Atlassian

Documentation for Confluence 5.9

Contents

Confluence Documentation Home

Confluence is where you create, organize and discuss work with your team. Learn more or try it for free.

Capture the knowledge that's too often lost in email inboxes and shared network drives in Confluence instead – where it's easy to find, use, and update. Give every team, project, or department its own space to create the things they need, whether it's meeting notes, product requirements, file lists, or project plans, you can get it done in Confluence.

Confluence Server 5.9 (Latest)

Get started

Get started

Pages and Blogs

Search

Popular

Links

Tables

Macros
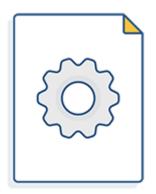
Integrations

Integrate with JIRA applications

Integrate with HipChat

Get add-ons from the Marketplace

Installation

Install Confluence

Supported Platforms

Database Configuration

Release Notes

Tutorials

Navigate Confluence

Space Ace

Admin

Manage Users

Customise your Site

Get started

Welcome to the Confluence getting started documentation. In this section, you'll find tutorials and other information that'll be useful for evaluating Confluence, and getting to know it when you're starting out.

Teams in Space

For each tutorial in this section, we'll use a fictional organisation known as 'Teams in Space'. Their mission is to:

> "Perform flight research and technology integration to revolutionize aviation and pioneer aerospace technology. Also, land the first humans on Mars by 2020."

You're an astronaut in the 'See Space EZ' team, which is working on the upcoming colonisation of Mars.

Go ahead – dive into the tutorials, and let us show you around Confluence and some of its handy features.

- Tutorial: Navigate Confluence
  - The dashboard
  - The space directory
  - The space sidebar
  - Keyboard shortcuts
  - Complete your mission
- Tutorial: Space ace
  - Create a project space
  - Create your personal space
  - Create the team's PR space
  - Delete and archive spaces
- Confluence Mobile

Tutorial: Navigate Confluence

Confluence is pretty simple to use, once you get to know it. This tutorial aims to get you acquainted with the Confluence user interface, and show you how and where to perform some common tasks.

Teams in Space

In this tutorial,you'll be working with some new Teams in Space recruits. Let's get to know them.

Alana Baczewski
Tech Lead

Emma Silvetti
Launch Specialist

William Vladinov
Aerospace Engineer

Now that you've met your team, let's take a look at your mission.

Mission brief

Your mission commander has thrown you a curveball: this week you'll be training new recruits at Teams in Space HQ on your collaboration tool – Confluence. You just need to know the basics, so we'll go through the main things you need to know to complete your mission.

Your mission is broken up into the following components:

- Get to know the dashboard
- Find your way in the space directory
- Master the space sidebar
- Impress everyone with keyboard shortcuts

Those new recruits will be here tomorrow; we better get started!

Let's go!

The dashboard

1. The dashboard
2. The space directory
3. The space sidebar
4. Keyboard shortcuts
5. Complete your
   mission

The dashboard is the hub of your Confluence site, providing you with access to information and updates that are important to you. It's also the first thing your new recruits will see, so you need to make a good impression on this one.

You can get to the dashboard from anywhere in Confluence by choosing the **site logo** at the left of the Confluence header.

The dashboard has a collapsible sidebar that helps you get around:

- **Discover**
  Watch the action unfold in real time with **All updates or** check out pages with lots of likes and activity in the **Popular** feed.
- **My Work**
  Get lightning fast access to your recently created and edited pages in **Recently worked on**, get back to that page you stumbled across yesterday in **Recently viewed**, and have mission critical pages on speed dial under **Favourites**.
- **My Spaces**
  This is where you can keep links to the spaces that you hop in and out of several times a day.
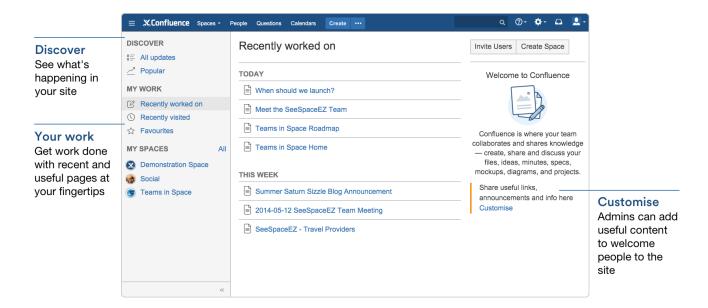
If you're a Confluence admin you can give the dashboard some personality by adding useful announcements, links, or a photo from your last mission (or office party). The the whole right hand column is ready and waiting for you to customise.

Discover
See what's happening in your site

Your work
Get work done with recent and useful pages at your fingertips

Customise
Admins can add useful content to welcome people to the site

Try clicking one of the spaces on the sidebar, then return to the dashboard by clicking the **site logo**. Even when your shuttle is spinning out of control, the dashboard is there to orient you.

You'll discover more about the dashboard as you get to know Confluence, but, for now, let's move on to the space directory.

> **Pro tips**
> - You can choose to set any page as your personal home page
> - You can always get to the dashboard at `https://yoursite.com/wiki/dashboard.action`
> - Your Confluence admin can customise the global dashboard that all users see

Next
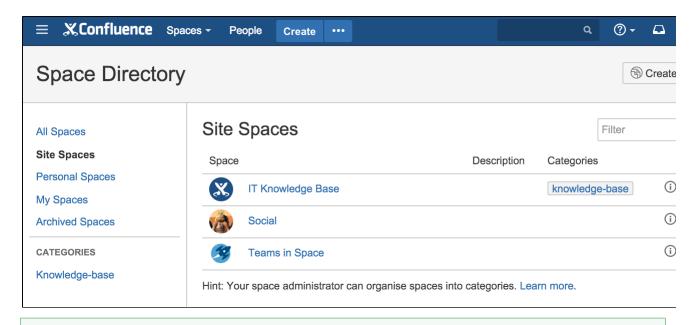
The space directory
1. The dashboard
2. The space directory
3. The space sidebar
4. Keyboard shortcuts
5. Complete your mission

The space directory won't let you look up ET in the intergalactic phone book, but it *will* let you see and filter all the spaces in your Confluence site. Spaces are places to collect pages with a common theme – you can create as many spaces as you like – and you can find them all in the space directory.

Here are some tasks to get you comfortable using the space directory:

1. **Visit the space directory** – To get to the space directory, choose **Spaces** > **Space directory** in the Confluence header.
2. **Choose the spaces you'll use the most** – No doubt there'll be a space or two that you'll use on a regular basis. Click the star to the right of a space to make it appear under **My Spaces** on the dashboard.
3. **Choose space categories** – Once you're there, you'll see a list of all the spaces in your Confluence site. Choose the 'My Spaces' category on the left to see only the spaces you marked with a star. Then choose all spaces again.
4. **Filter the list of spaces** – Type part of a space name in the **Filter** field at the top right. That'll quickly narrow down the list of spaces if there are a lot of them.

> You can also categorise spaces with labels you create yourself. We're not going to cover that here, but, if you'd like to know more, you can check out Use Labels to Categorise Spaces.

Understanding and using the space directory will make it much easier to find pages and blog posts that are relevant to you.

You're ready to impress those new recruits with your knowledge of Confluence's space directory; now it's time to sneak a peek at the space sidebar.

Next

The space sidebar

1. The dashboard
2. The space directory
3. The space sidebar
4. Keyboard shortcuts
5. Complete your mission

What's in the sidebar?

The sidebar is a feature of every Confluence space; it's where you'll find the page tree (a hierarchical list of pages in the space), customisable space shortcuts, and a link to the space's blog.
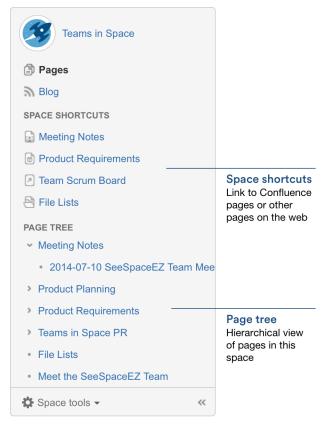
The See Space EZ team will find their meeting notes, decisions, requirements, and other pages in the sidebar. Basically any page you create in the space will appear in the sidebar by default.

> When you use certain page templates, like meeting notes, Confluence will automatically add an index page to your space shortcuts. The index page is just a place where you can view all pages of the same type – meeting notes in this example – in one place.

The space's blog is great for announcements and what's new-type updates.

Configure the sidebar

You can expand or collapse the sidebar using the left square bracket (**[**) on your keyboard, or by dragging it with your mouse.

> Teams in Space
>
> 📄 **Pages**
> 🔊 Blog
>
> **SPACE SHORTCUTS**
> 📋 Meeting Notes
> 📋 Product Requirements
> ↗ Team Scrum Board
> 🖨 File Lists
>
> **PAGE TREE**
> ⌄ Meeting Notes
> • 2014-07-10 SeeSpaceEZ Team Mee
> › Product Planning
> › Product Requirements
> › Teams in Space PR
> • File Lists
> • Meet the SeeSpaceEZ Team
>
> ⚙ Space tools ▾   «

Space shortcuts
Link to Confluence pages or other pages on the web

Page tree
Hierarchical view of pages in this space

There are lots of things you can change in the sidebar, but we'll stick to the basics – just enough to train your new recruits.

> You need to be a space admin to complete this task. Take a look, but feel free to skip it if you're not an admin of any space.

Try this out:

1. Choose **Space tools** > **Configure sidebar** from the bottom of the sidebar (or the cog menu if your sidebar is collapsed)
2. Add a space shortcut by clicking **+Add link**
   Shortcuts can be to Confluence pages or spaces, or to any other content on the web. Try linking to this blog post, which mentions Teams in Space (we're always after a plug at Teams in Space HQ!)

You can also hide things – like the space's blog – in the sidebar, if they're of no use in the space.

The sidebar is pretty easy, right? You'll be schooling those recruits in no time. Next up: Impress them with your knowledge of keyboard shortcuts.

Next

Keyboard shortcuts

1. The dashboard
2. The space directory
3. The space sidebar
4. Keyboard shortcuts
5. Complete your mission

Give a person some space food, and they'll eat for a day; teach a person to rehydrate their own space food, and they're set for the whole mission. Or something like that. Keyboard shortcuts fall into this basket. We could give you a list here, but then you'd need to keep referring to this page. The best way to go is to show you how to find the list of keyboard shortcuts *within Confluence*.

Your recruits also need to work fast, so you'll need to pass this wisdom on to them.

Find the keyboard shortcuts

To open the list of keyboard shortcuts in Confluence, do any of the following:

- Choose the help icon
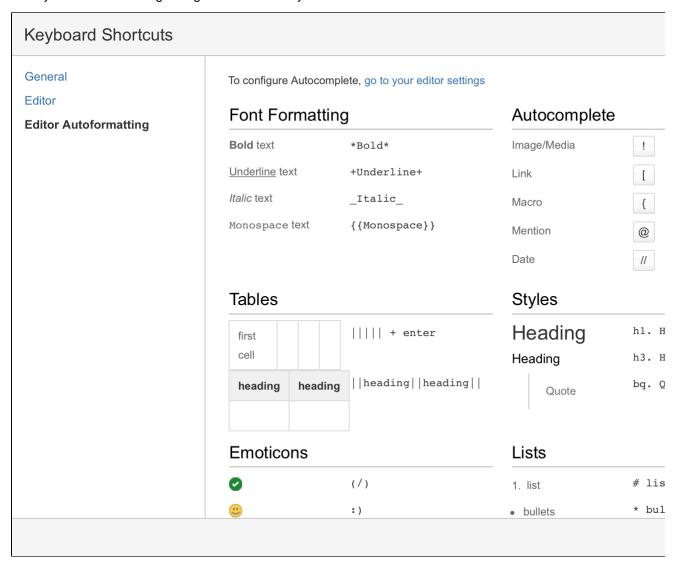
  at top right of the screen, then choose **Keyboard Shortcuts**
- When viewing a page, press **Shift+?**
- While editing a page, choose the question mark icon

  in the editor toolbar

What you'll see is a dialog listing the available keyboard shortcuts in Confluence.

The keyboard shortcuts are broken up into 3 categories:

- **General** – Global, page and blog post shortcuts.
- **Editor** – Text editing and formatting shortcuts.
- **Editor Autoformatting** – Wiki markup and autoformatting shortcuts.

> You can turn the 'General' keyboard shortcuts off when you visit the **General** tab in the keyboard shortcuts dialog.

Take some time to open the dialog and take a look at the shortcuts, and maybe find some you'll use a lot.

Then, start practising!

Next, we'll wrap up this mission and give you some ideas about where to head from here.

Next

Complete your mission

1. The dashboard
2. The space directory
3. The space sidebar
4. Keyboard shortcuts
5. Complete your mission

Well done, astronaut, you've acquitted yourself admirably. I'm sure those new recruits will be mightily impressed with your knowledge of Confluence.

In this tutorial, we've:

- Explored the anatomy of the dashboard
- Navigated using the space directory and favourited a space
- Taken a look at and customised the space sidebar
- Found a handy list of keyboard shortcuts to help you work faster

Just look at your team's adoring faces...

If you'd like to take things to the next level, check out our tutorial on becoming a space ace.

Tutorial: Space ace

This tutorial will take you on a journey through Confluence to create and customise spaces, and delete them if you want to, so you can achieve the rank of 'Space Ace'!

You'll need to have the 'Create space' and 'Create personal space' permission to complete this tutorial. If you've just set up Confluence, you won't have any trouble; if you're using an existing instance and you're not an admin, speak to your Confluence admin to make sure you have the right permissions.

Teams in Space

In this tutorial you're a new recruit on the Teams in Space crew, but, even though you're new, you'll be given a lot of responsibility. You need the power to go with it.

TEAMS IN SPACE

Mission brief

You're in charge of organising information and resources for the planned mission to Mars. There's going to be plenty of important information, and it must be readily available to the people who need it. Some information, though, will be sensitive, and may be 'for your eyes only.' You'll use the power and flexibility of Confluence spaces to organise information, and make sure it's visible to the right people.

Your mission is broken up into the following components:

- Create a space to house all of the important information related to the mission
- Create your own space to keep yourself organised
- Create a public relations space, where you'll introduce your team the world

What's a space?

Well, being an astronaut, I hope you know what *space* is, but what's a Confluence space all about? It's really just a place to put related things, like information pages and files. But spaces also give you a place to collaborate with groups of people, whether that's your team, people working on a common project, or the whole world.

Every space has its own permissions, allowing you to grant access and other privileges to the right people. They also have a blog, so you can post important messages and updates to whoever can see the space. You can have as many Confluence spaces as you like, and you can archive or delete spaces when you no longer need them.

Enough about that; let's begin.

[Start the mission!](#)

Create a project space

1. Create a project space
2. Create your personal space
3. Create the team's PR space
4. Delete and archive spaces

The Mars colonisation crew needs a place to put all their mission-critical information and resources, and you're charged with setting it up. This part is going to be easy, because the information needs to be viewable by the entire Teams in Space organisation. That means we can set up the project space without any special permissions.

If you haven't done so already, open up Confluence and log in so we can get started.

Create the space

1. Choose **Spaces** > **Create space** from the Confluence header
2. Select the **Blank space** option and choose **Next**
3. Enter a **Space name** – for this space, we'll call it 'Mars Colony', as it's being used for the Mars colonisation project.
4. Change the **Space key** to 'MARS' – this step isn't absolutely necessary, but it helps people if they're navigating to this space by name. The space key forms part of the URL, so making it a word or name makes it much easier to associate with your project.
5. Click **Create**
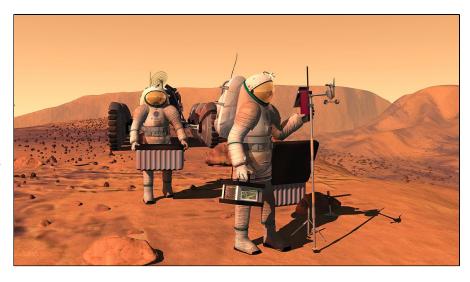
You now have a space set up for the Mars colonisation project. Because everyone at Teams in Space HQ needs access to the information in this space, you don't need to do anything with the space's default permissions. It's visible to everyone in your organisation, but not to the general public.

Every space has a default home page, which you can customise to suit your needs. Add the following image and text to your space's home page to get things started. Just click **Edit** (or press **E** on your keyboard) to edit the home page, and copy and paste the text. For the image it's best to drag it to your desktop and save it there, then drag it into your page. That'll make sure the image is attached directly to the page.

Hit **Save** when you're happy with the home page.

A **human mission to Mars** ha
s been the subject of [science fiction](#), [engineering,](#) and

scientific proposals throughout the 20th century and into the 21st century. The plans comprise proposals to land on Mars, eventually settling on an d terraforming the planet, while exploiting its moons, Phobos a nd Deimos.

Exploration of Mars has been a goal of national space programs for decades. Preliminary work for missions that would involve human explorers has been undertaken since the 1950s, with planned missions typically being cited as taking place 10 to 30 years in the future when they are drafted. The list of manned Mars mission plans in the 20th century shows the various mission proposals that have been put forth by multiple organizations and space agencies in this field of space exploration.

Your 'Mars colony' space is ready for your team to add pages to. If you want to find it again, choose **Spaces** i n the Confluence header, and select it from the list; to add a page to it, navigate to the space and hit **Create** i n the header.

Now it's on to your personal space; a place where you can work in peace, and be sure that no one's looking over your shoulder.

Next

Create your personal space
1. Create a project space
2. Create your personal space
3. Create the team's PR space
4. Delete and archive spaces
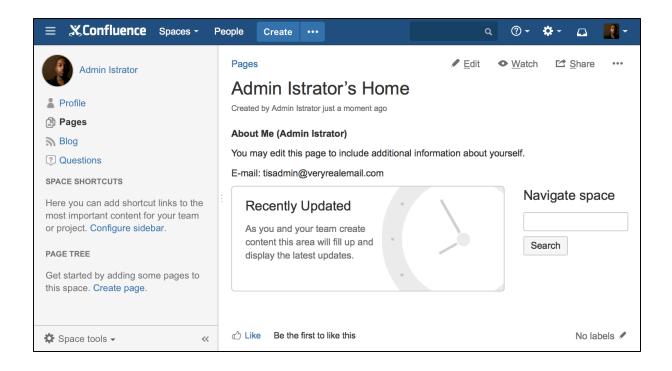
As a newbie on the team, you might want to keep some work to yourself until you're ready to present it. There's always the chance your mission commander will also send you some information that's 'for your eyes only,' and you'll need to keep that in a safe place.

For this part of the mission, we'll be creating a special type of space: a personal space. We'll be using your personal space like a sandbox, at least to start with – somewhere you can play around, draft pages, try out features, and generally see what spaces are capable of.

Create your personal space

1. Choose your profile picture at the right side of the Confluence header
2. Select **Add Personal Space...**
3. Hit **Create**

   You've now got a space that you can call your own. But we still need to lock it down to make sure it's only visible to you.

4. Choose **Space tools** > **Permissions** from the bottom of the sidebar
5. Hit **Edit Permissions** (enter your password if prompted)

   You should see the `confluence-users` group listed under **Groups**. To the left of the list of permissions is the **View** permission, which determines whether everyone in that group can see your space.

6. Uncheck **View** and hit **Save all** at the bottom of the page

You're now the only one that can view this space. Feel free to try anything in this space, and store super secret stuff here.

Next

Create the team's PR space

1. Create a project space
2. Create your personal space
3. Create the team's PR space
4. Delete and archive spaces

Now it's time to go public; the world needs to know about the mission and its brave participants.

In this step, we'll create a team space and open it up to everyone. That's right – you can open Confluence spaces up to anonymous (not logged in) users.
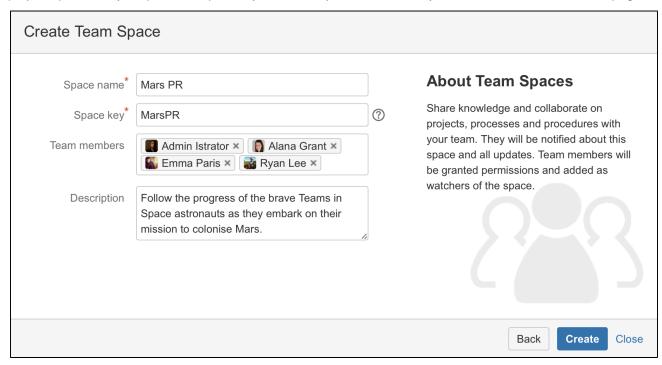
> In order to allow anonymous access to your Confluence site, a site admin needs to grant anonymous users the 'Use Confluence' permission. Don't worry if you can't do that, or if it's not done; it's just something to note if you're opening up your Confluence site for real.

Create a Team space

1. Choose **Spaces** > **Create space** from the header
2. Select **Team Space** and hit **Next**
3. Enter a **Space name** (let's call it 'Mars PR')

4. Change the **Space key** to 'MarsPR'
5. If there are other people using Confluence with you, feel free to add them as **Team members** (you can remove them later), or just stick with yourself for now
6. Paste this in as the **Description:** Follow the progress of the brave Teams in Space astronauts as they embark on their mission to colonise Mars.

Great! You now have a team space, again with its own home page. This home page is a little different to the project space and your personal space – you'll see any team members you added, listed on the home page.

Each space you create also has its own blog, so your social media team will be able to create posts in this space and speak directly to all those adoring fans. But none of those fans can see this space. Yet.

Allow anonymous access

It's time to let the world in by changing the permissions on this space.

1. Choose **Space tools** > **Permissions** from the bottom of the sidebar
2. Scroll down until you see **Anonymous**, then hit **Edit Permissions**
3. Tick the **View** permission for anonymous users and hit **Save all**

That's it. You can now share the space's URL, which will be http(s)://<your_confluence_site>/display/MarsPR. Visitors to that space don't need to log in, or have a licence for Confluence.

Next

Delete and archive spaces

1. Create a project space
2. Create your personal space
3. Create the team's PR space
4. Delete and archive spaces

We hope you've had a successful mission, and learned a bit about the power and versatility of Confluence spaces. Flash forward 18 months, and just look how well the colony is coming along!

If you need to clean up old spaces (or destroy the evidence of a failed mission!), you can either archive or delete a space. Archiving just means it won't show up in the regular search, whereas deleting is obviously a lot more permanent.

