go back to step 1 and make more requests.
 - b. **Limited**: If you get a 429 message, go back to step 2 and double the initial timeout. You can stop once you reach a certain threshold, like 20 minutes, if thats enough to make your requests work.

With this strategy, youll deplete tokens as quickly as possible, and then make subsequent requests to actively monitor the rate limiting status on the server side. It guarantees youll get a 429 if your rate is above the limits.

2. Specific timed backoff

This strategy is a bit more specific, as its using the retry-after header. Were considering this header an industry standard and plan to use it across the Atlassian suite, so you can still be sure the same code will work for Bitbucket and Confluence, Server and Cloud, etc. This strategy makes sure that you will not be limited, because youll know exactly how long you need to wait before youre allowed to make new requests.

✓Universal, works with any rate limiting system within the Atlassian suite (and other products using retry-after) Bitbucket and Confluence, Server and Cloud, etc.

♥Doesnt require too much knowledge about limits or a rate limiting system.

⊗High impact on a Confluence instance because of concurrency. Were assuming most active users will send requests whenever theyre available. This window will be similar for all users, making spikes in Jira performance. The same applies to threads most will either be busy at the same time or idle.

Summary of this strategy

Heres the high-level overview of how to adjust your code:

- Active: Make requests and observe the retry-after response header, which shows the number of seconds you need to wait to get new tokens. Keep concurrency level to a minimum to know exactly when the rate limit kicks in.
 - a. Success: If the header says 0, you can make more requests right away.
 - b. **Limited**:If the header has a number greater than 0, for example 5, you need to wait that number of seconds.
- 2. **Timeout**: If the header is anything above 0, start the timeout with the number of seconds specified in the header. Consider increasing the timeout by a random fraction, up to 20%.
- 3. Retry: After the timeout specified in the header has passed, go back to step 1 and make more requests.

With this strategy, youll deplete tokens as quickly as possible, and then pause until you get new tokens. You should never hit a 429 if your code is the only agent depleting tokens and sends requests synchronously.

3. Rate adjustment

This strategy is very specific and expects particular response headers, so its most likely to work for Confluence Data Center only. When making requests, youll observe headers returned by the server (number of tokens, fill rate, time interval) and adjust your code specifically to the number of tokens you have and can use.

It can have the least performance impact on a Confluence instance, if used optimally.

- ✔Highly recommended, especially for integrations that require high-volume traffic.
- Safe, as you can easily predict that all requests that must go through will in fact go through. It also allows for a great deal of customization.
- ⊗Very specific, depends on specific headers and rate limiting system.

Summary of this strategy

Heres the high-level overview of how to adjust your code:

- 1. Active: Make requests and observe all response headers.
- 2. **Adjust:** With every request, recalculate the rate based on the following headers:
 - a. x-ratelimit-interval-seconds: The time interval in seconds. You get a batch of new tokens every time interval.
 - b. x-ratelimit-fillrate: The number of tokens you get every time interval.
 - c. retry-after: The number of seconds you need to wait for new tokens. Make sure that your rate assumes waiting longer than this value.
- 3. **Retry:** If you encounter a 429, which shouldnt happen if you used the headers correctly, you need to further adjust your code so it doesnt happen again. You can use the retry-after header to make sure that you only make requests when the tokens are available.

Customizing your code

Depending on your needs, this strategy helps you to:

By following the headers, you should know how many tokens you have, when you will get the new ones, and in what number. The most useful headers here are x-ratelimit-interval-seconds and x-ratelimit-fillrate, which show the number of tokens available every time interval. They help you choose the perfect frequency of making your requests.

You can wait to perform complex operations until youre sure you have enough tokens to make all the consecutive requests you need to make. This allows you to reduce the risk of leaving the system in an inconsistent state, for example when your task requires 4 requests, but it turns out you can only make 2. The most useful headers are x-ratelimit-remaining and x-ratelimit-interval-seconds, which show how many tokens you have right now and how long you need to wait for the new ones.

With all the information returned by the headers, you can create more strategies that work best for you, or mix the ones weve described here. For example:

If youre making requests once a day, you can focus on the max requests you can accumulate (x-ratelimit-limit), or lean towards the remaining number of tokens if a particular action in Confluence triggers your app to make requests (x-ratelimit-remaining).

If your script needs to work both for Confluence Data Center and some other application, use all headers for Confluence and focus on the universal retry-after or request codes if the app detects different software.

Running Confluence Data Center on a single node

Data Center allows you to run Confluence in a cluster with multiple nodes, or on a single server (also known as non-clustered, or standalone Data Center).

This page outlines the architecture and requirements of a non-clustered Confluence Data Center deployment, as well as some of the benefits and considerations.

Architecture

The deployment architecture of a non-clustered Data Center deployment is the same as a Server installation. Heres what a typical setup looks like:



As you can see, Confluence Data Center deployed on a single node looks just as a Server installation, and consists of:

- Confluence Data Center, running on a single node
- A database that Confluence reads and writes to

SeeGetting started as a Confluence administrator to learn more about single server Confluence installations.

Requirements

Non-clustered Confluence Data Center installations have the same minimum requirements as a Confluence Server installation. Check our Confluence System Requirements guide for a full overview of the supported platforms and hardware youll need.

Benefits of running a non-clustered Data Center deployment

There are a range of reasons you may choose a single node Data Center. Some of the benefits include:

Keeping your existing infrastructure

Running on a single node means that you can upgrade from Server to Data Center without adding to your infrastructure. In most cases, moving to Data Center will be as simple as updating your license.

Accessing Data Center-only features

Your Data Center license unlocks a suite of additional security, compliance, and administration features to help you easily manage enterprise-grade Confluence site like SAML single sign-on, advanced permission management, rate limiting, and more. See the complete list.

As non-clustered Confluence Data Center installations are cluster-compatible, you can still enable and configure clustering whenever youre ready to scale. Learn more about setting up a cluster.

Considerations

Non-clustered Data Center is the simplest setup, but it has some limitations. Just like a Server installation, youll still have the application server as a single point of failure, so it cant support high availability or disaster recovery strategies.

Some deployments start to experience performance or stability issues once their size profile hits Large or XLarge. Most clustered deployments provide you the flexibility to scale up your infrastructure to address heavy loads (or even scale down to save costs during light loads). On AWS or Azure, you can also quickly address most stability issues by replacing misbehaving nodes with fresh ones.



For more information about size profiles, see Data Center performance sizing. We also explain our own strategies for managing our clustered deployments in How Atlassians monitor their enterprise deployments.