To archive a space:

1. Choose **Space tools** > **Overview** from the bottom of the sidebar
2. Click **Edit Space Details**
3. Change the **Status** from 'Current' to 'Archived' and hit **Save**

To delete a space:

1. Choose **Space tools** > **Overview** from the bottom of the sidebar
2. Select the **Delete Space** tab

What next?

If you'd like to know more about spaces and the permissions that govern them, check out Spaces and Permissions and Restrictions in the Confluence documentation.

Teams in Space HQ, signing off.

Confluence Mobile

When you access Confluence on a mobile device, you'll see a version of Confluence which is optimised for mobile viewing. Confluence chooses the mobile or desktop interface based on your device, but you can still switch to the desktop site on your mobile by choosing menu

then choosing **Switch to desktop version**.

You can also swap from the desktop view to the mobile view if you're on a mobile device, by choosing **Switch to Confluence Mobile** at the top of the page.

On your supported mobile device, you can:

- View the Confluence dashboard, pages, blog posts, and user profiles.
- Add or reply to a comment on a page or blog post.
- Like a page, blog post or comment.
- Watch a page or blog post.
- See your notifications and tasks.

You can't add or edit pages or blog posts, or edit existing comments, using the mobile interface.

The dashboard – the first thing you see

Choose a tab to see:

- Popular content – what people like in your wiki.
- Recent blogs – the latest blog posts.
- Network – updates by people in your network.

Tap the links to view the full content of a page, blog post or comment.

Searching for content and people

Tap the menu icon

to open the menu panel on the left of the page. Then type text or a person's name in the **Search** box. The mobile interface offers the quick navigation search, which returns matches on page title only. (See Search.) To use the full search, switch to desktop mode.

Viewing pages, blog posts and comments

Tap a link on the dashboard or on any other page. Confluence will display the linked page, blog post or comment.

You can:

- View the content, tap a link to move to another page, and interact with the page using the standard functionality supported by mobile browsers.
- Like or unlike a page, blog post or comment.
- Watch or stop watching a page or blog post.
- Add or reply to a comment.

Viewing people's profiles
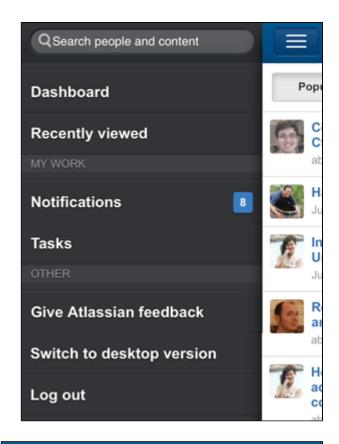
Search for a person's name, then view that person's user profile. Tap the options to phone, SMS or email a colleague directly from your mobile device.
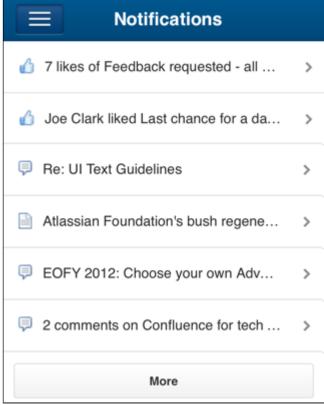
Following up on notifications

You can view and respond to your notifications on your phone or other mobile device. Tap the menu icon

to open the menu panel on the left of the page. Choose **Notifications**, and tap a notification to see its details. You can reply, watch or like via the inline actions. Tap **Open** to open the page or blog post in a new page. For full details, see Workbox Notifications.

Viewing tasks

You can view and manage your tasks on your phone or other mobile device too. Tap the menu icon

to open the menu panel on the left of the page. Choose **Tasks** then tap a task to see its details.

More things you may need to know
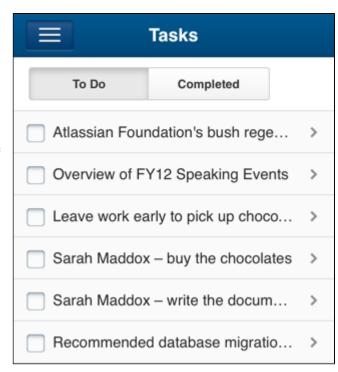
Some macros may not appear

Confluence macros are not yet fully supported in
Confluence mobile. If you view a page that contains
an unsupported macro, you will see a message
inviting you to click through to the desktop version of
the page.

*Screenshot: Macro not rendered in Confluence
mobile*

**Administrators can disable Confluence mobile on
your site**

The mobile functionality is provided by a plugin called the 'Confluence Mobile Plugin'. To remove the
functionality from your site, you can disable the plugin.

Confluence mobile is a web interface, not a native app

Atlassian does not supply a native mobile application that you can download and install onto your mobile
device. Confluence mobile is a web-based user interface, which Confluence displays when it detects a
mobile client.See Supported Platforms for supported mobile browsers.

Spaces

Confluence Spaces are containers for pages and blo
g posts with related content, and they come in two
main varieties:

- **Site spaces** – Sometimes called 'global'
  spaces, these are areas where you can
  create content and collaborate with other
  users.
- **Personal spaces** – You, and other
  Confluence users, can set up a personal
  space. You can keep it private, or open it up
  for other users to view or edit. Personal
  spaces are listed in the People Directory.

Create a space any time you need a place to store
related content by choosing **Spaces** > **Create space**
in the Confluence header. Two of the main ways
people use Confluence spaces are for teams and/or
projects.

Create a space for the Marketing team, Dev team, IT
team, HR team, and any team that needs to
collaborate on work, and store information and files.
If you have projects, big or small, that would benefit
from having a place people can work together and

On this page:

- What's a space admin?
- How is content arranged in
  spaces?
- View all spaces in Confluence

Related pages:

- Create a Space
- Create a Personal Space
- Space Permissions Overview
- Configure the Sidebar
- Export Content to Word, PDF,
  HTML and XML
- Delete a Space
- Administer a Space

store related files, make a space for them too. There's no limit to the number of spaces you can create in Confluence.

What's a space admin?

Every space has at least one space admin – usually the person that created the space – and that person can grant permissions to that space, including granting space admin privileges to others. Space admins don't have to be Confluence admins; they can be a regular Confluence users, and may only have special permissions for a single space. For example, you'll be the space admin for your personal space.
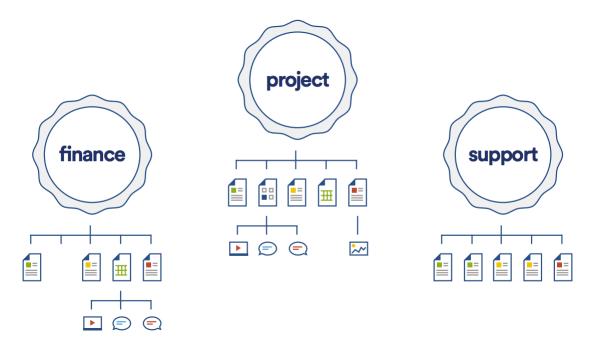
How is content arranged in spaces?

Each space you create, including your personal space, is automatically created with a home page – the first page you'll see when you navigate to the space – which you can edit to suit your needs, and you can create any number of child pages to store the content and information you need. Spaces don't nest – you can't have a space within a space – but pages do, so you can create as many levels of hierarchy as you need using pages.

Each space also has its own blog, allowing you to share news and make announcements. Blog posts are a great way to keep people involved in what's going on in your team or project.

You can set different levels of access for each space using its permissions.

Diagram: Arrangement of spaces, pages, and attachments

View all spaces in Confluence

There are two main ways to view spaces in Confluence:

- **The space directory** - choose **Spaces > Space directory** for a list of all the site and personal spaces you have permission to see. Filter the list of spaces by selecting from the categories on the left of the space directory.
- **The dashboard** – you can make your most useful spaces appear under **My spaces** on the dashboard. Choose the star icon in the space sidebar or space directory to add a space to **My spaces** .

The **Spaces** menu in the header also displays a list of your recently viewed spaces, allowing you to quickly navigate to the things you view most often.

The Spaces List macro allows you to display a list of spaces on a Confluence page, and filter them by category, if you need to.

Create a Space

When you create spaces in Confluence, you can choose to set up a space for each team, project, or a mix of both depending on your needs. You can always move pages and blog posts between spaces later if you need to.

Confluence allows you to set up a blank space, or use the handy space blueprints to make the process of setting up team, documentation, or knowledge base spaces quick and easy.

Create a space

1. Choose **Spaces** > **Create space** in the Confluence header
2. Select the type of space you would like to create, then choose **Next**
3. Enter details of your space including **Name** and **Space Key** , and any other information you're prompted for depending on the blueprint you choose
4. Choose **Create**

Each space you create will automatically have a home page, which you can customise to display relevant information for people viewing the space. Space blueprints have a home page with extra features like team members in the case of the 'team space' blueprint.

On this page:

- Create a space
- Space keys
- Space blueprints
- Space permissions
- Notes
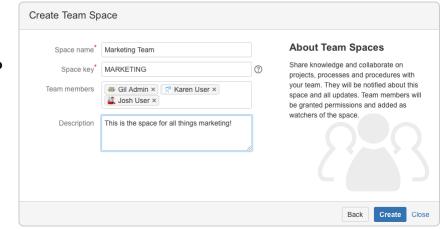
Related pages:

- Spaces
- Create a Personal Space
- Configure the Sidebar
- Edit Space Details
- Delete a Space

You can also set up a personal space to store your own work, which you can share with others or make private. Choose **your pro file picture** at top right of the screen, then choose **Add Personal Space...**. See Create a Personal Space for more information.

Create Team Space

Space name*	Marketing Team
Space key*	MARKETING ⑦
Team members	🐗 Gil Admin × 🐾 Karen User × 👤 Josh User ×
Description	This is the space for all things marketing!

About Team Spaces

Share knowledge and collaborate on projects, processes and procedures with your team. They will be notified about this space and all updates. Team members will be granted permissions and added as watchers of the space.

Back   Create   Close

Space keys

Each Confluence space has a **sp ace key**, which is a short, unique identifier for a space, and forms part of the URL for that space. Each time you create a space, Confluence will create the space key for you, but you can override the default space key if you want to make it something more memorable.

For example, you might give your marketing team's space a key of 'marketing'. You can then navigate directly to the space using a URL like this:

```
http://my.confluenceSite.com/display/marketing/
```

Choosing a space key

Each space key:

- Must be unique
- Can contain any alphanumeric character (A-Z, a-z, 0-9)
- Can be up to 255 characters long

You can't change the space key after you create your space, so choose your space key carefully! Personal spaces use your username as the space key.

Space blueprints

A space created using a space blueprint will have a customised homepage, sidebar and may contain page blueprints or sample content that's specific to that type of space.

When you create a space using a blueprint, a wizard will prompt you to enter information to help set up your space. For example, the Team Space blueprint asks for the members of your team, and displays them on the homepage.

The types of space blueprints available are:

- **Team space** – A great building block if you are using Confluence as an intranet or to manage teams. Team spaces highlight the members of the team, and grant permissions to those users accordingly.
- **Knowledge Base space** – This space blueprint uses search and page labels to make content easier to find, right from the space homepage. It also contains two page blueprints for creating how-to and troubleshooting articles. The templates used in these page blueprints are completely customisable to meet your needs. The Knowledge Base space blueprint also integrates with JIRA Service Desk.
- **Documentation space** – This space blueprint displays the full page tree in the sidebar and hides other sidebar features including blogs and shared links. The homepage uses search and page labels to make content easy to find. Add the 'featured' label to any page you want to highlight on the homepage. This space does not include any page blueprints but you can create and promote templates for your documentation authors to use.

Space permissions

Each space is created with a set of default permissions. The user who created a site space is automatically granted 'space admin' permissions for that space, meaning they can then grant permissions to other users and groups. See Space Permissions Overview for more information.

ℹ️ To create a site space you need the 'Create Space' global permission; to create a personal space you need the 'Personal Space' global permission.

Note: System Administrators can edit the permissions of spaces in their Confluence site at any time.

Notes

- If you've set a specific theme (such as the Documentation or other third party theme) for your whole site, spaces will be created with that theme. You may not see some space blueprint-specific sidebar customisation if you're not using the default theme.
- Administrators can disable individual space blueprints - see Administering Site Templates.
- Spaces can't be nested (you can't have a space within a space), but can be grouped using space categories.
- You can export a whole space, or part of a space, to PDF, HTML or XML .

Create a Personal Space

Your personal space is a place where you can create your own pages and publish your own blog posts. Once you've set up your personal space, Confluence users can reach it by clicking your name in the People Directory. Visit your personal space by choosing your profile picture at the right side of the Confluence header and choosing **Personal Space**.

Create your personal space

ℹ️ To set up your personal space, you need the 'Personal Space' permission which is assigned by a Confluence administrator. See Permissions and Restrictions and Global Permissions Overview.

To create your personal space:

1. Choose **your profile picture** at top right of the screen, then choose **Add Personal Space...**
2. Choose **Create**
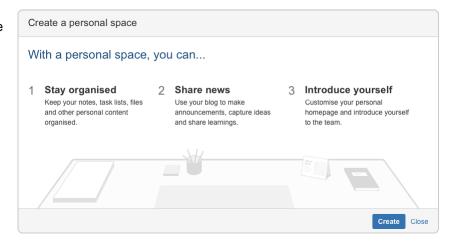
You can now go ahead and update your space by customising the home page, creating more pages, and adding blog posts.

Change the permissions on your space at any time to determine who can and can't access the content. So if you want it to be a private sanctuary, that's no problem.

Create a personal space

With a personal space, you can...

1 **Stay organised**
Keep your notes, task lists, files and other personal content organised.

2 **Share news**
Use your blog to make announcements, capture ideas and share learnings.

3 **Introduce yourself**
Customise your personal homepage and introduce yourself to the team.

Create Close

Change the look and feel of your space

For a *really* personal space, apply a different theme, or modify its colour scheme. You can also add and arrange links in the sidebar.

Space Permissions Overview

Each space in Confluence has its own set of permissions, which can be granted and revoked by a space administrator. The user that creates a space, like when you create your personal space, is automatically an administrator of that space, and they can add other space admins if they want to by adding the **Space Admin** permission to another user or group.

How do you find out who the space admin is for a particular space? Choose **Spaces** > **Space directory** in the Confluence header, then choose the **Space Details**

ⓘ

icon beside a space.

If you're one of those lucky space admins, you can assign permissions for the space to individual users, groups, or anonymous users. To change permissions for a space, choose **Space tools** > **Permissions** from the bottom of the sidebar, then choose **Edit Permissions** to change permission settings.

If your space is using the Documentation theme, choose **Browse** > **Space Admin** > **Permissions**.

ℹ️ Space admins can also manage watchers for that space by choosing **Watch** at the top-right of any page in the space.

Permissions Summary

The following permissions can be assigned for each space:

Permission		The user or group can:
View		View the space's content, including the space's details, its pages, and blog posts.
Pages	**Add**	Create and edit pages in the space.
	Restrict	Apply page-level restrictions.
	Delete	Delete pages in the space.
Blog	**Add**	Add and edit blog posts in the space.
	Delete	Delete blog posts in the space.
Comments	**Add**	Make comments in the space.
	Delete	Delete comments from the space.
Attachments	**Add**	Add attachments to pages in the space.
	Delete	Delete attachments from pages in the space.
Mail	**Delete**	Delete individual mail items.
Space	**Export**	Export content from the space via the space-level export screens. Note that this permission doesn't affect the exporting of a single page's content. Anyone who has permission to view the page also has permission to export its content.
	Admin	Administer this space, including granting permissions to other users. This permission can't be granted to anonymous users (users who aren't logged in).

Notes

ℹ️ Confluence administrators aren't necessarily space administrators.

- A user who has the 'Confluence Administrator' global permission isn't automatically a space admin for a particular space. In order for them to be a space admin, they must belong to a group which has space admin and view permission for the space, or their username must be specifically granted space admin and view permission for the space.
- If you deny all admin access to a space by mistake, so that nobody has access to administer the space any more, you'll need to ask someone with Confluence Administrator global permission to restore the permissions for you.
- A user who has the 'System Administrator' global permission *does* automatically have space admin permissions for all spaces.

Assign Space Permissions

You can assign space permissions to groups or to individual users. You need to be a space administrator to assign space permissions. In addition, Confluence administrators can set the default permissions that will be applied to new spaces as described below.

To access the permissions for a space:

1. Go to the space and choose **Space tools** > **Permissions** from the

bottom of the sidebar
2. Choose **Edit Permissions**

⌄ Using the Documentation theme?

1. Go to the 'Space Permissions' page:
   - Choose **Browse** > **Space Admin** from the header

     Note: The **Space Admin** option appears only if you have s pace admin permissions, or if you're part of the 'confluence-administrators' group.
   - Choose **Permissions** from the space administration options.

2. Choose **Edit Permissions**

The Edit Space Permissions page is divided into the following sections:
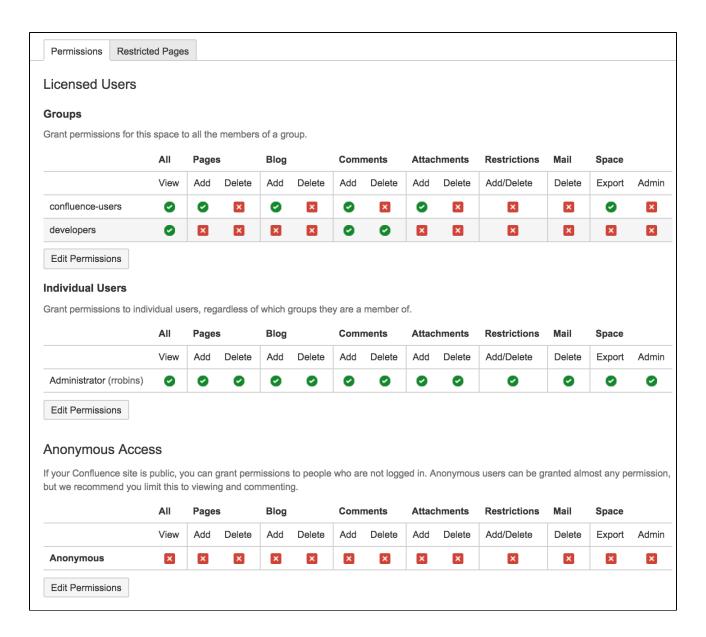
- **Licensed Users -** this is where you grant permissions to groups and individual users.
- **Anonymous Access** - this is where you grant permissions to users who are not logged in (essentially making the space public)

On this page:

- Assign space permission s
- Remove space permission s
- Anonymou s users
- Set default space permission s
- Manage and recover space admin permission s
- Managing unlicensed users from JIRA Service Desk

Related pages:

- Space Permission s Overview
- Global Permission s Overview

Permissions | Restricted Pages

Licensed Users

Groups

Grant permissions for this space to all the members of a group.

	All	Pages		Blog		Comments		Attachments		Restrictions	Mail	Space	
	View	Add	Delete	Add	Delete	Add	Delete	Add	Delete	Add/Delete	Delete	Export	Admin
confluence-users	✓	✓	✗	✓	✗	✓	✗	✓	✗	✗	✗	✓	✗
developers	✓	✗	✗	✗	✗	✓	✓	✗	✗	✗	✗	✗	✗

Edit Permissions

Individual Users

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

	All	Pages		Blog		Comments		Attachments		Restrictions	Mail	Space	
	View	Add	Delete	Add	Delete	Add	Delete	Add	Delete	Add/Delete	Delete	Export	Admin
Administrator (rrobins)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Edit Permissions

Anonymous Access

If your Confluence site is public, you can grant permissions to people who are not logged in. Anonymous users can be granted almost any permission, but we recommend you limit this to viewing and commenting.

	All	Pages		Blog		Comments		Attachments		Restrictions	Mail	Space	
	View	Add	Delete	Add	Delete	Add	Delete	Add	Delete	Add/Delete	Delete	Export	Admin
Anonymous	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗

Edit Permissions

Assign space permissions

To add a new user or group to the permissions list, search for either a group or user in their respective sections and choose **Add**. The group or user will appear in the list of groups; select the appropriate checkboxes to assign them permissions, then choose **Save all**.

If you're unsure of a group or user's name, you can search by choosing the

🔍

icon. You can use an asterisk '*' as a wildcard if you need to.

To bulk assign or revoke permissions, choose either **Select All** or **Deselect All**.

There's no way to change permissions for multiple spaces at once. Permissions are managed on a space by space basis.

Remove space permissions

To remove a user or group from the space permissions list, deselect all the checkboxes for that user or group and save the changes. The user or group won't appear in the list once you save.

Anonymous users

Anonymous users are users that aren't signed in to your Confluence site, and, by default, they won't have access to any of your content. If you want to grant access to a space for anonymous users, you can edit the

'Anonymous Access' section at the bottom of the permissions page.

Tick or un-tick the checkboxes to suit your preference for anonymous users, then choose **Save All** to apply the changes.

ⓘ You can't grant space administration rights or page restriction rights to anonymous users.
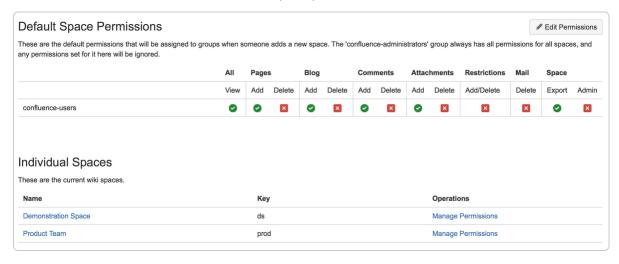
Set default space permissions

If you're a Confluence Administrator, you can set the default permissions that will be applied to new spaces. Default permissions are configurable for groups only, not for individual users or anonymous users.

To set the default space permissions:

1. Go to
   ⚙ ▾
   > **General Configuration** > **Space Permissions**
2. Choose **Edit Permissions**

Screenshot: Default space permissions

Manage and recover space admin permissions

Users with System Administrator permissions are able to manage permissions for spaces, including adding or removing Space Admin permissions for a space.

To manage space permissions:

1. Go to
   ⚙ ▾
   > **General Configuration** > **Space Permissions**.
2. Locate the space in the individual spaces list and choose **Manage Permissions**.

There may be some instances where a space administrator has removed Space Admin permissions from all other users and groups for a space, meaning no other user can administer the space. Users with Confluence Administrator permissions can recover permissions for the space in this instance.

To recover Space Admin permissions:

1. Go to
   ⚙ ▾
   > **General Configuration** > **Space Permissions**.
2. Locate the space in the Individual Spaces list and choose **Recover Permissions**

The user will then be able to choose **Manage Permissions**, and add any appropriate permissions to the space. Requests to recover permissions are recorded in the Confluence log files.

Managing unlicensed users from JIRA Service Desk

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

When a space is accessible to all active users, you'll see the following message in the space permission screen.

Permissions	Restricted Pages

🔓 **Any active user can view this space**

Any active user can view pages in this space, including users who don't have a Confluence license. This was enabled through JIRA Service Desk so people can view knowledge articles when raising a help request.

Edit this permission

It's important to note that this permission overrides all existing space permissions, so any logged in Confluence user will also be able to see the space (regardless of their group membership). This is due to the way Confluence inherits permissions.

You can **edit this permission** at any time to revoke access to a space. It can only be re-enabled from JIRA Service Desk.

Active users, who don't hold a Confluence license, have very limited access to Confluence. They can view pages, but can't like, comment, edit, view the dashboard, use the space directory, see user profiles or search your full site.

See Use JIRA applications and Confluence together for more information about JIRA Service Desk integration.

Change the Look and Feel of a Space

You can customise the 'look and feel' of a space on your Confluence site through options available in the **Space Administration** menu. By default, the look and feel of a space is based on site-wide settings configured from the Administration Console.

You need to be a space administrator to change the look and feel of a space.

- Change the Space Logo
- Apply a Theme to a Space
- Edit a Space's Colour Scheme

> **Related pages:**
>
> - Changing the Look and Feel of Confluence
> - Spaces
> - Styling Confluence with CSS

Change the Space Logo

In Confluence, you can replace the default space logo with an image of your choice. If you have team spaces, you can use an icon that represents the team, or if it's a space for a client, add their company logo to really impress them.

The instructions below apply to site spaces. For your personal space, your p rofile picture is used as the space icon.

You need to be a space administrator to replace a space's logo.

> **Related pages:**
>
> - Configure the Sidebar
> - Change the Look and Feel of a Space

To change a space's logo, in spaces using the default theme:

1. Go to the space and choose **Spa ce tools** > **Configure sidebar** fro m the bottom of the sidebar
2. Choose the edit icon

   next to the space name
3. Choose **Upload an image**
4. Select an image from your computer
5. Adjust the size of the image to fit within the highlighted circle
6. Choose **Save**

To change a space's logo, in spaces using the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **Change Space Logo** in the left-hand panel
3. Click **Choose File**
4. Select an image from your computer
5. Choose **Upload Logo**

Notes

- **Minimum dimensions:** Space logos are set at 48 x 48px. Logos less than these dimensions will be centred with whitespace around them.
- **Space logos in the default theme are circular.**
- **Dimensions of image when using the Documentation theme:** The documentation theme does not provide an option to resize or crop your image. Uploading a square image will give the best results.

Apply a Theme to a Space

Themes allow you to personalise the 'look and feel' of Confluence. You can apply a theme to your entire Confluence site or to individual spaces. Choose a specific theme if you want to add new functionality or significantly alter the appearance of Confluence.

Confluence comes with a selection of themes. In addition, a site administrator can install new themes as plugins via the Confluence Administration Console. Provided that the theme is installed into your Confluence site, any space administrator can apply a theme to a space.

By default when you create a new space, the space will have the Confluence default theme.

Related pages:
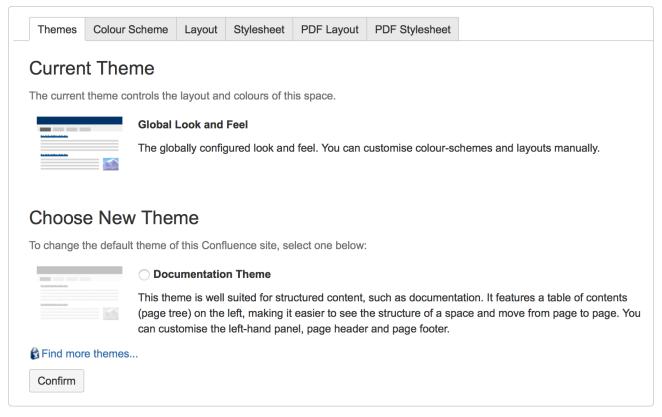
- Applying a Theme to a Site
- Edit a Space's Colour Scheme
- Change the Look and Feel of a Space

To apply a theme to a space:

1. Go to the space and choose **Space tools** > **Look and Feel** from the bottom of the sidebar
   You'll need Space Admin permissions to do this.
2. Choose **Themes** and select a theme option
3. Choose **Confirm**

Screenshot: Applying a theme

If your space is using the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **Themes** from the space administration options
3. Select a theme option
4. Choose **Confirm**

Want to customise Confluence or make it even more beautiful? Try a Confluence theme from the Atlassian Marketplace.

The Documentation Theme

The Documentation theme is going away in Confluence 6.0.
Read more about this change.

The Documentation theme is one of the themes bundled with Confluence. It provides an inbuilt table of contents for your space, a configurable header and footer, and text styles suited to documentation. You can also configure it to restrict the search results to the current space, rather than searching the entire Confluence site.

⚠ Please note, the Documentation theme doesn't support the JIRA links feature.

On this page:

- Customise the Document ation Theme
- Search the space or the site
- Hints and tips
- Notes

Related pages:

- Configure the Document ation Theme
- Space Jump Macro
- Move and Reorder Pages
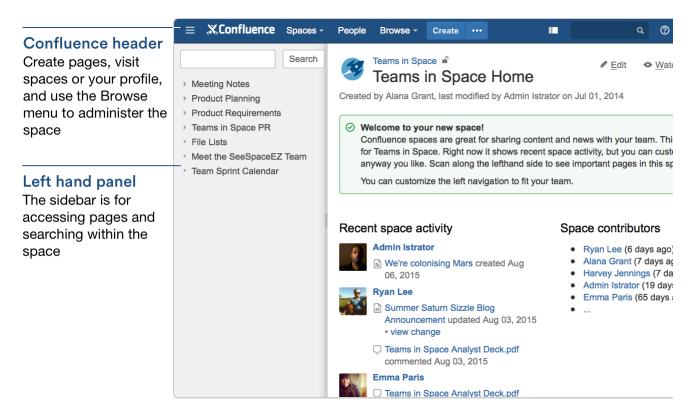
Quick guide to using the Documentation theme:

- The left-hand panel contains a page tree. This is a table of contents that shows the pages in the space.
- Pages with arrow symbols to their left have child pages. Choose the arrow(s) to show and hide child pages.
- Drag the thick vertical bar between the left-hand panel and the page to change its width.
- Choose the left-hand panel icon

  in the header, next to the search field, to hide or show the left-hand panel.
  Alternatively, press '**[**' on your keyboard to show/hide the left-hand panel.
- Overriding a space-restricted search: If your theme is configured to restrict the search to the current space, you can enter '**all:**' and your search term to search the entire site.

Confluence header
Create pages, visit spaces or your profile, and use the Browse menu to administer the space

Left hand panel
The sidebar is for accessing pages and searching within the space

Customise the Documentation Theme

A space administrator can customise the Documentation theme as follows:

- Change the content in the left-hand panel and add a header and footer to the page.
- Restrict the search so that it will show results from the current space only, not from the entire Confluence site.

For more, see Configure the Documentation Theme

Search the space or the site

When using the Documentation theme, the Confluence search offers a few options as described below.

Using the search box at the top right of the page:

- By default, the main Confluence search is configured to search the entire Confluence site.
    - You will see the words '**Search Confluence**' in the search box at top right of the page.
    - The Confluence search will look for matches in the entire Confluence site. This is the default behaviour for other themes too.
- A space administrator can configure the Documentation theme to restrict the search to the current space.
    - You will see the words '**Search this space**' in the search box at top right of the page.
    - The search will return results from the current space only.
    - You can override the search restriction. Enter '**all:**' and your search term to search the entire site. For example, enter the following into the search box at top right of the page to search the entire site for 'technical writing':

```
all: technical writing
```

Using the search box in the left-hand panel:

- By default, the Documentation theme's left-hand panel includes a search box. Enter your search term there, to restrict the search to the current space. Specifically, this will search only the pages that are children of the space's home page.
- If your administrator has restricted the main search to the current space, there will not be a search box in the left-hand panel.

Hints and tips

Below are some hints that you may find useful when using the Documentation theme.Where can I see a working example of the Documentation theme?

The Confluence documentation uses the Documentation theme. If you are reading this documentation online on the Atlassian documentation wiki, you are seeing a working version of the Documentation theme.

Which pages will appear in the Documentation theme's table of contents?

The theme constructs the page tree in the left-hand panel from all pages that are child pages of the space's home page. Each space has a single page designated as the 'Home' page. You can specify the home page in the space administration section.

Help, my pages do not appear in the Documentation theme table of contents

Cause 1: Your pages are not under the space's home page. The most probable reason why your pages do not appear in the page tree in the left-hand panel is this: The theme constructs that table of contents from all pages that are child pages of the space's home page. If your pages are above the home page in the page tree, they will not appear in the left-hand panel.

There are two ways to fix the problem:

- You can change the designated space home page in the space administration section.
- Or you can drag and drop all your pages to make them children of the current home page. You can drag and drop pages in the 'Pages' section of the space 'Browse' screen. See Move and Reorder Pages.

Cause 2: Problem with upgrade from Confluence 3.1, with Documentation theme as plugin, to Confluence 3.2 or later with Documentation theme bundled. If your existing Confluence installation already has the Documentation theme plugin installed, you may find that after upgrading to Confluence 3.2 the left-hand navigation bar is empty in the spaces that use the theme. The fix is to enable all modules of the Documentation theme plugin. See the knowledge base article.

Help, my left-hand panel has disappeared

If your entire left-hand panel has disappeared when using the Documentation theme, this is probably because you have clicked the sidebar icon

at top right, next to the search box. Click the icon again to restore the panel.
Can I change the default width of the left-hand navigation panel?

The Confluence user interface does not offer a way to change the default width of the left-hand navigation panel supplied by the Documentation theme. Users can change the width by dragging the middle bar, but the default width is not configurable. This post on Atlassian Answers gives a way to do it with CSS: Documentation Theme - Default Width (in px) of left panel.

Notes

- The Confluence default theme supplies a **sidebar**, which is different to the left-hand panel in the Documentation theme. For information about the sidebar, see Configure the Sidebar.
- The Documentation theme supplies a **Browse** menu in the Confluence header, which gives access to the space administration and advanced options.

Configure the Documentation Theme

The Documentation theme is going away in Confluence 6.0.
Read more about this change.

The Documentation theme is designed for spaces containing technical documentation, but you may find it useful for other types of structured content. It provides a table of contents for your space, a configurable header and footer, and text styles suited to documentation. (See features below.)

Quick guide to applying and customising the Documentation theme:

- Go to the space and choose **Space tools** > **Look and Feel** from the bottom of the sidebar
- Choose **Themes**.
- Select **Documentation Theme** and choose **Confirm**.
- If you want to customise the theme, choose **Configure theme**.
  - Select or deselect the default page tree.
  - Select or deselect the space-restricted search.
  - Enter the text and wiki markup for your custom left-hand panel header and footer.
  - Choose **Save**.

The rest of this page gives more details of the above procedure.
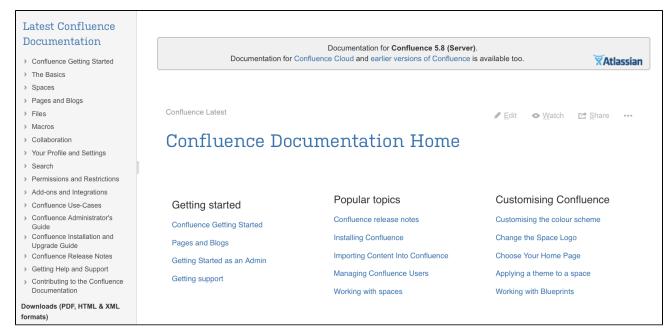
Screenshot: The Documentation theme showing the space sidebar replaced by a navigation sidebar.
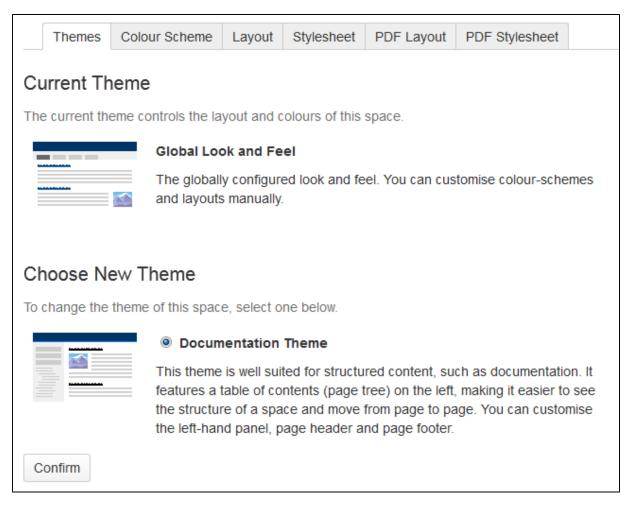
Applying the Documentation theme to your space

Follow the steps below to apply the Documentation theme to your space. All pages in the space will start using the theme immediately.

To apply a theme to a space:

1. Go to the space and choose **Space tools** > **Look and Feel** from the bottom of the sidebar
   You'll need Space Admin permissions to do this.
2. Choose **Themes** and select a theme option
3. Choose **Confirm**

   Screenshot: Applying a theme

If your space is using the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **Themes** from the space administration options
3. Select a theme option
4. Choose **Confirm**

Applying the Documentation theme to your site

If you have site administrator permissions, you can apply the theme at site level. It will then be the default theme for all spaces in the site. See the administrator's guide to applying a theme.

Customising the Documentation theme

The theme works well without any customisation. If necessary, you can restrict the search to just one space, change the content of the left-hand navigation panel and add your own header and footer. The following instructions assume you have already applied the Documentation theme.

To customise the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **Themes** in the left-hand panel under the heading 'Look and Feel'.
3. Choose **Configure theme** in the yellow area of the 'Current Theme' section at the top of the page. See screenshot below. The 'Documentation Theme Configuration' screen appears. See screenshot below.
4. Select or deselect the **Page Tree** check box. This determines whether your space will display the default search box and table of contents (page tree) in the left-hand panel.

5. Select or deselect the **Limit search results to the current space** check box.
   - If you select the check box:
     - The Confluence search will look for matches only in the current space by default. Users can override this restriction when entering their search term. See The Documentation Theme.
     - The default page tree in the left-hand panel will not include a search box.
   - If you do not select the check box:
     - The Confluence search will look for matches in the entire Confluence site. This is the default behaviour for other themes too.
     - The default page tree in the left-hand panel will include a search box, which restricts search results to the current space.
6. Enter text, images, macros and other wiki markup into any or all of the three text boxes.
   🛈 You can use the Include or Excerpt Include to include re-usable content into your footer. See hint below.
   - 'Navigation' – This text box contains content for the left-hand panel.
     - If the **Page Tree** check box is selected, the navigation panel contains the default search box and page tree. Any content you enter into the 'Navigation' text box will appear *above* the page tree and search box.
     - You can include your own content *underneath* the page tree as well as above. See hint below. In summary: Deselect the **Page Tree** check box. Insert your own page tree using the Pagetree macro, then add your own content under the macro.
   - 'Header' – This text box contains content for a page header that will appear above the page title on all pages in the space. See example screenshot below.
   - 'Footer' – This text box contains content for a page footer that will appear after the comments and above the site footer, on all pages in the space. See example screenshot below.
7. Choose **Save**.

Screenshot: The 'Configure theme' option

Current Theme

The current theme controls the layout and colours of this space.

Documentation Theme

This theme is well suited for structured content, such as documentation. It features a table of contents (page tree) on the left, making it easier to see the structure of a space and move from page to page. You can customise the left-hand panel, page header and page footer.

Configure theme - further customise this theme's options

Choose New Theme

To change the theme of this space, select one below.

○ Global Look and Feel

The globally configured look and feel. You can customise colour-schemes and layouts manually.

 Find more themes...

Confirm

Screenshot: Customising the Documentation theme

Navigation

☐ Page Tree
   Show the page tree in the sidebar.

☑ Limit search results to the current space
   Limit search results returned by the top right search field to the current space.

Navigation
```
*[Latest Confluence
Documentation|Confluence
Documentation Home]*

{pagetree:root=Confluence
Documentation
Home|searchBox=false}

*[Downloads (PDF, HTML &
```
Add your own links and content to the sidebar using wiki markup.

Messages

Header
```
{section}
{column:width=1%}{column}
{column}{panel}
*{include:_MajorReleaseNumb
erConfluence|nopanel=true}*
 -- This documentation is
for *Confluence 5.0*
developer (EAP) releases
only. Not for you? *[See
```
Display a banner or text at the top of every page in this space by entering wiki markup here.

Footer
```
{include:_Documentation
Footer|nopanel=true}
{include:ALLDOC:_Copyright
Notice|nopanel=true}
```
Display a banner or text at the bottom of every page in this space by entering wiki markup here.

[Save]   Cancel

Customising the theme at site level

If you have site administrator permissions, you can apply and customise the theme at site level. The customisation options are the same as the space level options, as described above.

Features of the Documentation theme

Here's a summary of the features that the Documentation theme provides:

- By default, the left-hand panel contains a search box and a table of contents (page tree) showing all the pages in your space. Specifically, it shows the pages that are children of the space's home page.
- The left-hand panel is fully customisable. You can choose to include or exclude the search box and page tree. You can enter your own text, images and wiki markup.
- People viewing the page can drag the thick bar between the left-hand panel and the content, to increase or decrease the width of the panel. They can also remove the panel altogether, by clicking the sidebar icon

  at top right, next to the search box.
- The left-hand and right-hand panels scroll independently of each other.
- The page title is neatly above the page content, and not uncomfortably above the navigation panel as tends to happen when you insert the navigation panel yourself.
- Because the left-hand panel is part of the theme, it will be upgraded whenever Confluence is upgraded. There is no need to remove and then re-apply your customisations on each upgrade, as you would do if you added your own left-hand navigation bar.
- The theme also provides a customisable header and footer, which will appear at the top and bottom of every page.
- The text and heading styles are designed to enhance the content typically found in a documentation space.
- You can use the Space Jump macro to link from a page in one wiki space to a page with the same name in another space, without knowing the name of the page when you create the link.
- You can limit the Confluence search results to the current space. If you choose this option, the Confluence search will look for matches only in the current space by default. Users can override the restriction. See The Documentation Theme.

Hints and tips

Below are some hints that you may find useful when using the Documentation theme.Where can I see a working example of the Documentation theme?

The Confluence documentation uses the Documentation theme. If you are reading this documentation online on the Atlassian documentation wiki, you are seeing a working version of the Documentation theme.

Which pages will appear in the Documentation theme's table of contents?

The theme constructs the page tree in the left-hand panel from all pages that are child pages of the space's home page. Each space has a single page designated as the 'Home' page. You can specify the home page in the space administration section.

Help, my pages do not appear in the Documentation theme table of contents

Cause 1: Your pages are not under the space's home page. The most probable reason why your pages do not appear in the page tree in the left-hand panel is this: The theme constructs that table of contents from all pages that are child pages of the space's home page. If your pages are above the home page in the page tree, they will not appear in the left-hand panel.

There are two ways to fix the problem:

- You can change the designated space home page in the space administration section.
- Or you can drag and drop all your pages to make them children of the current home page. You can drag and drop pages in the 'Pages' section of the space 'Browse' screen. See Move and Reorder Pages.

Cause 2: Problem with upgrade from Confluence 3.1, with Documentation theme as plugin, to Confluence 3.2 or later with Documentation theme bundled. If your existing Confluence installation already has the Documentation theme plugin installed, you may find that after upgrading to Confluence 3.2 the left-hand navigation bar is empty in the spaces that use the theme. The fix is to enable all modules of the Documentation theme plugin. See the knowledge base article.

Help, my left-hand panel has disappeared

If your entire left-hand panel has disappeared when using the Documentation theme, this is probably

because you have clicked the sidebar icon

at top right, next to the search box. Click the icon again to restore the panel.
Can I change the default width of the left-hand navigation panel?

The Confluence user interface does not offer a way to change the default width of the left-hand navigation panel supplied by the Documentation theme. Users can change the width by dragging the middle bar, but the default width is not configurable. This post on Atlassian Answers gives a way to do it with CSS: Documentation Theme - Default Width (in px) of left panel.

Hiding pages from the left-hand table of contents

You can 'hide' pages by putting them at the same level as or higher than the space home page.

Each space has a single page designated as the 'Home' page. You can specify the home page in the space administration section.

The theme constructs the page tree in the left-hand panel from all pages that are child pages of the space's home page.

If your documentation pages are at the same level as the space home page, they will not appear in the left-hand navigation bar. So you can 'hide' pages by putting them at the same level as or higher than the space home page. The pages will show up in the search results and people can see the content if they open the page, but the pages will not appear in the left-hand panel.

More detail: The theme uses the Pagetree macro to produce the table of contents. When entering the Pagetree macro, you can choose the top page in the page tree. The Documentation theme chooses the space home page as the top page.

Using reusable content in your header, footer or sidebar

You can use any text or wiki markup in your theme header, footer or left-hand panel. One useful hint is to use the Include or Excerpt Include to include re-usable content into your footer.

The screenshot above shows the theme customisation options, with examples of the macros used to include content from other pages. And the example screenshot also above shows the resulting header, footer and left-hand panel.

Adding content below the page tree in your sidebar

If you want to include your own content underneath the page tree, you can deselect the 'Page Tree' check box, add your own page tree using the Pagetree macro in the 'Navigation' text box, and then add your own content under the macro.

The screenshot above shows the theme customisation options, with the default page tree deselected and a custom page tree inserted, along with additional content in the left-hand panel. The example screenshot also above shows the resulting left-hand panel.
Adding an expanding All Versions section to the sidebar

If you want to include a an expanding list of links to other spaces, as we have included in this space (see 'Docs for all Confluence releases' in the sidebar), you can use an Expand Macro and an Include Page Macro in the 'Navigation' text box.   For example:

```
*[Docs for all Confluence releases|_Latest Versions of Confluence
Documentation]*
{expand:Choose a version...}
{include:_Latest Versions of Confluence Documentation}
{expand}
```

If you would like this to display below the page tree, follow the steps above.

Adding underlines to your links

By default, the Documentation theme does not underline hyperlinks. If you prefer to have your links underlined, you can edit the CSS stylesheet for your space and add the following CSS code:

```
.wiki-content a:link, .wiki-content a:visited, .wiki-content a:active
{
  text-decoration: underline;
}
```

To edit a space's CSS style sheets:

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

If your space uses the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Click **Stylesheet** in the left-hand panel under the heading 'Look and Feel'.
3. Choose **Edit**.
4. Paste your custom CSS into the text field.
5. Save your changes. The new CSS will be visible on all content pages in the space.

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

⚙ ▾

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

Jumping to the same page in another space

The {spacejump} macro is provided along with the Documentation theme. You can use space jumping to link from a page in one wiki space to a page with the same name in another space, without knowing the name of the page when you create the link. When a reader is viewing a page and chooses the link provided by the macro, they will go to a page with the same name, but in the space specified in the macro. See more about the Space Jump macro.
Hiding the left-hand panel completely

It's not possible to remove the left-hand panel entirely, using the user interface supplied by the theme. There is an improvement request here:

CONF-25923 - Provide option to hide the left-hand panel entirely in the Documentation theme
RESOLVED   If you like, you can comment on and/or vote for that request. In the meantime, Atlassian Answers is a good place to ask the question and see if other people can help you with custom code to remove the panel.

Notes

- The Confluence default theme supplies a **sidebar**, which is different to the left-hand panel in the Documentation theme. For information about the sidebar, see Configure the Sidebar.
- The Documentation theme supplies a **Browse** menu in the Confluence header, which gives access to the space administration and advanced options.

Switch to the Default Theme

If you use the Documentation theme in your spaces, it's time to consider switching to the Confluence default theme.
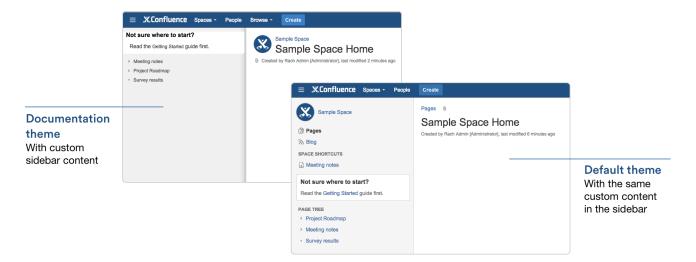
A growing number of great features (such as JIRA links button, space shortcuts and sticky table headers) are not available in the Documentation theme, and we plan to retire this theme altogether in the near future.

This page will show you how to turn on the default theme, and keep any existing header, footer or sidebar content.

Here's an example of the documentation theme, and default theme with the same custom content side by side:

On this page:

- Understanding global look and feel
- Switch your space to the default theme
- Make the default theme available across the whole site
- Hints and tips for the sidebar
- Edit the header, footer or custom sidebar content
- Documentation theme features that aren't available in the default theme
- Troubleshooting

Documentation theme
With custom sidebar content

Default theme
With the same custom content in the sidebar

Understanding global look and feel
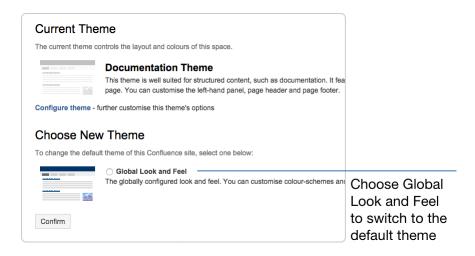
Before you start, it's important to understand the idea of "global look and feel". Confluence allows you to apply most themes to an individual space, or to the whole site. When a space says it is using the 'global look and feel' it means it is using whichever theme has been applied for the whole site, and inheriting any customisations from the site level.

The Confluence default theme, however, can only be applied to a space if it is currently applied to the entire site (it must be the current global look and feel).

Switch your space to the default theme
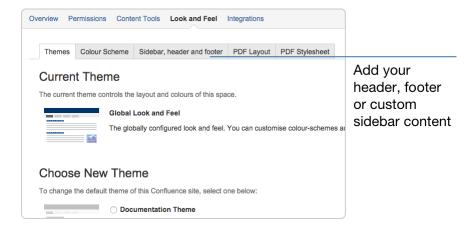
To turn on the default theme and keep any existing header, footer or custom sidebar content:

1. Go to **Browse** > **Space Admin** > **Theme** and then **Configure theme**.
2. Copy the wiki markup from the **Navigation**, **Header** and **Footer** fields to somewhere safe (like a document or text file) then **Cancel**.
   (If these fields are empty, you can go straight to the next step).
3. Select **Global look and feel** (that's the default theme) and choose **Confirm**.

Current Theme

The current theme controls the layout and colours of this space.

Documentation Theme

This theme is well suited for structured content, such as documentation. It fea[...]
page. You can customise the left-hand panel, page header and page footer.

Configure theme - further customise this theme's options

Choose New Theme

To change the default theme of this Confluence site, select one below:

○ **Global Look and Feel**

The globally configured look and feel. You can customise colour-schemes an[...] — Choose Global Look and Feel to switch to the default theme

Confirm

(If **Global Look and Feel** is already listed as your current theme you'll need to follow the steps below to change your Site theme first).

4. The default theme is now enabled - you'll notice the space admin options have moved.
   Choose the **Header, Footer and Sidebar** tab.

Overview   Permissions   Content Tools   **Look and Feel**   Integrations

Themes   Colour Scheme   Sidebar, header and footer   PDF Layout   PDF Stylesheet

Current Theme

The current theme controls the layout and colours of this space.

Global Look and Feel

The globally configured look and feel. You can customise colour-schemes a[...] — Add your header, footer or custom sidebar content

Choose New Theme

To change the default theme of this Confluence site, select one below:

○ **Documentation Theme**

5. Paste the wiki markup for your custom sidebar content (navigation), header and footer into the fields provided then **Save**.
6. Use the space name link in the sidebar to get to the space homepage and check the appearance of your space.

≡ ✕**Confluence**   Spaces ▾   Pe[...]

✕ Sample Space — Back to your space's home page

📄 Pages
🔊 Blog

SPACE SHORTCUTS

📄 Meeting notes — Customisable shortcuts and links

Not sure where to start?
Read the Getting Started guide first.

PAGE TREE
› Project Roadmap
› Meeting notes
• Survey results

⚙ Space tools ▾   « — Space admin options

The default theme includes some extras in the sidebar, such as shortcuts and index pages for blogs and blueprints. Your custom content will appear between the shortcuts and page tree.

To choose what you'd like to show or hide:

1. Go to **Space Tools** (at the bottom of the sidebar) and select **Configure Sidebar**.
2. Use the

   ⊖

   and

   ⊕

   icons to show and hide each section.
3. This is where you can change the site logo too.

You can't hide the space name, logo or page tree.

Make the default theme available across the whole site

If the Documentation theme is enabled for the entire site, you won't see an option to switch to the default theme in each space. This means that before you can switch each space over to the default theme, you'll need to make it available for the whole site.

To make the default theme available across your whole site:

1. Go to

   ⚙ ▾

   > **General Configuration**> **Themes**
2. Choose **Configure theme** beside the Documentation theme
3. Copy any wiki markup from your header, footer and sidebar fields to somewhere safe (like a document or text file) then **Cancel**.
   (If these fields are empty, the Documentation theme has not been customised for the entire site)
4. Select the **Default theme** and choose **Confirm**.
   The default theme is now the global look and feel. If you did not need to copy over any wiki markup, the process is now complete.
5. If you have wiki markup to migrate, go to

   ⚙ ▾

   > **General Configuration** > **Sidebar, header and footer**
6. Paste the wiki markup for your custom sidebar content (navigation), header and footer into the fields provided then **Save**.

All new spaces will now be created with the default theme.

The impact on existing spaces depends on how the theme is applied:

- All spaces that used the Documentation theme because it was the global look and feel will now use the default theme. You don't need to do anything.
- All spaces that have the Documentation theme explicitly applied (for example to enable different wiki markup in the sidebar, header or footer) the documentation theme will still be applied, and you'll need to migrate each one separately.

Hints and tips for the sidebar

Here's a few hints and tips for making your sidebar migration as easy as possible.

Do you need to use wiki markup?

If the wiki markup in the Documentation theme for your space was simple, for example just adding a page tree and a couple of links to useful spaces or external sites, chances are you don't need to continue to customise the sidebar at all. The default theme comes with a page tree option, and you can add links to pages, spaces or websites. These space shortcuts are much easier to keep up to date (and are safe from page renames).

Are you using space CSS to style the theme?

If you're overriding the space CSS to change the appearance of the doc theme sidebar, you'll probably need

to make some changes to your CSS, as some class and ID names are different between the two themes.

For example, instead of specifying `#splitter-sidebar`, you'll need to use `.acs-side-bar`.

Edit the header, footer or custom sidebar content

You can update the wiki markup for the sidebar, header and footer at any time. Go to **Space Tools** > **Look and Feel** > **Sidebar, Header and Footer**.

These fields use wiki markup, check our guide to wiki markup for help, or check out some common customisations below.

˅ Add a search field to your sidebar...

To add a search field to the sidebar:

1. Go to **Space Tools** > **Look and Feel** > **Sidebar, Header and Footer**
2. In the **Sidebar** field add the wiki markup for either the Livesearch or Page Tree Search macros

```
{livesearch}
```

```
{pagetreesearch}
```

3. The search field will appear above the page tree.

˅ Add a panel with some content...

To add a panel with some custom content to the sidebar:

1. Go to **Space Tools** > **Look and Feel** > **Sidebar, Header and Footer**
2. In the **Sidebar** field add the wiki markup for the Panel Macro then add some text:

```
{panel}This is some custom content to appear in the sidebar{panel}
```

3. Your content will appear above the page tree.

Documentation theme features that aren't available in the default theme

Not all documentation theme features are available in the default theme.

Limit search results to the current space

We've decided not to add the ability to restrict the Confluence search on the header to a single space for now. You will still be able to add the Live Search or Page Tree Search macros to the sidebar or homepage of your space, and set these macros to only return results from the current space.

Space jump macro

This macro is part of the documentation theme. It is still available for now, but is likely to be removed when we remove the Documentation theme.

Child pages below the page content

In the Documentation theme, children of the current page are listed below the page content. This is not available in the default theme, however you could choose to add a Children Display Macro to the footer in the default theme.

View pages alphabetically

In the Documentation theme you could choose to view all pages in a site alphabetically. This isn't available in the default theme.

Troubleshooting

Find out which spaces are using the Documentation theme

There's no simple way to see a list of spaces using the Documentation theme in Confluence itself, however if you have a very large site, your Confluence Administrator can use the following query to get a list of spaces directly from the database.

```
SELECT *
FROM BANDANA B, SPACES S
WHERE B.BANDANAKEY='atlassian.confluence.theme.settings'
AND S.SPACEKEY=B.BANDANACONTEXT
AND B.BANDANAVALUE LIKE ('%documentation%')
ORDER BY S.SPACENAME;
```

This query will only find spaces that are explicitly using the documentation theme. It doesn't include spaces using the global look and feel (these spaces automatically change when you change the Site Theme, you wont need to change the theme space by space).

Check for layout changes

If the underlying layout templates for your spaces have been customised, you may find that your space looks strange or broken when you apply the default theme.

If this happens, you can reset the templates back to the default. You'll need to be a Site Administrator to do this.

1. Switch back to the Documentation theme temporarily.
   (If you're unable to use the space navigation, use this URL, replacing YOURSPACEKEY with the space key for the space.

   ```
   http://<yoursite>/spaces/choosetheme.action?key=YOURSPACEKEY
   ```

2. Go to **Browse** > **Administration** > **Layouts** or use the following link if you're unable to get there via the space navigation.

   ```
   http://<yoursite>/spaces/listdecorators.action?key=YOURSPACEKEY
   ```

3. Choose **Reset Default** next to any template that have been customised.
4. Return to the **Themes** page and try applying the default theme again.

Edit a Space's Colour Scheme

Confluence allows you to customise the colour scheme of a space. By default, a space's colour scheme is based on global settings configured from the Administration Console.

You need to be a space administrator to edit a space's colour scheme.

To change the colour scheme for a space:

1. Go to the space and choose **Space tools** > **Look and Feel** from the bottom of the sidebar
2. Choose **Colour Scheme**
3. Choose **Select** next to a scheme listed under **Custom Colour Scheme** (if not already selected)
4. Choose **Edit**
5. Enter standard HTML/CSS2 colour codes or use the colour-picker

Related pages:

- Change the Look and Feel of a Space
- Apply a Theme to a Space

[icon] to choose a new colour from the palette provided

If your space uses the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **Colour Scheme** under the heading **Look and Feel**
3. Follow the steps above to select a custom colour scheme and edit the colours

The colour scheme allows you to edit the colours of UI elements including the top bar, tabs and backgrounds.

Some UI elements below are for specific themes, and colour changes may not take effect for other themes.
- **Top Bar** - the top navigation bar background
- **Top Bar Text** - the text on the top navigation bar
- **Header Button Background** - buttons on the top navigation bar (e.g. Create button)
- **Header Button Text** - the text on buttons on the top navigation bar
- **Top Bar Menu Selected Background** - background colour of top navigation bar menu items when selected (e.g. spaces)
- **Top Bar Menu Selected Text** - text colour of top navigation bar menu items when selected
- **Top Bar Menu Item Text** - text on top navigation bar drop down menus (e.g. help or cog)
- **Menu Item Selected Background** - highlight colour on top navigation bar drop down menu items
- **Menu Item Selected Text** - text colour on highlighted top navigation bar drop down menu items
- **Search Field Background** - the background colour of the search field on the header
- **Search Field Text** - the colour of the text in the search field on the header
- **Page Menu Selected Background** - the background colour of the drop down page menu when selected
- **Page Menu Item Text** - the text of the menu items in the drop down page menu
- **Heading Text** - all heading tags throughout the space
- **Space Name Text** - the text of the current space name located above the page title
- **Links** - all links throughout the space
- **Borders and Dividers** - table borders and dividing lines
- **Tab Navigation Background** - the background colour of the tab navigation
- **Tab Navigation Text** - the text of the tab navigation when highlighted
- **Tab Navigation Background Highlight** - the background colour of the tab navigation when highlighted
- **Tab Navigation Text Highlight** - the text of the tab navigation elements when highlighted

Screenshot: Editing the colour scheme

Custom Colour Scheme

A custom colour scheme which can be edited.

Selected

Edit

The following colours can be customised for this colour scheme.

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

Reset   Save   Cancel

Handy Hint
If you mess things up, just choose **Reset** then try again.

Administer a Space

If you're a space admin, for either a site or personal space, there are various things you can change about that space. For example, you can change the space's name, description, look and feel, and permissions. You can even cr eate templates to speed up page creation in the space.

The space's sidebar is another area you can customise to suit your needs, by adding shortcut links or changing the space logo.

On this page:

- Administer a space
- Administer a space that uses the Document ation theme

Related pages:

- Configure the Sidebar
- Spaces
- Space Permission s Overview
- Assign Space Permission s

Administer a space

To view the space tools page:

1. Go to the space and choose **Space tools** > **Permissions** from the bottom of the sidebar
2. The 'Space Permissions' page displays.

The administration options are divided into the following categories. Some options may not be applicable for personal spaces.

- **Overview**
  - Space Details - see Edit Space Details
  - Space Categories - see Use Labels to Categorise Spaces
  - Delete Space - see Delete a Space
- **Permissions**
  - Permissions - see Assign Space Permissions
  - Restricted Pages - see Page Restrictions
- **Content Tools**
  - Templates - see Page Templates
  - Orphaned Pages - see Orphaned Pages
  - Undefined Pages - see Undefined Page Links
  - Attachments - see Manage Files
  - Trash - see Delete or Restore a Page.
  - Export - see Export Content to Word, PDF, HTML and XML
  - RSS Feeds - see Subscribe to pre-specified RSS feeds
  - Import - see Import a Text File. *Not available for personal spaces.*
- **Look and Feel**
  - Themes – see Apply a Theme to a Space
  - Colour Scheme – see Edit a Space's Colour Scheme
  - Layout - see Customise Space Layouts
  - Stylesheet - see Styling Confluence with CSS
  - PDF Layout – see Customise Exports to PDF
  - PDF Stylesheet – see Customise Exports to PDF

- **Integrations**
  - Application Links – see Linking to Another Application
  - Mail Accounts – see Mail Archives. *Not available for personal spaces.*
  - Mailbox Import – see Import Mail from an mbox. *Not available for personal spaces.*
  - Mail – see Mail Archives. *Not available for personal spaces.*

You can also customise the sidebar, including changing the space logo and adding shortcuts to other spaces - see Configure the Sidebar for more information.

Administer a space that uses the Documentation theme

In this theme, the standard Confluence sidebar is replaced by a page navigation sidebar. As a result, the layout of the space administration options are slightly different.

To view the space admin page:

1. Go to the space you wish to manage
2. Choose **Browse** > **Space Admin** from the header

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

The 'Space Details' page displays. For spaces using the Documentation theme, the administration options are divided into the following categories:

- **General**
  - Space Details – see Edit Space Details.
  - Space Categories – see Use Labels to Categorise Spaces.
  - Templates – see Page Templates.
  - Delete Space – see Delete a Space.
  - Trash – see Delete or Restore a Page.
- **Security**
  - Permissions – see Assign Space Permissions.
  - Restricted Pages – see Page Restrictions.
  - Application Links – see Linking to Another Application.
- **Mail**
  - Mail Accounts – see Mail Archives
  - Mailbox Import – see Import Mail from an mbox
- **Look and Feel**
  - Themes – see Apply a Theme to a Space.
  - Colour Scheme – see Edit a Space's Colour Scheme.
  - PDF Layout – see Customise Exports to PDF.
  - PDF Stylesheet – see Customise Exports to PDF.
  - Change Space Logo – see Change the Space Logo.

  Import
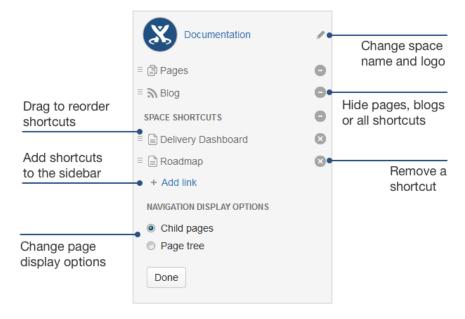
  - Import Pages from Disk – see Import a Text File.

Configure the Sidebar

If you're an admin for a space, you can make changes to the space's sidebar like changing the space's name, logo, or the way the page hierarchy is displayed. You can also add shortcut links to the sidebar to help you and others navigate to important content quickly.

ℹ The default and documentation themes in Confluence both feature a sidebar. If you're using a third party theme, it may not feature a sidebar.

On this page we'll cover how to customise the sidebar in the

Drag to reorder shortcuts

Add shortcuts to the sidebar

Change page display options

Change space name and logo

Hide pages, blogs or all shortcuts

Remove a shortcut

default theme. If you want
information on the sidebar in the Documentation theme, see Configure the Documentation Theme.

Configure a space's sidebar

To start configuring the sidebar, choose **Space tools** at the bottom of the sidebar, then choose **Configure sidebar**.

From there you can:

- Change the space name and/or space logo:
  1. Choose the edit icon

     next to the space name
  2. Type in a space **Name** and/or choose **Upload an image** to change the space logo
  3. Choose **Save**
- Configure the **Pages** and **Blog** links:
  - Choose the icons to **hide** or **show** the 'Pages' or 'Blog' link.
    For example, you may want your space to be used primarily as a blog (hide the 'Pages' link), or you may not need a blog in your space (hide the 'Blog' link).
  - Drag the links to a different sequence within the section (you can't move a link from one section to another).
    ℹ Add-ons may add other links in this section of the sidebar. For example, the Team Calendars add-on may put a link in this location.
- Add or remove the **shortcut links**:
  - Choose **Add link** to add a shortcut link to the sidebar. This can be a link to an important page for your team, or to an external site, for example.
  - Choose the icon to **remove** a shortcut link.
  - Drag the links to a different sequence within the section. Note that you cannot move a link from one section to another.
  - Choose the **hide** or **show** icon beside the 'Space Shortcuts' heading to show or hide all shortcuts on the sidebar.
- Change the **navigation display options**:
  - Choose **Child pages** to see the current page and its children in the sidebar.
  - Choose **Page tree** to see the page tree for the entire space, expanded to the current page.

Edit Space Details

Space details are the name, description, home page, and archived status of a space, which you can edit if you're an administrator of the space.

You can also view the space key and the space creator's name, but you can't edit them.

Related pages:

- Configure the Sidebar
- Spaces
- Archive a Space
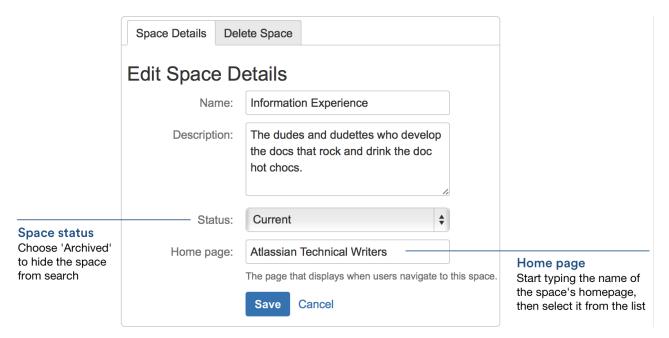
To edit a space's details:

1. Go to the space and choose **Space tools** > **Overview** from the bottom of the sidebar
   If your space uses the Documentation theme, choose **Browse** > **Space Admin** in the header.
2. Choose **Edit Space Details**
3. Update any of the following:

   - **Name**
   - **Description**
   - **Status** – Set the status to 'Archived' if you want to archive the space.
   - **Home page** – Start typing the name of a page in the space, then select it to set it as the new home page. This is the page you'll see when you navigate to the space. If you set this field to blank (no selection) the default home page will be the 'Pages' page.

ℹ The space fields don't accept wiki markup; if you enter wiki markup in these fields, it'll be displayed as plain text. You can also change the name of a space via the sidebar.

Space Details | Delete Space

Edit Space Details

Name: Information Experience

Description: The dudes and dudettes who develop the docs that rock and drink the doc hot chocs.

Status: Current

Home page: Atlassian Technical Writers

The page that displays when users navigate to this space.

Save Cancel

Space status
Choose 'Archived' to hide the space from search

Home page
Start typing the name of the space's homepage, then select it from the list

Use Labels to Categorise Spaces

A **space category** is a label that you can apply to a space for the purpose of grouping your spaces in the the space directory, and in the recent activity area of the dashboard.

For example, if you have a space for each of your projects, you can add a category of 'project' to each of those spaces. It'll mean they're easy to find if your Confluence site has a lot of spaces, which are a mix of project, team, personal, and other spaces.

Add as many space categories as you think you need; it's just like adding labels to a page or blog post.

You need to be an administrator of the space to add categories to it.

On this page:

- Categorise a space
- Categories in the space directory
- Remove a space from a category

Related pages:

- Add, Remove and Search for Labels
- Spaces

Categorise a space

1. Go to the space and choose **Space tools** > **Overview** from the bottom of the sidebar
2. Choose **Edit** next to **Space Categories**
3. Under **Space Categories**, enter your category name and choose **Add**
Alternatively, choose a category in the list of '**Suggested Space Categories**
4. Choose **Done**

˅ Is your space using the Documentation theme?

If your space is using the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **Space Labels** in the space administration options
3. Under **Space Categories**, enter your category name and choose **Add**
Alternatively, choose a category in the list of **Suggested Space Categories**.
4. Choose **Done**

Categories in the space directory

Once you've categorised some spaces, you can view your categories by choosing **Spaces** > **Space directory** in the Confluence header, then choosing one of your categories from the list on the left.

You can also view spaces by category in the **Space Categories** tab of the recent activity section on the dashboard, or embed the Spaces List Macro on any page and allow filtering by category.

Remove a space from a category

To remove a space from a category, follow the above steps to add a space category but, instead of adding a new category, choose the **x** next to the space category you want to remove. If you remove all spaces from a category, the category will also be removed.

Screenshot: Space categories

Space Details	Remove Space

Space Categories

Space Categories will appear inside the categories tab on your dashboard and as tabs in the space directory. They can be used to group together related spaces.

Labels confluence ×   development ×

[   ]   Add

Delete a Space

Deleting a space permanently removes the space and all of its contents. To delete a space you must be an administrator of the space.

⚠️ Deleting a space is permanent. If you're unsure, always create an XML backup of the space before proceeding. You can also choose to backup the attachments if you need to.

Once you've deleted the space, there is no way to restore it unless you've made an XML space backup.

See Restoring a Space.

Related pages:
- Archive a Space
- Export Content to Word, PDF, HTML and XML

To delete a space:

1. Go to the space and choose **Space tools** > **Overview** from the bottom of the sidebar
2. Choose **Delete Space**
3. Choose **OK**

If your space is using the Documentation theme:

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

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

2. Choose **Delete Space** in the space administration options
3. Choose **OK**

Space Tools

Overview   Permissions   Content Tools   Look and Fee

Space Details	Delete Space

✏️ Edit Space Details

Name	Project Space
Key	PS
Home page	🔗 Project Space Home
Created by	Gil Admin
	(2015-06-04)
Categories	(None) [Edit]
Description	

Delete space
Choose the Delete space tab to permanently delete a space

Members of the `confluence-administrators` group can also delete spaces, including personal spaces.

Archive a Space

You can archive a space, so that its content is less visible but it's still available in your Confluence site. You need to be an administrator of the space to archive it.

If you want the space to be fully visible again, you can change its status from archived to current at any time.

Archive a space

1. Go to the space and choose **Space tools** > **Overview** from the bottom of the sidebar
2. Choose **Edit Space Details**
3. Select **Archived** in the **Status** dropdown menu
4. Choose **Save**

If your space uses the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
   The 'Edit Space Details' screen appears.
2. Select **Archived** from the **Status** dropdown menu
3. Choose **Save**

The effect of archiving a space

If a space is archived:

- The pages and other content don't appear in the Confluence quick navigation aid, which drops down when you enter text in the search box.
- By default, the pages and other content don't appear in the Confluence search results. If a Confluence site contains any archived spaces, the search screen will provide an option to include them in the search results.
- The space and its pages don't appear on the dashboard.
- Updates to the space's content don't appear in activity streams, such as the recent updates section of the dashboard.
- The space doesn't appear in space-selection dropdown menus. Similarly, its pages and other content don't appear in any dropdown menus in the Confluence user interface.
- In the space directory, the archived space won't appear in the general space lists. Archived spaces appear in the **archived spaces** tab. They'll also appear in the category tabs, as determined by their labels.

These functions remain available for archived spaces:

- You can view the content as usual, by following a link or typing in a URL belonging to the archived space.
- You can edit the content as usual, as determined by the space permissions.

- RSS feeds, watches and notifications remain active.

ℹ Archiving a space has no effect on search results of external search engines. For example, the space will still appear in Google search results.

View Space Activity

Space activity information is **disabled by default**, and the 'Activity' tab won't be visible unless the Confluence Usage Stats plugin 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:
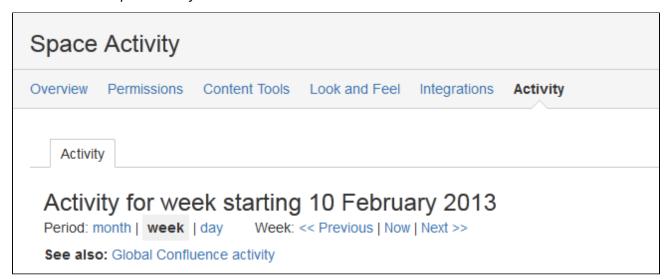
1. Go to the space and choose **Space Tools** at the bottom of the sidebar
2. Choose **Activity**

Related pages:

- Page History and Page Comparison Views
- Watch Pages, Spaces and Blogs
- How do I get more statistics from Confluence?

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

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)	Most active content (Edits)	Most active contributors (Edits)
1. Angry Nerds Home (6)	1. Ideas for Angry Nerds 2 - revenge of the... (2)	1. Josh User (5)
2. Which nerd are you? (3)	2. Why I loved Angry Nerds (1)	2. Ewan User (3)
3. Step 1 - download the game (2)	3. Which nerd are you? (1)	3. Rach Admin (1)
4. Why I loved Angry Nerds (2)	4. Step 2 - Play (1)	
5. Step 2 - Play (2)	5. Status update (1)	
6. Ideas for Angry Nerds 2 - revenge of the... (2)	6. Games as a communication medium (1)	
7. Status update (1)		
8. Games as a communication medium (1)		

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 in Confluence Cloud, and is **disable d by default**. System administrators can enable this plugin (go to **add-ons**, select **System add-ons** a nd search for 'Confluence Usage Stats').

- The plugin collects data only when it's activated.
- If you're using Confluence Data Center, 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.
- The Activity tab does not appear in spaces that use the Documentation theme.

View Recently Updated Content

The 'Recently Updated' view is a useful way of keeping track of the changes being made in a space. It displays links to the most recently added or modified content within the space, including pages, blog posts, mail messages and comments.

If you want to display a list of recently updated content on a page, try the Re cently Updated Macro.

To view the recently updated content in a space, go to the space and choose **Pages** in the sidebar. If the space is using the Documentation theme, choose **Browse** > **Pages** in the header, then choose **Recently Updated**.

Related pages:

- Recently Updated Macro
- View Space Activity
- Favourite Pages

You'll see a list of the most recently added or modified content in the space. Choose any of the links to open the corresponding content.

Customise Space Layouts

You can modify Confluence's look and feel by editing the layout files. This page tells you how to customise the layout files for a space. You will require need space administrator permissions for that space.

If you're a Confluence system administrator, you can also customise the layout of your entire Confluence site. For more information, see Customising Site and Space Layouts. Site layout customisations modify the default layout of all spaces in the Confluence site.

Any space layout customisations will override the equivalent site customisations.

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

> **Related pages:**
> - Change the Look and Feel of a Space
> - Apply a Theme to a Space
> - Styling Confluence with CSS

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.

To edit the layout of Confluence, you will need to modify these decorator files. A decorator file is a .vmd file and is written in a very simple programming language called **Velocity**. Learn more about Velocity. Once you become familiar with Velocity, you can edit the decorator files to personalise the appearance of Confluence.

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

- **Site layouts**: These are used to define the controls that surround each page in the site. For example, if you want to make changes to the header and the footer, you will need to modify these layouts.

- **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 they allow you to alter the way the surrounding comments or attachments are shown.

- **Export layouts**: These control the appearance of spaces and pages when they are exported to HTML. If you are using Confluence to generate a static website, for example, you will need to modify these layouts.

Learn more about using decorators.

To edit a decorator file:

1. Go to the space and choose **Space tools** > **Look and Feel** from the bottom of the sidebar
2. Choose **Layout** (**Layout** is displayed only if you are a Confluence system administrator.)
   You'll see a list of the layouts for the space
3. Click **Create Custom** to edit the default vmd file
   This will open up the vmd file in edit mode. If you only want to view the vmd file, click **View Default**.
4. Make any changes and click **Update**

> You can't update the layout templates if your space is using the Documentation theme.

Screenshot : Edit Layouts Example

Content Layouts

Content layouts control the appearance of content such as pages and blog posts: they don't change the way the pages themselves are displayed, but they allow you to alter the way the surrounding comments or attachments are shown.

Decorator	Operations
Space Layout Controls how a space is displayed.	< default > · Create custom · View Default
Page Layout Controls how a page is displayed.	< default > · Create custom · View Default

Export Content to Word, PDF, HTML and XML

You can export all or part of a Confluence space to various formats, including Microsoft Word, HTML, PDF and XML.

To use the space export functionality, you need the 'Export Space' permission. See the guide to space permissions.

Export single pages to Word or PDF

If you need to send content to people who don't have access to Confluence, you can export a single page as a Word document or a PDF. You can also export a single blog post to PDF.

If you've got permission to view the page in Confluence, you'll be able to export it in this way; go to the page and choose either **Tools** > **Export to Word** or **Tools** > **Export to PDF**.

On this page:

- Export single pages to Word or PDF
- Export to HTML, XML, or PDF
- Customising the appearance of PDF exports

Related pages:

- Customise Exports to PDF

When you export a single page to PDF, the PDF stylesheet customisations are applied, but any PDF layout customisations aren't. To make your PDF layout customisations apply to a single page exported to PDF, you'll need to use the 'multiple page' method described below to export the single page. See Customise Exports to PDF.

Export to HTML, XML, or PDF

If you want to export a space – or selected pages in a space – to HTML, XML, or PDF, Confluence can create a zipped archive of the HTML or XML files, or a single, downloadable PDF file.

The HTML export is useful if you want convert your space into a static website, or you can use the XML export option if you need to import the space into another Confluence site, or use the data from the space in another application. You might use the PDF option if you're producing a user manual from your technical documentation in Confluence.

To export pages to HTML, XML, or PDF:

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar
2. Choose **Export**
3. Select either **HTML**, **XML**, or **PDF**, then choose **Next**
4. Select either a normal or custom export for HTML or PDF, or a Full or Custom XML export:
   - **Normal Export (HTML and PDF)** – to produce an HTML file containing all the pages that you have permission to view.
   - **Full Export (XML)** – to produce an XML file containing all the pages in the space, including those that you do not have permission to view.
   - **Custom Export** – if you want to export selected pages only, or if you want to exclude comments from the export.
5. Choose **Export**

ⓘ Blog posts aren't included when exporting to PDF using this method, and comments are never included when exporting to PDF.

When the export process has finished, you can download the zipped archive or PDF.

If your space uses the Documentation theme, choose **Browse** > **Space Operations**, then choose either **HTML, XML,** or **PDF Export** from the left menu, and follow the steps above to export the Space.

Customising the appearance of PDF exports

You can add a title page, table of contents, and customised headers and footers to the PDF output. For more

advanced customisations, you can apply Cascading Style Sheet (CSS) modifications. These customisations are specific to each space, and you need the 'Space Administrator' permission to apply these customisations. For more information, see Customise Exports to PDF.

Notes on PDF exporting

- To export a PDF containing international text, see Create a PDF in Another Language
- Confluence's PDF export feature is designed to handle a wide variety of content, but on rare occasions the PDF Export process may fail due to an unrecognised customisation. If that happens, the PDF export screen will indicate the title of the page in which the problem occurred, to help you diagnose the cause of the failure.

Notes on HTML exporting

- In the zip file, page attachments are placed in individual folders with names in the following format:
  `...\download\attachments\xxxxxx`
  where '`xxxxxx`' is the page ID of the page containing the attachments.
- Blog posts aren't included in the HTML export. See the feature request:

  CONF-14684 - Export news to HTML
  OPEN

- To customise the HTML output, you'll need to modify the file `confluence-x.y.z-jar/com/atlassian/confluence/pages/Page.htmlexport.vm`. To learn how to repackage this file, see How to Edit Files in Confluence JAR Files

Notes on XML exporting

- See Restoring a Space for notes on restrictions when **importing** a space.
- If you're doing the export for backup purposes, consider another means of backup. See Production Backup Strategy.
- If you are running Confluence behind Apache HTTP Server and are facing timeout errors, please consider creating the export directly from Tomcat, instead of going through Apache. This will speed up the process and prevent timeouts

Customise Exports to PDF

Confluence administrators and space administrators can customise the PDF exports for individual spaces.

Please note:

- PDF customisations are unique to each space.
- PDF customisation only applies to space exports (not single page exports via **Tools** > **Export to PDF**)
- Confluence's PDF customisations use a combination of HTML and CSS (Cascading Style Sheets).
- To achieve a particular requirement in the exported PDF file, you make changes in one or both of the following:
  - The **PDF Layout**, where HTML is used to define the structure of the exported content, including features such as the title page, headers and footers.
  - The **PDF Stylesheet**, where CSS is used to define the style of elements in

the exported content, such as page size and margins, font, colour, and list numbering.

For further customisations, see the Advanced PDF Export Customisations page.

Customising the PDF Layout

You can add your own HTML to customise the title page, page headers and page footers in the PDF output.

Setting a global PDF layout

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **PDF Layout**. You can enter HTML and CSS that will be rendered on the PDF title page, as well as page headers and footers.

Setting the PDF layout at space level

1. Go to the space and choose **Space tools** > **Look and Feel** from the bottom of the sidebar You'll need Space Admin permissions to do this.
2. Choose **PDF Layout**.
3. Choose **Edit**.

If your space uses the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
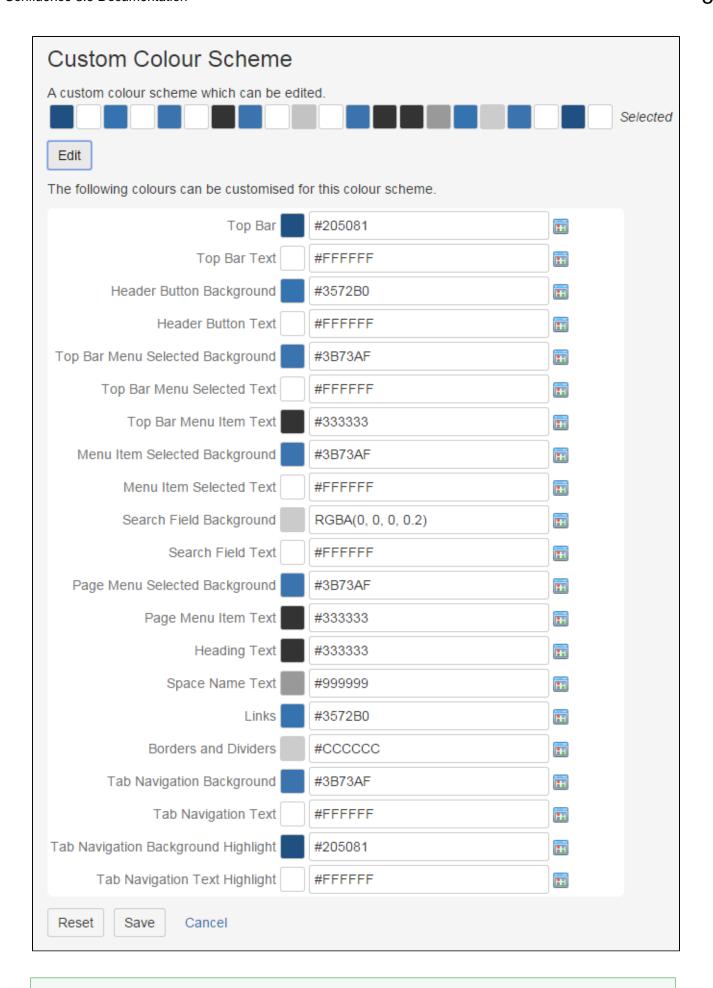2. Choose **PDF Layout** in space administration options.
3. Choose **Edit**.

You can customise the HTML that is generated for the following areas when the space is exported to PDF:

- **PDF Space Export Title Page** – title page.
- **PDF Space Export Header** – page headers.
- **PDF Space Export Footer** – page footers.

Enter your customisations into each text box as required. The PDF layout section accepts HTML code. You can include inline CSS in the HTML too.

Customising the PDF Stylesheet

You can adjust the appearance of the PDF pages by customising the CSS in the PDF Stylesheet screen.

To get started, download the default CSS rules for the PDF stylesheet - confluencedefaultpdf.css.

Any rule defined in this file can be customised and added to the PDF Export Stylesheet section. Your customisations override any default CSS rule. If no customisations are defined, the default CSS rules will be applied.

By default, the export does not include a title page, headers or footers. You can define these in the PDF layout.

To customise the PDF Stylesheet:

Setting a global PDF Stylesheet

1. Choose the **cog icon**
   ⚙ ▾
   , then choose **General Configuration** under Confluence Administration
2. Choose **Look and Feel** > **PDF Stylesheet**. The following screen allows you to enter and save CSS code that will render content on each page.

Setting a space PDF stylesheet

1. Go to the space and choose **Space tools** > **Look and Feel** from the bottom of the sidebar You'll need Space Admin permissions to do this.
2. Choose **PDF Stylesheet**.
3. Choose **Edit**.
4. Enter your customisations.

If your space uses the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **PDF Stylesheet** in the space administration options.
3. Choose **Edit**.
4. Enter your customisations.

The PDF Export Stylesheet field accepts Cascading Style Sheets (CSS) rules. The 'PDF Export Stylesheet' page shows the current (e.g. customised) contents of your PDF stylesheet.

Example Customisations

This section provides examples of typical customisations that you can add. See also Advanced PDF Customisations.

Page Size

The default page size is based on the location of your Confluence server. For example, if this server is located in the US then the default paper size of your PDF export will be US Letter (8.5 inches wide by 11 inches long). If the server is located in Australia, the default paper size will be A4 (210 mm wide by 297 mm long). More information about paper sizes can be found on Wikipedia.

To modify the page size to A4, edit the PDF Stylesheet to add a `size` property to the `CSS@page` rule, like this:

CSS - PDF Stylesheet

```
@page
{
/*The A4 paper size is 210 mm wide by 297 mm long*/
size: 210mm 297mm;
}
```

Page Orientation: Landscape or Portrait

To change the page orientation of your PDF document, simply reverse the order of the values declared in the `@page` rule's `size` property. The first and second values of this property represent the width and height of the page, respectively.

For example, to generate PDF space exports in A4-sized landscape orientation, your `@page` rule might look like this:

CSS - PDF Stylesheet

```
@page
{
/*A4-sized pages in landscape orientation are 297 mm wide by 210 mm long*/
size: 297mm 210mm;
}
```

Page Margins

To set all margins to 15 mm, with a paper size of A4, edit the `CSS @page` rule in the PDF Stylesheet, like this:

CSS - PDF Stylesheet

```
@page
{
size: 210mm 297mm;
margin: 15mm;
}
```

To set the margins independently, edit the `@page` rule as follows:

CSS - PDF Stylesheet

```
@page
{
margin-top: 2.54cm;
margin-bottom: 2.54cm;
margin-left: 1.27cm;
margin-right: 1.27cm;
}
```

To set margins to provide a gutter for binding a printed document, use the `:left` and `:right` pseudo-classes, as follows:

CSS - PDF Stylesheet

```
@page :left
{
margin-left: 4cm;
margin-right: 3cm;
}
@page :right
{
margin-left: 3cm;
margin-right: 4cm;
}
@page :first
{
margin-top: 10cm /* Top margin on first page 10cm */
}
```

Note the use of the `:first` pseudo-class in the example above to define distinct margins for a cover or title page.

Page Breaks

By default, Confluence pages are exported without page breaks, so that shorter pages will appear on the same PDF page.

To make each Confluence page appear on a separate page in the PDF file, add the following rule in the PDF Stylesheet:

CSS - PDF Stylesheet

```
.pagetitle
{
page-break-before: always;
}
```

Title Page

You can add a title page to your PDF document by adding HTML to the Title Page section of the PDF Layout screen. The following example creates the title page and adds a title:

HTML - PDF Layout: Title Page Section

```
<div class="fsTitlePage">
 <img src="/download/attachments/590719/titlepage.png" />
<div class="fsTitle">Planning for Confluence 4.0</div>
</div>
```

Use CSS rules in the PDF Stylesheet to control the appearance of the title page and the title text:

```
CSS - PDF Stylesheet

.fsTitlePage
{
margin-left: auto;
margin-top: 50mm;
margin-right: auto;
page-break-after:always
}

.fsTitle
{
font-size: 42px;
font-weight: bold;
margin: 72px 0 4px 0;
text-align:center;
}
```

Adding an Image to the Title Page

In the example above, the title page includes an image called 'titlepage.png', centred in the middle of the page. The "590719" in the `src` attribute is the ID number of the page to which the image is attached.

Follow these instructions to include an image on your title page:

1. Attach the image to a Confluence page.
2. View the list of attachments on that page, then right-click the image and copy its location.
3. Paste the link into the appropriate `src=""` attribute in your PDF Stylesheet, as shown above.
4. Edit the image URL so that it is relative, by removing the first part of the URL before `/download/....`

Table of Contents

By default, a table of contents will be generated after the title page, or at the beginning of the PDF document if no title page is defined in the PDF Layout. To see the default CSS rules applied to the table of contents, download the default CSS rules ([confluencedefaultpdf.css](#)) and examine the specific rules with `toc` in their name.

To make changes to the appearance of the table of contents, define CSS rules in the PDF Stylesheet.

Disabling the Table of Contents

To prevent the table of contents from being generated in your PDF document, add the `div.toc-macro` rule to the PDF Stylesheet and set its `display` property to none:

```
CSS - PDF Stylesheet

div.toc-macro
{
display: none;
}
```

Changing the Leader Character in the Table of Contents

The leader character is used to link a heading in the table of contents with its page number. By default, the leader character is the '.' (dot) character. Leader values of `dotted`, `solid` and `space` are allowed. You can also use a string, for example `leader(". . . ")`.

To change the leader character to a solid line, modify the `leader()` value on the `content` property of the CSS rule as follows:

CSS - PDF Stylesheet

```
span.toclead:before
{
content: leader(solid);
}
```

Headers and Footers

You can add headers and footers to your PDF pages using the 'Header' and 'Footer' sections of the PDF Layout screen. By default, headers and footers only apply to a space export and not to exports of single pages (however, see Advanced PDF Export Customisations). The following example adds a simple copyright notice.

HTML - PDF Layout: Footer Section

```
Copyright © 2013, Atlassian Pty Ltd.
```

Page Numbering

To add page numbering, you need to combine customised HTML in the PDF Layout with customised CSS in the PDF Stylesheet.

1. PDF Layout HTML: In the Footer section (or the Header section), use an empty span element with a unique ID, for example `pageNum`, to act as a place holder for the page number.

   HTML - PDF Layout: Footer Section

   ```
   <span id="pageNum"/>
   ```

2. PDF Stylesheet CSS: Create the following CSS selector rule for the empty span:

   CSS - PDF Stylesheet

   ```
   #pageNum:before
   {
   content: counter(page);
   }
   ```

Analysing the above CSS selector rule in more detail:

- The `#pageNum` rule selects the HTML element with the specified ID of "pageNum", which is the `span` element we created for the header or footer.
- The `:before` part of the selector is a pseudo class that allows the insertion of content before the `span` element is processed.
- The `counter(page)` is a function that returns the current page number as its content.
- The `content` property tells the CSS processor that dynamic content (that is, an incrementing page number) is to be inserted at the span tag.

Wrapping Long Words

In order to break long words or words that are not separated by whitespace, add a selector to the PDF stylesheet containing the `word-wrap` property with a value of `break-word`:

```
CSS - PDF Stylesheet
div
{
word-wrap: break-word;
}
```

General Formatting

You can use the stylesheet to customise the output of just about anything else that will influence the look and feel of the final document. This includes fonts, tables, line spacing, macros, etc. The export engine works directly from the HTML output produced by Confluence. Therefore, the first step in customising something is to find a selector for the HTML element produced by Confluence or the Confluence macro. Then add a CSS rule to the PDF stylesheet. Your customisation will appear in the PDF export.

Notes

- **Demotion of heading elements: `h1, h2, and so on.`** Due to the hierarchical manner in which a space is exported, Confluence will modify the heading elements to generate a uniform appearance for the entire space export. This means that headings will be demoted. This will affect the application of custom PDF Stylesheets. It is possible to calculate the amount by which a heading will be demoted in order to have the correct CSS styling applied. A heading will be demoted by the value of its depth in the export tree. A page at the first level will be demoted by 1 (all `<h1>` elements will become `<h2>` elements, and so on). A page at the second level will be demoted by 2, and so on.

- **Atlassian support for PDF customisation is limited.** We support the mechanism for customising the PDF layout with HTML and CSS, and we will help if the mechanism is broken or does not work as we say it should in our published examples. But, since custom HTML and CSS offer potentially limitless possibilities, Atlassian will not support issues that are caused by or related to PDF customisations.

Advanced PDF Export Customisations

This page provides information about 'advanced' PDF export customisations. These expand upon the regular customisations described in Customise Exports to PDF.

⚠ The information below is for advanced users. Be aware that the advanced customisations described below require knowledge of certain parts of Confluence, and of CSS and HTML. Customisations are not supported by Atlassian, so our support engineers won't be able to help you with these modifications.

On this page:

- Header and Footer
  - Adding Headers and Footers to Single Page Exports
  - Adding Images to Headers and Footers
- Incorporating Other Fonts
- Adding a Dynamic Title to the Title Page
- Hiding Text from the PDF Output
- Indexing
- Notes

Related pages:

- Customise Exports to PDF
- Create a PDF in Another Language

Header and Footer

Adding Headers and Footers to Single Page Exports

Single page exports don't support adding HTML headers and footers via the PDF Layout page, but you can

use CSS rules in the PDF Stylesheet page (**Space tools** > **Look and Feel** > **PDF Stylesheet**) to produce headers and/or footers for a single page export.

For custom headers, define any of the following rules within your `@page` rule: `@top-left`, `@top-center`, a nd `@top-right`. These rules allow you to define the `content` of the left-hand side, centre and right-hand side of your page's header area, respectively.

For custom footers, define `@bottom-left`, `@bottom-center` and `@bottom-right` rules within your `@pag e` rule.

For example, the following rules add a document title at the centre of the header and a page number at the centre of the footer:

CSS - PDF Stylesheet

```
@page
{
    @top-center
    {
        content: "Document Title Goes Here"; /* This is the content that will
appear in the header */
        font-family: ConfluenceInstalledFont, Helvetica, Arial, sans-serif;
        font-size: 8pt;
    }
    @bottom-center
    {
        content: "Page " counter(page); /* This is the content that will appear
in the footer */
        font-family: ConfluenceInstalledFont, Helvetica, Arial, sans-serif;
        font-size: 8pt;
    }
    /* Any other page-specific rules */
}
```

Notes:

- The `font-family` and `font-size` properties ensure that the header and footer text is rendered in the same default font style used for the body text, based on the default CSS rules.
- It is not possible to use this method to insert images (stored as attachments within your Confluence instance) into the headers and footers of single page exports.

Adding Images to Headers and Footers

To insert an image into the header or footer, add HTML to the Header or Footer section of the PDF Layout screen.

The following example uses an HTML `img` element with `src` attribute to add an image to the left of the header. The `src` attribute refers to an image attached to a Confluence page. The image element is usually placed within a `div` element container.

HTML - PDF Layout: Header Section

```
<div style="margin-top: 10.0mm;">
        <img
src="https://confluence.atlassian.com/download/attachments/12346/header-image.png
" />
    </div>
```

In the example above, the header includes an image called 'header-image.png'. The "12346" in the `src` attrib

ute is the ID number of the page to which the image is attached.

Follow these instructions to include an image on your page:

1. Attach the image to a Confluence page.
2. View the list of attachments on that page, then right-click the image and copy its location.
3. Paste the link into the appropriate `src=""` attribute in your PDF Stylesheet, as shown above.
4. Edit the image URL so that it is relative, by removing the first part of the URL before `/download/....`

Notes:

- This example uses an inline CSS property `margin-top` in the `style` attribute to force the image away from the top of the page by 10mm. This comes in handy when your header image is large enough to touch or spill over the top of the page.
- Likewise, for footers, you can use the `margin-bottom:XXmm` property to force an image away from the bottom of the page by 'XX' mm.
- Very large images can spill over into the body of a page or alter the position of text or other elements used within a header or footer. In such situations, it is recommended that you reduce the size of the image and then attach it to your Confluence page again. If you prefer to keep the image size and want to move the content lower instead, you can do so by configuring the `margin-top` properties in the `@page` CSS rule.
- By default, a header or footer image is aligned to the left-hand side of the page. However, you can align this image to the centre or right-hand side of a page by adding either the `text-align:center` or `text-align:right` properties to your `style` attribute. For example, to align the header image to the right-hand side of the page, your `style` attribute would look similar to this: `style="margin-top:10mm; text-align:right"`.

Incorporating Other Fonts

By default, Confluence provides Times New Roman, Helvetica or Courier fonts for use in PDF exports. You can use your own fonts for PDF exports by declaring them in a `@font-face` CSS rule in your PDF Stylesheet.

The following CSS rule example shows how to declare the Consolas font and apply it to some elements for your PDF export:

CSS - PDF Stylesheet

```
@font-face { src: url(file:///usr/share/fonts/Consolas.ttf); -fs-pdf-font-embed:
embed; } .code pre, .preformatted pre, tt, kbd, code, samp { font-family:
Consolas, monospace; font-size: 9pt; }
```

The font path specified in the CSS must be the path to the font on the Confluence server.

Adding a Dynamic Title to the Title Page

When you export an arbitrary set of pages from Confluence, you may like to have a corresponding title added to the cover (or title) page automatically. This can be done (in a somewhat irregular way) by using the top level item from the default table of contents as the title. This method relies on having the exported pages structured as sub-pages of the top-level page. In other words, the pages to be exported should consist of a page (at the top-level) and all of its child pages. The result is that the title that appears on the cover page changes depending on the top-level page that is used for the export.

The CSS below moves, and styles, the top-level TOC item for use as the title on the cover page, and turns off the leader and page number normally associated with this item in the TOC.

CSS - PDF Stylesheet

```
.fsTitlePage { position:relative; left:0px; } /* Turn off the default section
numbering for this TOC item */ .toclvl0:before { content: &quot; &quot;;
counter-reset: chapter 0; } /* Hide the default page numbering for this TOC item
*/ .toclvl0 .tocnum { display: none; } /* Move and style this TOC item */
.toclvl0 { position:absolute; top:250px; font-size: 42px; font-weight: bold;
margin: 72px 0 4px 0; text-align:center; }
```

Hiding Text from the PDF Output

This section describes a way to hide text from your PDF export. In other words, you can have text on the Confluence page that will not appear in the PDF export.

There are three steps:

1. Follow the instructions to define the NoPrint user macro.
2. Use the NoPrint macro to mark some text on a Confluence page.
3. Add the following CSS to your PDF stylesheet to make the PDF export recognise the NoPrint macro:

CSS - PDF Stylesheet

```
.noprint { display: none ; }
```

Indexing

To obtain an index at the end of the exported PDF file, consider using the Scroll Wiki PDF Exporter plugin that is produced by K15t Software GmbH.

Notes

If styling is not working as expected, it is useful to look at the intermediary HTML source to which the CSS is applied. This intermediary HTML is created whenever you create an HTML export that contains **multiple pages**, and is stored in the `temp` directory in Confluence's home directory. For example:

```
/temp/htmlexport-20110308-154047-1/export-intermediate-154047-2.html
```

Create a PDF in Another Language

To export a Confluence page written in a language other than English, you'll need the necessary font for that language.

Related pages:

- Export Content to Word, PDF, HTML and XML

Upload a Font File to Confluence

1. Find the appropriate font file:
   - **Windows users:**All font files in Windows are stored in a directory called:

     ```
     C:\WINDOWS\Fonts
     ```

   - **Unix users:**All font files in Unix are stored in:

```
/usr/share/fonts
```

- Microsoft True Type core fonts such as Verdana can be downloaded from this page: http://coref onts.sourceforge.net/
2. Copy the font file into a temporary folder, for example a folder on your desktop.
3. Choose the **cog icon**

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

   then choose **PDF Export Language Support**.
4. Upload the file you copied in step 2.
5. Choose **Install**.

Notes

- The only font files supported are *true type fonts* and *true type collections*. The accepted file extensions are *.ttf and *.ttc.
- Confluence can only store **one** font file at any one time. Please create a collection to install more than one *.ttf files.
- We recommend that you use Unicode font Verdana for correct character encoding and exporting to PDF.
- If the font file size is bigger than your current attachment size limit, you will not be able to upload it. Please increase the attachment size limit temporarily and re-upload again. An improvement of the error messaging is tracked at **CONF-24706** **RESOLVED**
- To make use of an installed font in your PDF Export style sheet (CSS) refer to it by the font-family ConfluenceInstalledFont.

Pages and Blogs

When you create any content in Confluence – whether it's taking down some quick notes from a meeting, writing a requirements page, or letting your teammates know about the company's latest marketing push – you'll either be creating it as a page or a blog post. You may be shocked to know that this page you're reading now, is, in fact, a Confluence page!

Pages and blog posts are the two things that allow you to capture and share information in Confluence, and which one you use depends on what you want to do with your content. If you want the information to last, and possibly evolve over time, then you likely want to create a page; if the information is specific mainly to the current time-frame, and isn't going to change over time, then you'll most likely want to create it as a blog post. These aren't hard-and-fast rules; they're just pointers to give you a place to start.

Related pages:

- Create and Edit Pages
- Blog Posts
- The Editor
- Page Templates
- Delete or Restore a Page
- Spaces

Each Confluence space, including your personal space, allows you to create pages in it, and has its own blog where you can create posts. If you're not sure what a space is, or what you can do with spaces, check out our page on Spaces.

Take a look at the below pages to learn more about pages and blog posts in Confluence.

- Create and Edit Pages
- Blog Posts
- The Editor
- Move and Reorder Pages
- Copy a Page
- Delete or Restore a Page
- Add, Remove and Search for Labels
- Drafts
- Page Restrictions
- Links

- Anchors
- Tables
- Add, Assign, and View Tasks
- Autocomplete for links, files, macros and mentions
- Page Layouts, Columns and Sections
- Create Beautiful and Dynamic Pages
- Page Templates
- Blueprints
- Import Content Into Confluence
- Orphaned Pages
- Undefined Page Links
- View Page Information
- Page History and Page Comparison Views
- Confluence Markup

Create and Edit Pages

Create a page

You can create a page from anywhere in Confluence; just choose **Create** in the header and you're ready to go. Pages are the place to capture all your important (and unimportant) information; start with a blank page and use it like a word processor to add rich text, tasks, images, macros and links, or use one of the useful blueprints to capture meeting notes, decisions, and more.
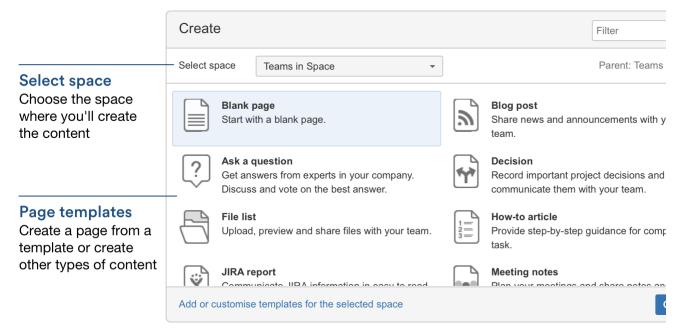
On this page:

- Create a page
- Collaborate or restrict
- Organise and move
- Other page actions

If you want to quickly create a blank page, hit the **Create** button in the header; if you want to create a page from a template, hit the **Create from template** button.

Create blank page   Create ••• Create from template
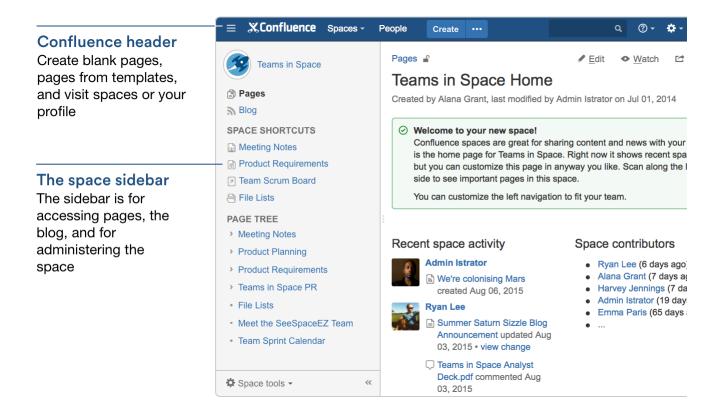
Screenshot: The create dialog

Create   Filter

Select space   Teams in Space   ▼   Parent: Teams

Select space
Choose the space where you'll create the content

Blank page — Start with a blank page.

Blog post — Share news and announcements with y team.

Ask a question — Get answers from experts in your company. Discuss and vote on the best answer.

Decision — Record important project decisions and communicate them with your team.

Page templates
Create a page from a template or create other types of content

File list — Upload, preview and share files with your team.

How-to article — Provide step-by-step guidance for comp task.

JIRA report — Communicate JIRA information in easy to read

Meeting notes — Plan your meetings and share notes an

Add or customise templates for the selected space

Once you decide on a blank page or blueprint, you'll be taken straight into the Confluence editor. The editor

is where you'll name or rename your page, add the content, and format it to look great. When you've added some content, choose **Preview** to take a peek at what your finished page will look like, and choose **Save** wh en you've finished your edits.

After you save you'll see the page in 'view' mode. You can re-enter the editor any time by choosing **Edit** or pressing **E** on your keyboard.

Confluence header

Create blank pages, pages from templates, and visit spaces or your profile

The space sidebar

The sidebar is for accessing pages, the blog, and for administering the space

> Another useful way to create a page is to use the Create from Template Macro. This macro allows you to choose a page template, and adds a button to the page allowing one-click page creation. If you want others to create pages using this template, this is a great option.

Collaborate or restrict

Once you've created a page, you can decide if you want to keep it private, using restrictions, or collaborate on it with others using @mentions, sharing, and comments.

Organise and move

You can also organise pages in a hierarchy, with child and/or parent pages for closely related content. When you navigate to a Confluence page and choose the **Create** button in the header, the page you're creating will by default be a child of the page you're viewing. Have as many child pages and levels in the hierarchy as you need to, and move pages if you want to change their location.

If you want to view all pages in a Confluence space, choose **Pages** in the sidebar, or choose **Browse** > **Pag es** at the top of the screen if you're using the Documentation theme. If the space is using the Default theme, you'll see recent updates to pages and a page tree displaying all pages in the space; if it's using the Documentation theme, you can choose either **Recently Updated**, **Alphabetical**, or **Tree** view of the pages in the space.

> Each time you create a page, you're creating it in a space. Spaces are containers used to contain pages with related content, so you can set them up for each team in your organisation, for projects, a

combination of both, or for any reason you want to group pages together. See Spaces for more information.

Other page actions

- Copy a page
- Delete a page or remove a specific version of a page
- Monitor page updates and other activity through page notifications
- View page history, and manage and compare versions of a page
- Search page content, including attachments
- Export pages to Word, PDF, HTML or XML
- Like a page

⚠ We recommend you don't use special characters in page or attachment names, as the page or attachment may not be found by Confluence search, and may cause some Confluence functions to behave unexpectedly.

If you rename a page, Confluence will automatically update all relative links to the page, except in some macros. Links from external sites will be broken, unless they use the permanent URL. See Working with Links for more information.

Blog Posts

Blog posts are an easy way to share announcements, journal entries, status reports, or any other timely information. Others can join in by commenting on and/or liking your blog post and, if you get enough of either, your post might make it to the popular feed on the dashboard!

Each space in Confluence, including your personal space, has it's own blog. To view a space's blog, go to a space and choose **Blog** in the sidebar; if you're using the Documentation theme, choose **Browse** > **Blog** at the top of the screen. You'll see a list of the latest blog posts, and you can click through to earlier posts via the navigation area in the sidebar.

Create a blog post

You can follow the same process to create a blog post as when you create a Confluence page.

1. Navigate to the space where you want to create your blog post
2. Choose **Create** in the Confluence header and select **Blog post**
3. Add your content and choose **Publish**

You can create blog posts from the Dashboard, but you'll need to make sure you choose the space it's going to appear in in the create dialog.

On this page:

- Create a blog post
- Move a blog post
- Restrict a blog post
- Delete a blog post
- Export a blog post

Related pages:

- Subscribe to RSS Feeds within Confluence
- Blog Posts Macro
- Collaboration
- Export Content to Word, PDF, HTML and XML

Blog posts can be attractive and engaging in the same way a page can be, so go ahead and add images, YouTube clips (preferably of cats), and tables to your post to really grab your audience.

To create a blog post, you need the 'Add Blog' permission. See Space Permissions.

Move a blog post

If you create a blog post in the wrong space, or want to reorganise your spaces, you can move an individual blog post to another space.

To move a blog post, go to the post and choose
•••
 > **Move** and select the new destination space.

You'll need the 'Delete blog' permission in the current space, and 'Add blog' permission in the new (destination) space to do this.

Restrict a blog post

You can restrict a blog post so that it is only available to specific users or groups. Blog post restrictions work in the same way as page restrictions.

To restrict a blog post prior to publishing it, choose the **Unrestricted** button in the footer and apply your restrictions. To restrict a blog post after publishing, choose
•••
 > **Restrictions** and apply your restrictions.

Notes:

- Notifications are sent at the point a blog post is created - removing restrictions **does not** trigger a new notification.
- As a blog post has no parent, restrictions aren't inherited.

Delete a blog post

To delete a blog post, choose
•••
 > **Delete**. Deleting a blog post follows the same rules as deleting a page.

Export a blog post

You can export individual blog posts to PDF. This is useful, for example, if you want to email an internal blog post to people outside your organisation.

See Export Content to Word, PDF, HTML and XML for more information on exporting blog pages to PDF.

The Editor

The Confluence editor is what you'll use to create and edit Confluence pages, blog posts, and comments. You can enter content as you would in a Word document, apply formatting, and embed other content and files on the page.

Note: To edit a page, you need the 'Add Pages' permission for the space. See space permissions. Someone may also apply page restrictions that prevent you from editing the page.

Edit a page or blog post

You'll be taken to the editor whenever you create a new page or blog post, or add a comment. To edit an existing page or blog post, choose **Edit** at the top of a page or press **E** on your keyboard.

Confluence automatically saves drafts of your page as you work. If another user begins editing the same page as you, Confluence will display a message and will try to merge the changes when you save your page. To see changes between different versions of the page, look at the history of the page.

The editor

The editor allows you to enter or change the title of your page; insert content including text, images, and links; and format your content using the toolbar.
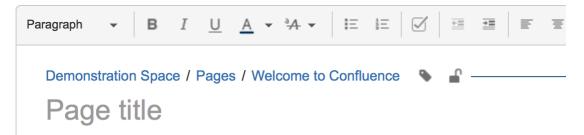
If you're renaming your page, there are some things you should take into account.

Paragraph | B | I | U | A | ᵃA | ☰ | ☰ | ☑ | ☰ | ☰ | ☰ | ☰

Demonstration Space / Pages / Welcome to Confluence

Page title
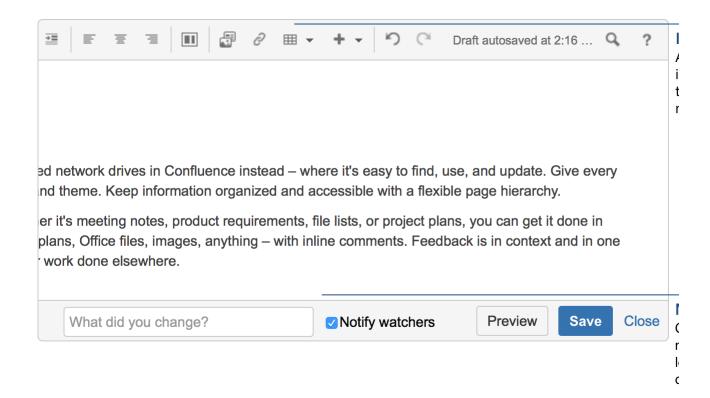
Add your page title and body text

Capture the knowledge that's too often lost in email inboxes and shared network driv team, project, or department its own space with its own permissions and theme. Kee

Create anything with a rich content editor in your web browser. Whether it's meeting Confluence.Target feedback on any Confluence page or file – project plans, Office fil place, so you don't waste time connecting feedback in email with your work done els

...ed network drives in Confluence instead – where it's easy to find, use, and update. Give every ...nd theme. Keep information organized and accessible with a flexible page hierarchy.

...er it's meeting notes, product requirements, file lists, or project plans, you can get it done in ...plans, Office files, images, anything – with inline comments. Feedback is in context and in one ...work done elsewhere.

Editor toolbar

The editor toolbar is where you format your page layout and text, and add links, tables, images, attachments and macros. You can also perform a find and replace, or get help using the editor by choosing the help icon
.

Screenshot: The editor toolbar

The Insert menu

The **Insert** menu is particularly useful. Use the **Insert** menu to include any of the following content types on your page:

- An image
- A link to another Confluence page or external URL, or a link to an attachment or image

- An emoticon or symbol, or a horizontal line
- A macro (choose a specific macro, or **Other Macros**, from the **Insert** menu)

You can also use keyboard shortcuts to insert links, images, and macros. Try out the shortcuts listed below:

- Type **[** (square bracket) to insert a link.
- Type **!** (exclamation mark) to insert an image or other media.
- Type **{** (curly bracket) to insert a macro.

Typing any of the above shortcuts will trigger the autocomplete functionality, prompting you with a list of suggestions to finish off the entry. For more shortcuts, click the help icon

?

on the editor toolbar.

Restrictions, labels, and notifications

When editing a page, you may want to set restrictions on who can view or edit the page, or add labels to the page to make it easily searchable.

Once you're ready to save, you can enter change comments to let others know what you've changed, and, if you like, send an email notification to anyone watching the page.

Things to help you work faster

Auto-formatting

You can type Confluence wiki markup directly into the editor to have Confluence auto-format your text as you type. To learn more, choose help icon

?

in the toolbar, then choose **Editor Autoformatting**.

Autoconvert for pasted links

When you paste certain URLs into Confluence, the editor will analyse what you're pasting and automatically convert it into something that will display well in Confluence. Examples include:

- YouTube videos
- JIRA issue queries
- Google Maps
- Confluence pages, blog posts, comments, user statuses, user profiles.
- Shared screenshot links from Skitch
- And more.

Drag-and-drop for external images and files

You can drag files, like images, multimedia, Office files and PDFs, from your computer and drop them directly into the editor. The contents of the file will be embedded into the page or blog post.

Drag-and-drop within the editor

In the editor panel, you can drag an image or a macro from one location to another on the page. Hover your cursor over the image or the macro placeholder and your cursor changes to a drag-and-drop icon

⬦

. Click the image or macro and drag it to a new location.

Note: For the drag-and-drop of images and macros in the editor, Confluence supports the following browsers: Chrome, Firefox, and Internet Explorer 10 (desktop mode).

Keyboard shortcuts

To view the available keyboard shortcuts, choose the help icon

[?]

in the editor toolbar.

Find and replace text

Click the

🔍
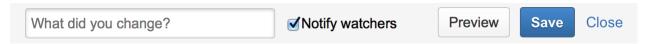
icon on the toolbar, or use the keyboard shortcut **Ctrl+F** (Windows) or **Cmd+F** (Mac OS).

Search matches are highlighted in yellow. You can step through the results one by one, replace the matching text strings one by one, or replace all matching strings at once. Find and replace works only within the current page.

Record change comments and notify watchers

When you finish editing a page, you can add a comment to let others know what you changed. Type a short message in the change comments field in the footer. The comment will be visible in the page history.

If you want to send a notification to people watching the page, select **Notify watchers**. The change comment will be included in the notification email. The **Notify watchers** checkbox remembers your last selection, so if you choose not to notify people, the checkbox will be deselected for you next time.

| What did you change? | ☑ Notify watchers | Preview | **Save** | Close |

Symbols, Emoticons and Special Characters

You can add various symbols and special characters to Confluence pages. You can also use them in other places that display content, such as blog posts, comments, and the dashboard welcome message.

Insert symbols and special characters

1. Edit the page (if you're viewing the page, press **E** on your keyboard)
2. Choose **Insert** > **Symbol**
3. Choose a symbol to insert it

Screenshot : Available symbols

On this page:
- Insert symbols and special characters
- Insert emoticons
- Prevent emoticons from appearing

Related pages:

- The Editor

Symbol

Insert emoticons
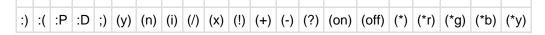
There are two ways to add an emoticon, or smiley, to your page.

By choosing an emoticon from those available:

1. Choose **Insert** > **Emoticon**
2. Choose an emoticon to insert it

By typing a character combination:
You can insert emoticons by typing commonly-used character combinations. The below table shows the emoticons (and other icons) available in Confluence, and the character combinations that create them:

:)	:(	:P	:D	;)	(y)	(n)	(i)	(/)	(x)	(!)	(+)	(-)	(?)	(on)	(off)	(*)	(*r)	(*g)	(*b)	(*y)

Prevent emoticons from appearing

To undo the conversion of a character combination into an emoticon, press **Ctrl+Z** (Windows) or **Cmd+Z** (Mac).

To prevent Confluence from converting text to emoticons automatically, disable 'Autoformatting' in your user profile. See Edit Your User Settings.

Move and Reorder Pages

The easiest way to set a page's location in Confluence is to navigate to the space where you want the page to live and, if necessary, find its parent page and choose **Create**. Sometimes though, you'll want to change a page's location either while you're creating it, or after it's been created.

You can also move and reorder pages in the page tree (hierarchy).

Set page location or move a page

1. Do either of the following:
   - **While creating a page** – choose the location icon

     at the top of the page
   - **Once a page is created** – choose

     •••

> **Move**

2. Use the tabs on the left of the 'Set Page Location' dialog to help you find the new space and/or parent page for your page (the **Current location** and **New location** breadcrumbs at the bottom of the dialog indicate the current parent page and new parent page)
3. Select **Reorder** if you want to move the page to a different position amongst the child pages (when you choose **Move** in the next step, you'll be able to reorder the page)
4. Choose **Move** (If you're reordering the child pages, choose the new position for the page and choose **Reorder**)

On this page:

- Set page location or move a page
- Reorder pages within a space
- Notes about permissions

Related pages:

- Create and Edit Pages
- Copy a Page
- Delete or Restore a Page
- Orphaned Pages

The page – along with any attachments, comments, and child pages – is moved to your chosen location. Confluence will automatically adjust all links to the moved pages, to point to the page(s) in its new location.

When completing the **New parent page** field, you need to select the page suggested by Confluence's autocomplete. Typing or pasting the page name (or using your browser's autocomplete) won't work.

Screenshot: Setting the location or moving a page

Set Page Location

Advanced

Search

Recently Viewed

Browse

Specify the new parent page for this page and its children by space a

New space: Confluence Development

New parent page: Confluence Development home

Start typing a page title to see a list of suggestions.

Current location: Dev Team › Dev Team › Product requirements ›

New location: Confluence Development › Confluence Develop

Reorder pages within a space

You can change the location of a page within its space, and reorder pages in the hierarchy. This allows you to:

- Move a single page, or a family of pages, to a different parent within the space.
- Reorder pages that are children of the same parent.

All links to the page are maintained. When you move a parent page, the entire hierarchy of child pages will move too.

To move or reorder a page:
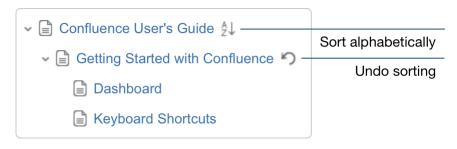
1. Go to the space and choose **Space tools** > **Reorder pages** from the bottom of the sidebar
2. Expand the branches to locate the page you want to move
3. Drag the page to a new position in the tree

> If you're using the Documentation theme go to **Browse** > **Pages** > **Tree** and then drag the pages to a new position as above.

Alternatively, you can choose to order a group of child pages alphabetically by choosing the **So**

rt Alphabetically (A-Z) icon. The **Sort Alphabetically** (A-Z) icon only appears next to the parent page if the page family is currently sorted manually.

If you change your mind, you can use the **Undo Sorting** icon to revert back to the previous manual page order. This option is only available immediately after sorting the page, while you're still on the **Reorder Pages** tab, and haven't performed any other action.

> ˅ 📄 Confluence User's Guide ⤋   ──── Sort alphabetically
> 　 ˅ 📄 Getting Started with Confluence ↺ ──── Undo sorting
> 　 　 📄 Dashboard
> 　 　 📄 Keyboard Shortcuts

Notes about permissions

To move a page, you need the following permissions:

* 'Add' permission on the page you're moving, and
* 'View' permission on the page's parent page. If you're moving the page to a different parent, you need 'View' permission on the new parent.

To move a page into a different space, you also need:

* 'Delete' permission on the space you're moving from, and
* 'Add' permission on the space you're moving to.

If the page has restrictions, and you want to keep the page restrictions in the new location, you'll also need 'Restrict' permission on the space you're moving to. Alternatively, remove the page restrictions before performing the move.

Copy a Page

If you need to duplicate the content of a page, the easiest way is to copy the page.

When you copy the page it'll need to be renamed if you intend to keep it in the same space – and Confluence will rename the page 'Copy of [originalPageName]' by default – as pages in a space must have unique names. If you plan on moving it to another space, you can keep the same name as long as you set the location while you're creating the copy. See Move and Reorder Pages for more information.

> **Related pages:**
>
> * Create and Edit Pages
> * Move and Reorder Pages

ℹ️ You need the 'Create Pages' permission in a space to copy pages within that space. See space permissions for more information.

To copy a page:

1. Go to a page in the space and choose
   ●●●
   > **Copy**
   Confluence will open a copy of the page in the editor. By default, Confluence will name the page 'Copy of [originalPageName]'.
2. Rename the page and make any other changes required in the body of the page
   If you need to move the new page to a different space or a different parent, choose **Location** at the bottom-left of the page. Refer to the instructions on moving a page.
3. Choose **Save**

Notes

* Copying a page will duplicate all of the original page's attachments and labels, but won't copy comments from the original page.
* This method of copying a page doesn't copy the child pages. Please add your vote to issue CONF-2814 if you'd like to see this improvement.
* You can copy an entire space, including all pages in the space (parents and children) — refer to these instructions on copying a space.

Delete or Restore a Page

When you delete a Confluence page, it moves to the space's trash much like an email moves to your email trash or a file moves to the trash on your computer. It's not permanently deleted, and is still retrievable, until you purg e the page from the trash.

Delete a page

When you delete a page in Confluence, you're deleting all versions of the page. If you only want to delete a specific version of a page, take a look at the instructions below for deleting a specific version.

To delete a page and all its versions:

Go to the page and choose
•••
> **Delete** (confirm you want to delete the page when prompted)
Note: The **Delete** option will only appear if you have permission to delete this page. See Space permissions.

If someone has applied page restrictions to the page, the restrictions may also prevent you from deleting the page.

To delete a specific version of a page:

If you want to delete a specific version of a page, you need to be a space administrator.

⚠ Deleting a page version is permanent. It won't be moved to the trash, so you can't restore a deleted version.

1. Go to the page and choose
   •••
   > **Page History**
2. Choose **Delete** next to the version you want to delete (confirm the action when prompted)

Once you've deleted a version, the other versions are re-numbered where necessary. For example, if you delete version 2, version 3 will become the new version 2.

Delete a page with children

If you delete a page that has any child pages, the child pages moves up to the nearest parent page. If you want to keep the child pages under another parent, it's a good idea to move them before deleting their former parent.

Screenshot: Warning when deleting a page with children

> **Warning**
>
> ⚠ This page has 1 child page(s). Removing this page will result in the parent page adopting all this pages child pages.
>
> [OK] Cancel

> If you want to remove a parent page *and* a large number of child pages:
>
> 1. Create a temporary new space
> 2. Move the parent page to the new space (the child pages will move too)
> 3. Delete the space

Restore a deleted page

When you restore a Confluence page, you're moving it from the trash to the root of the space. You may want to move the page to another location after you've restored it, as it won't be easy to find at the root of the space.

You need to be a space administrator to restore a deleted page.

To restore a deleted page:

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar
2. Choose **Content Tools** > **Trash**
   A list of deleted pages and blog posts for the space displays.
3. Choose **Restore** for the page you wish to restore

⌄ Is your space using the Documentation theme?

   If your space is using the Documentation theme:

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

      Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
   2. Choose **Trash** from the space administration options
      A list of deleted pages and blog posts for the space displays.
   3. Choose **Restore** for the page you wish to restore

To find the page after you've restored it, choose **Pages** to go to the root of the space (the page isn't restored to its original position in the page hierarchy). See Move and Reorder Pages for more information.

If a new page has already been created in that space with the same name as the page you're restoring, you'll be given an option to rename the page before it's restored.

Purge deleted pages

If you want to permanently get rid of a page, you need to purge it from the trash. That'll mean the page, and all its versions and attachments, are gone for good.

⚠ Attachments of deleted pages will remain in the database (allowing potential retrieval) until the trash is purged. Once the trash is purged, all attachments on the page(s) are permanently removed.

You need to be a space administrator to purge deleted pages.

To purge deleted pages:

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar

2. Choose **Trash**
   A list of deleted pages and blog posts for the space displays.
3. Choose **Purge** for a specific page or **Purge All** to permanently clear all deleted pages and news items

∨ Is your space using the Documentation theme?
   If your space is using the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **Trash** from the space administration options
   A list of deleted pages and blog posts for the space displays.
3. Choose **Purge** for a specific page or **Purge All** to permanently clear all deleted pages and news items

Add, Remove and Search for Labels

Labels are key words or tags that you can add to pages, blog posts and attachments. You can define your own labels and use them to categorise, identify or bookmark content in Confluence.

For example, you could assign the label 'accounting' to all accounts-related pages on your site. You can then browse all pages with that label in a single space or across the site, display a list of pages with that label, or search based on the label.

Because labels are user-defined, you can add any word that helps you identify the content in your site.

You can also apply labels (known as categories) to spaces, to help organise your Confluence spaces. See Use Labels to Categorise Spaces.

Label a page or blog post

Any user with permission to edit a page can add labels to it. Any existing labels appear at the bottom-right of the page, below the page content.

To add a label to a page or blog post:

On this page:

- Label a page or blog post
- Label an attachment
- The Labelled content page
- Search by label
- Search for labelled pages using a URL
- Remove labels
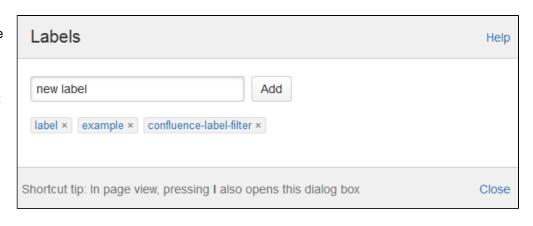
Related pages:

- Use Labels to Categorise Spaces
- Display Pages with Label Macros

1. Go to the page choose the edit icon 🖉 beside the list of labels or press **L** on your keyboard
2. Type in a new label (existing labels are suggested as you type)
3. Choose **Add**

> If you're editing or creating a page, and you want to add labels, choose the label icon
>
> 🏷
>
> at the top of the page.

Label an attachment

1. Do either of the following:
   - Go to the page that contains the attachment and choose Go to
     ●●●
     > **Attachments**
   - Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar, then choose **Attachments**
   You'll see a list of attachments, with any existing labels listed in the **Labels** column.
2. Choose the edit icon
   ✎
   beside the list of labels and type in a new label (existing labels are suggested as you type)
3. Choose **Add**

You can also add labels in a list of attachments displayed by the Attachments macro, by choosing the edit icon

✎
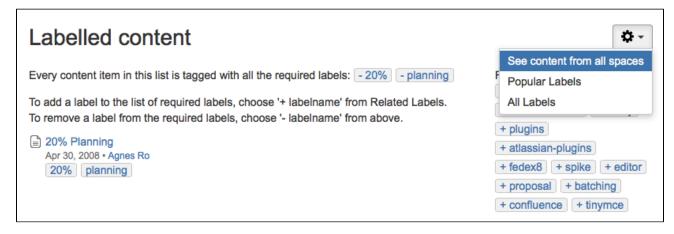
beside each label.

ℹ If you add one or more labels to a template, that label will be copied to the page when someone adds a page based on that template. See Creating a Template.

The Labelled content page

If you're viewing a page or post that has labels or displays the Attachments macro, you can choose any label to go to the **Labelled content** page for the space. Choose **Popular Labels** or **All Labels** from the cog at the top-right – to view the most-used labels or all labels in the space – or choose **See content from all spaces** from the cog to view labelled content from all spaces in your Confluence site.

Screenshot: The Labelled content page

The **Popular Labels** option displays a word cloud, where the bigger a label is displayed, the more popular it is. Choose any label to view content tagged with that label.

You can also navigate to the labels view for a space by entering the following URL (replace SPACEKEY with the space's key):

```
<your.Confluence.site>/labels/listlabels-alphaview.action?key=SPACEKE
Y
```

⌄ Is the space using the Documentation theme?

If the space is using the Documentation theme, access the **Labelled content** page by choosing **Browse > Labels**

Search by label

You can use the 'labelText:' prefix to search specifically for content that has a specific label. For example, if you're looking for pages with the label 'chocolate', type labelText:chocolate into the search field in the Confluence header. For more examples of searching by label, see Confluence Search Syntax.

Search for labelled pages using a URL

Entering a URL with an appended label or labels is another way to search for pages with particular labels.

In your browser's address bar, enter the following URL and press enter: http://<your.Confluence.sit e>/label/foo+bar

The **Labelled content** page will load, showing search results for pages with the **both** labels, 'foo' and 'bar'. Replace 'foo' and 'bar' with the label(s) you want to search for, and separate multiple labels with a + symbol.

Adding a label to your results:

Once you're on the **Labelled content** page, you can add more labels to your search by choosing them from the **Related Labels** list at the top-right of the page. Each label is listed with a plus (+) sign.

If you want to remove labels from your search, locate the included labels at the top of the page and choose the label(s) you want to remove. Each included label will be listed with a minus (–) sign.

Remove labels

When viewing page, blog post, or attachment labels, an **x** appears alongside each label. Choose the **x** to remove the label.

> If you have deleted pages that contain a label, you may need to purge the deleted pages from the space's trash to ensure that the label disappears too.

Display Pages with Label Macros

Using labels and macros, you can categorise pages and then display them in Confluence in a number of ways.

As an example, you could label all pages relevant to the marketing team with 'marketing', and then add more specific labels like 'online', 'mobile', and 'physical' to different pages where required.

You could then use the Content by Label Macro to display different combinations of pages with the marketing label. Some combinations you could use would be:

- All pages with the label 'marketing'.
- Pages with all of the following labels: 'marketing', 'mobile', and 'online'.
- Pages with either the 'mobile' or 'online' labels, in the Marketing space.

Related pages:

- Add, Remove and Search for Labels
- Macros
- Use Labels to Categorise Spaces

There are a lot of ways you can filter the content, making it easier for you to find content that's relevant to

you.

Other label macros

Here are some other macros that use labels, and can help you categorise and display your content.

Navigation Map macro

The Navigation Map macro renders the list of pages associated with a specified label as a navigation map.

Related Labels macro

The Related Labels macro lists labels commonly associated with the current page's labels.

Content by Label macro

The Content by Label macro displays a list of content marked with specified labels.

Content Report Table macro

The Content Report Table macro displays a set of pages and blog posts in tabular format, based on the specified labels.

Labels List macro

The Labels List macro lists all labels of a space, grouped alphabetically.

Recently Used Labels macro

The Recently Used Labels macro lists labels most recently used in a specified scope - global (site), space, or personal.

Popular Labels macro

The Popular Labels macro displays popular labels in a list or in a heatmap (also called a cloud).
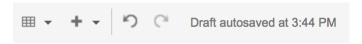
Drafts

A draft is a 'snapshot' of unsaved page content, which Confluence creates automatically at regular intervals while you're creating or editing a page or blog post. Drafts can minimise the loss of work if your Confluence site experiences a problem, as you can retrieve the page content from the last saved draft.

By default, Confluence saves a draft of your page once every thirty seconds, however, a Confluence administrator can configure how often drafts are saved. In addition, whenever you edit a page and then move to another screen, Confluence will automatically save a draft. Each new draft replaces the previously saved draft.

Each time Confluence saves a draft, it displays a message and the time of the last save, to the right of the **undo** and **redo** buttons in the editor toolbar. When you edit a page that wasn't successfully saved, Confluence will let you know that a version of the page you're editing wasn't saved, and will give you the option to resume editing it from the draft.

Screenshot: Message displaying the time when the draft was last saved

On this page:

- View drafts
- Resume editing a draft
- View your unsaved changes

Related pages:

- Configuring the Time Interval at which Drafts are Saved
- Concurrent Editing and Merging Changes
- Your Profile and Settings

More about drafts:

- You only have access to drafts of pages you've been working on, and whose content hasn't yet been saved.
- You can't create a draft explicitly.
- Your drafts are listed in the 'Drafts' tab of your profile.
- Once you've resumed editing a draft, or chosen to discard it, the draft is removed from your drafts tab.

View drafts

Your drafts are listed on the **Drafts** tab of your user profile. To see your drafts, choose **your profile picture** a t top right of the screen, then choose **Drafts**.

Resume editing a draft

There are two ways to resume editing an unsaved page or blog post:

1. Using your drafts view. To resume editing a draft from this view:
   a. Choose **your profile picture** at top right of the screen, then choose **Drafts**.
   b. Choose **Resume Editing** next to the appropriate draft to resume editing that draft.
      ⓘ If you hadn't entered a page title, the draft will be named '*Untitled*'.
2. If you create a new page or blog post and didn't save it, but a draft is saved, when you next add a page or blog post in that space, Confluence will ask you if you want to resume editing the page. If you choose **resume editing**, the draft and its unsaved content will be restored, allowing you to continue editing it.

What happens if I'm editing the draft of a page that has since been updated?

Confluence will display a message informing you that you're editing an outdated page. If there are no conflicts between the two versions, Confluence will give you the option to **Merge and Resume editing**.

If there are conflicts, Confluence will give you the option to **View the Conflict** or to **Discard** your changes.

View your unsaved changes

When you edit a page or blog post, you can view any 'unsaved' changes you've made since the last automatically saved draft, by clicking the **Draft autosaved at** message in the toolbar.

Screenshot: Segment of the unpublished changes window

Concurrent Editing and Merging Changes

Sometimes, another user may edit the same page as you're editing, at the same time you do. When this happens, Confluence will do its best to ensure nobody's changes are lost.

How will I know if someone else is editing the same page as I am?

If another user is editing the same page as you, Confluence will display a message above your edit screen letting you know who the other user is and when the last edit was made.

Screenshot: Notification of Simultaneous Page Editing

> (i) This page is also being edited by Ewan User. Your changes will be merged with theirs when you save. ✕

Related pages:
- Page History and Page Comparison Views
- Drafts

What happens if two of us are editing the same page and the other user saves before I do?

If someone else has saved the page before you, when you click **Save**, Confluence will check if there are any conflicts between your changes and theirs. If there are no conflicting changes, Confluence will merge the changes.

If there are conflicts, Confluence will display them for you and give you the option to:

- **Continue editing** - Continue to edit the page; useful if you want to manually merge the changes.
- **Overwrite** - Replace the other person's edits with yours (their edits will not be included in the latest version).
- **Cancel** - Discard your changes and exit the editor, keeping the other person's edits.

Example Scenario

For example, Alice and Bob both edit the same page at the same time.

If Alice clicks save before Bob, Bob is now effectively editing an out-of-date version of the page. When Bob clicks save, Confluence will examine his changes to see if any overlap with Alice's. If the changes don't overlap (i.e. Alice and Bob edited different parts of the page), Bob's changes will be *merged* with Alice's automatically.

If Bob's changes overlap with Alice's, Confluence will display an error message to Bob showing where Alice has changed the page, and giving Bob the options to overwrite Alice's changes with his own, to re-edit the document to incorporate Alice's work, or to cancel his own changes entirely, maintaining Alice's changes.

Page Restrictions

Page restrictions allow you to control who can view and/or edit individual pages in a space. So, if you're working on a page that shouldn't be viewed by just anybody, it's easy to lock it down to the people who need to know. You can add restrictions for individuals or for Confluence groups.

To add or remove page restrictions, you'll need to have permissions to edit the page and 'Restrict' or 'Admin' permission in the space.

On this page:
- Restrict a page or blog post
- Remove restrictions from a page
- Request and grant access to view a restricted page
- View all restricted pages in a space
- Notes

Restrict a page or blog post

To restrict who can view or edit a page or blog post:
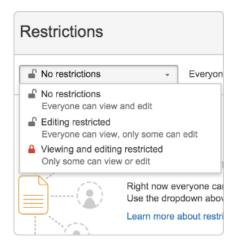
1. Chose the Restrictions
   🔓
   icon at the top of the page.
2. Choose whether you just want to limit only who can **Edit**, or who can **View and / or Edit**

3. Enter users or groups then click **Add** to add them to the list. If you chose **Viewing and Editing restricted** you can further specify for each person or group whether they can edit or just view the page.
4. **Apply** the restrictions.

You can add as many users and/or groups as you need.

In this example, some users and groups can view only, others can also edit, plus there are inherited restrictions that might impact who can view the page.

Speed it up
Apply the same restriction to multiple people and groups

Watch out
Restrictions on other pages can affect this one

Be specific
Choose exactly what each group or person can do

Who is 'everyone'?

When we say "everyone can view this page" *everyone* means all the people who can view the page by default. There's two things that can affect who can view a page - the space permissions, and view restrictions on any parent pages that are being inherited.

Restrictions don't override a person's space permission.   For example, if you say a person 'can view' in the restrictions dialog and they don't have 'view' permissions for the space, they won't be able to see the page.

How do inherited restrictions work?

View restrictions are inherited, which means a restriction applied to one page will cascade down to any child pages. Edit restrictions are not inherited, which means pages need to be restricted individually.

The restrictions dialog will tell you when there are inherited restrictions that might affect who can view your page.

Here's the basics:
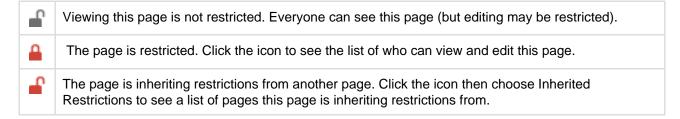
- If you restrict **viewing** to a person or group, only they will be able to see that page and all its child pages (unless there are further restrictions on the child pages)
- If you restrict **editing** to a person or group, they'll be able to see and edit that page, plus see its child pages.
- Parent pages (higher up in the page hierarchy) can have their own view restrictions that may prevent people from viewing your page.

If the person you've listed as a viewer or editor can't see the page, check to make sure:

- they have **View** space permission for that space, or
- there's no view restriction on a page higher up the page hierarchy that prevents them seeing any children of that page.

View current page restrictions

The restrictions icon at the top of the page gives you a clue that the page has restrictions:

🔓	Viewing this page is not restricted. Everyone can see this page (but editing may be restricted).
🔒	The page is restricted. Click the icon to see the list of who can view and edit this page.
🔓	The page is inheriting restrictions from another page. Click the icon then choose Inherited Restrictions to see a list of pages this page is inheriting restrictions from.
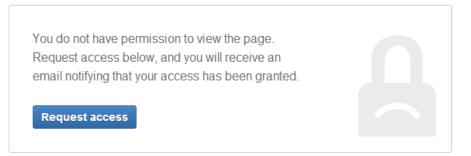
Remove restrictions from a page

Removing restrictions is easy. Choose **No restrictions** to remove all restrictions, or click **Remove** next to each person or group in the list if you want to change who can view or edit the page.

Request and grant access to view a restricted page

If you navigate to a page that you're not able to view because it has page restrictions applied (for example from a link or page URL) you may be able to request access to the page.

You do not have permission to view the page. Request access below, and you will receive an email notifying that your access has been granted.

Request access

To request access to a restricted page:
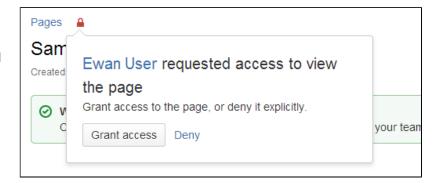
1. On the restricted page choose **Request access**
2. Wait for an email confirming that access has been granted

If the request access message doesn't appear, you're not able to request access for that particular page. This usually is because the page has inherited view restrictions from a parent page, or you may not have adequate space permissions.

To grant access to a restricted page:

1. In the request access email, choose **Grant access** You'll be taken to the restricted page, and a dialog will appear with the access request
2. Choose **Grant access**

Ewan User requested access to view the page

Grant access to the page, or deny it explicitly.

Grant access   Deny

The user will receive an email confirming that access has been granted.

This process is the same as navigating to

••• > **Restrictions** and adding a 'View' restriction for the user.

Who can grant access?

To grant access to a restricted page you will need to have permission to edit that page, and have the 'Restrict' or 'Admin' permission for the space.

Confluence will send an email to a user who can grant permissions. The sequence that Confluence will use to search for the appropriate user is:

1. The last person to edit the page
2. **All** non-admin users that can set permissions on the page (given that an admin user can always set permissions)
3. The page creator
4. **All** the space administrators

Confluence will try each of these roles in turn, emailing the first user that has appropriate permissions.

View all restricted pages in a space

You need space admin permissions to view the list of restricted pages in a space.

To view restricted pages:

1. Go to the space and choose **Space tools** > **Permissions** from the bottom of the sidebar
2. Choose **Restricted Pages**

If your space uses the Documentation theme:

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

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

2. Choose **Restricted Pages** in the space administration options

Screenshot: Restricted pages in a space

Permissions	Restricted Pages				
This report shows all the pages where access is restricted to specific groups or users by page level permissions.					
Title (Space)	**Type**	**Permitted User/Group**	**Creator**	**Created**	**Operations**
Induction plan - Sophie User (Stars)	**Edit**	leah	Leah Admin	Feb 18, 2013	🔓
Induction plan - Ewan User (Stars)	**Edit**	leah	Leah Admin	Feb 18, 2013	🔓
Draft status report (Stars)	**View**	admin	Rach Admin	Feb 18, 2013	🔓
	View	sophie	Rach Admin	Feb 18, 2013	🔓

Notes

- **You can't exclude yourself**
  When you apply a restriction, Confluence will automatically add you to the list. You can't remove yourself from this list.
- **Space Admin and System Administrator access to restricted pages**
  Users with 'Admin' permissions in a space, or users with the System Administrator global permission can remove restrictions from pages, even if the page restriction prevents them from viewing the page. Go to **Space Administration** > **Restricted Pages**.

Links

You can create links to pages, blog posts, anchors, attachments, external

websites, JIRA issues and more. Links can be text or images, and can be added in many different ways.

Links to pages within your Confluence site are relative, which means that you can move pages and rename pages without breaking links.

This page explains the most common ways to create links.

Insert a text link

To insert a link on a page:

1. Select some text or position your cursor where you want to insert the link
2. Choose **Link** on the toolbar or use the keyboard shortcut **Ctrl+K**
3. Select a page, blog post or attachment, or enter an external URL (see below for how to link to particular types of content)
4. Enter or modify the link text (this is the text that will appear on the page. If this field is left blank, the page name or URL will be used as the link text.)
5. Choose **Insert**

Alternatively you can type **[** followed by the page or attachment name. Autocomplete will suggest matching items for you.

You can also paste a URL directly onto your page. Confluence will automatically create the link, and if the URL is for a page in the current site, the page name will be set as the link text.
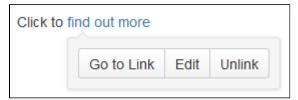
Insert an image link

1. Select an image on your page
2. Choose **Link** on the Image Properties toolbar
3. Select a page, blog post or attachment, or enter an external URL (see below for how to link to particular types of content)
4. Choose **Insert**

Modify a link

1. Select the link text or image
2. Choose **Edit** from the link properties toolbar
3. Modify the link and choose **Save**

Remove a link

1. Select the link text or image
2. Choose **Unlink** from the properties toolbar

Link to specific types of content

Confluence supports many methods for creating links. Some of the common ones are listed here.

Type of link	Ways to do this
Link to a page	Choose **Link** > **Search** then enter part of the page name.
	Choose **Link** > **Recently viewed** and select a page from the list.
	Type **[** and enter part of the page name then select the page from the list.
	Paste the URL of the page onto your page (Confluence will automatically create the link).
Link to a page in another space	Choose **Link** > **Search** enter part of the page name and select **All Spaces** from the drop down.
	Choose **Link** > **Advanced** then enter the space key followed by the page name `spacekey:mypage`.
	Type **[** and enter part of the page name then select the page from the list. (you can hover over each suggestion to see which space the page is from).
Link to a blog post	Choose **Link** > **Search** and enter part of the blog post name.
	Type **[** and enter part of the blog post name then select the blog post from the list.
Link to an attachment or image on this page	Choose **Link** > **Attachment** then upload or select an attachment from the list.
	Type **[** and enter part of the attachment file name then select the attachment from the list.
Link to an attachment on another page	Choose **Link** > **Search** and enter part of the attachment name.
	Type **[** and enter part of the attachment file name then select the attachment from the list (you can hover over each suggestion to see which space the page is from).
Link to a website	Choose **Link** > **Web Link** then enter the website URL.
	Type or paste the URL onto the page (Confluence will automatically create the link).
Link to an email address	Choose **Link** > **Web Link** then enter the email address.
	Type or paste the email address onto the page (Confluence will automatically create a 'mailto:' link).
Link to an anchor on a page	Choose **Link** > **Advanced** then enter the anchor name in one of the formats below.
	For an anchor on this page: `#anchor name`.
	For an anchor on another page in this space: `page name#anchor name`.
	For an Anchor on another page in another space: `spacekey:page name#anchor name`.
	See Anchors for more information on using anchors.

Link to a heading on a page	Choose **Link** > **Advanced** then enter the heading in one of the formats below. Heading text is case sensitive and must be entered without spaces.
	For a heading on this page: `#MyHeading`.
	For a heading on another page in this space: `Page Name#MyHeading`.
	For a heading on another page in another space: `spacekey:Page Name#MyHeading`.
	Be aware that these links will break if you edit the heading text. Consider using the Table of Contents macro or an Anchor instead.
Link to a comment on a page	Go to the comment, right click the **Date** at the bottom of the comment and copy the link. Paste the link directly onto your page or choose **Link** > **Web Link** and paste in the URL.
	Type **[$** then enter the Comment ID ('12345' in this example): `[$12345]`
Link to an undefined page (a page that does not exist yet)	Choose **Link** > **Advanced** then enter the new page name (a page will be created on click).
	Type **[** then enter the new page name then choose **Insert link to create page** .
	See Undefined Page Links for more information on undefined pages.
Link to a personal space or user profile	Choose **Link** > **Search** then enter the user's name and select their personal space homepage or their profile from the list.
	Type **[** then enter the user's name and select their personal space homepage or their profile from the list.
Link to a JIRA issue (where Confluence is connected to JIRA)	Paste the JIRA issue URL - Confluence will automatically create a JIRA Issue macro.

Link to Confluence pages from other websites

The best way to link to a Confluence page from outside Confluence, for example on a website or in an email message, is to use the tiny link which is a permanent URL. This ensures that the link to the page is not broken if the page name changes.

To access the permanent URL for a page:

1. View the page you wish to link to.
2. Choose **Tools** > **Link to this page**.
3. Copy the **Tiny Link**.
4. Use the tiny link in your website or email message.

You do not need to use the tiny link to link to pages within your Confluence site. Confluence automatically updates links when you rename or move a page to another space.

If you want to link to specific content such as anchors, headings or comments you need to use the following link syntax. Note that there are no spaces in the page name, anchor name or heading text.

In the examples below, the anchor name is 'InsertLinkAnchor' and the heading text is 'Insert a link'.

Purpose	Link syntax

Link to an anchor (from an external website)	`http://myconfluence.com/display/spacekey/pagename#pagename-anchorname`
	Example from this page:
	`https://confluence.atlassian.com/display/DOC/Working+with+Links#Workingw`

Link to a heading (from an external website)	`http://myconfluence.com/display/spacekey/pagename#pagename-headingtext`
	Example from this page:
	`https://confluence.atlassian.com/display/DOC/Working+with+Links#Workingw`

Link to a comment (from an external website)	`http://myconfluence.com/display/spacekey/pagename?focusedCommentId=commer`
	Example from this page:
	`https://confluence.atlassian.com/display/DOC/Working+with+Links?focusedCo`

Some things to note when linking to anchors from a website or email message:

- The page name is repeated in the URL, after the # sign. The second occurrence of the page name is concatenated into a single word, with all spaces removed.
- There is a single dash (hyphen) between the concatenated page name and the anchor name.
- The anchor name in the full URL is concatenated into a single word, with all spaces removed.
- The anchor name is case sensitive. You must use the same pattern of upper and lower case letters as you used when creating the Anchor.

Link to a comment

You can add a link to a comment by using the comment URL (a permanent link), or by using wiki markup to link to the Comment ID.

To find out the comment URL and comment ID:

1. Go to the comment you wish to link to
2. Choose the **Date** at the bottom of the comment and examine the URL

The number after 'comment-' is the Comment ID. An example is shown here.

```
https://confluence.atlassian.com/display/DOC/Working+with+Links?focus
edCommentId=368640803#comment-368640803
```

You can use wiki markup directly in the editor to link to a comment. Enter **[$** followed by the Comment ID, for example `[$12345]` where '12345' is the Comment ID.

Using shortcut Links

If you have configured shortcut links on your Confluence site, then you can link to an external site using a shortcut link that looks like this:`CONF-17025@jira`.

Our Confluence site (where this documentation is housed) is configured to allow shortcut links to our JIRA site, using the shortcut `@jira`. So the shortcut link `CONF-17025@jira` produces this link.

To add a shortcut link using the 'Insert Link' dialog:

1. Choose **Link** > **Advanced** and enter or paste the shortcut link into the **Link** field (shortcut links are case-insensitive)
2. Modify or enter link text (this is the text that will appear on the page)
3. Choose **Insert**

You can also type '[' and choose **Insert Web Link** > **Advanced** to enter a shortcut link.

See Configuring Shortcut Links for more details.

Trackback

Trackback enables two sites can stay informed each time one site refers to the other using trackback 'pings'.

In Confluence, Trackback can be enabled by a site administrator in the Administration Console. When Trackback is enabled, any time you link to an external webpage that supports Trackback Autodiscovery, Confluence will send a Trackback ping to that page to inform it that it has been linked to.

Confluence pages also support Trackback Autodiscovery, and when Trackback is enabled, can receive trackback pings sent by other sites.

To see who has sent a Trackback ping to a Confluence page:

1. Go to the page
2. Choose **Tools** > **Page Information**

Any Trackback pings the page has received will be listed under the page's **Incoming Links**.

Confluence incoming trackback pings only work with referenced pages that are public (anonymously viewable).

Anchors

You can use anchors to enable linking to specific locations on a page, and they can be especially useful for allowing your readers to navigate to specific parts of a long document. Anchors are invisible to the reader when the page is displayed.

There are two steps to using an anchor:

Step 1: Create the anchor

Step 2: Create a link to the anchor

On this page:

- Step 1: Create the anchor
- Step 2: Create a link to the anchor
- Notes

Related pages:

- Links

Step 1: Create the anchor

Add the Anchor Macro to mark the location you want to link to:

1. Do either of the following in the Confluence editor:
    - Choose **Insert** > **Other Macros**, then find and select the Anchor macro
    - Type **{** and the beginning of the macro name, then select the Anchor macro
2. Enter the **Anchor Name** (For example, 'bottom' or 'important information')
3. Choose **Insert**

Macro options (parameters)

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description

Anchor Name	None	This is the anchor name that you will use when creating the link.
		- The anchor name can include spaces. Confluence will remove the spaces automatically when building a URL that points to this anchor. - The anchor name is case sensitive. You must use the same pattern of upper and lower case letters when creating the link as you used when creating the Anchor macro.

Step 2: Create a link to the anchor

You can link to an anchor from:

- A page on the same Confluence site. The link may be on the same page as the anchor, another page in the same space, or a page in another space on the same Confluence site.
- Another web page or another Confluence site, using a specifically formatted URL.
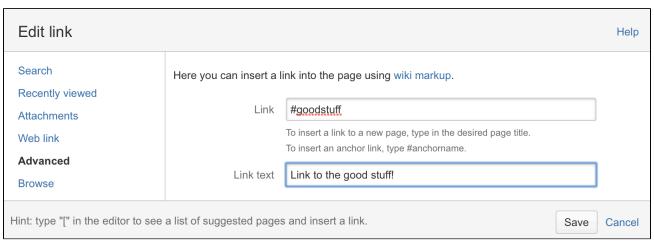
Link to an anchor on the same Confluence site:

1. Select some text or position your cursor where you want to insert the link
2. Choose **Link** in the toolbar or press **Ctrl+K**
3. Choose **Advanced** and enter the anchor name in the **Link** field, following the format below.

Anchor location	Link syntax for anchor	Examples
Same page	#anchor name	#bottom #important information
Page in same space	page name#anchor name	My page#bottom My page#important information
Page in different space	`spacekey:page name#anchor name`	DOC:My page#bottom DOC:My page#important information

ⓘ The anchor name is case sensitive, so you need to use the same pattern of upper and lower case letters as you used when creating the Anchor macro. You need to enter page and anchor names with spaces when you link to them in the same Confluence site.

4. Enter or modify the **Link Text** (this is the text that will appear on the page. If this field is left blank, the page name or URL will be used as the link text)
5. Choose **Save**

Screenshot: The 'Advanced' option in the link dialog

Edit link   Help

Search
Recently viewed
Attachments
Web link
Advanced
Browse

Here you can insert a link into the page using wiki markup.

Link: #goodstuff

To insert a link to a new page, type in the desired page title.
To insert an anchor link, type #anchorname.

Link text: Link to the good stuff!

Hint: type "[" in the editor to see a list of suggested pages and insert a link.   Save   Cancel

Link to an anchor from another web page or another Confluence site:

Use a full URL in the following format:

Link syntax	Examples
`http://myconfluence.com/display/spacekey/pagename#pagename-anchorname`	http://myconflu...
	http://myconflu...

Notes about the full URL:

- The page name is repeated in the URL, after the # sign. The second occurrence of the page name is concatenated into a single word, with all spaces removed.
- There is a single dash (hyphen) between the concatenated page name and the anchor name.
- The anchor name in the full URL is concatenated into a single word, with all spaces removed.
- The anchor name is case sensitive.

Notes

- **Table of contents on a page:** Consider using the Table of Contents Macro to generate a list of links pointing to the headings on the page. The list of links will appear on the page, and will be automatically updated each time someone changes the wording of a heading.
- **Linking to headings:** You can link directly to the headings of a page. See Links. However, if someone changes the wording of a heading, those direct links will be broken. Use the Anchor macro to ensure a lasting link within the body of a page.
- **Site welcome message:** If you are adding an anchor to a page that you are using in the site welcome message, you can only link to that anchor from another page. Internal links within that page will not work.
- **Templates:** When you are previewing a template, a link to an anchor is displayed as a 'broken' link. However, when you create a page using the template the resulting page will have the correct link.

Tables
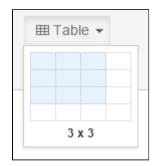
Confluence pages can include multi-row and multi-column tables. You can highlight cells, rows and columns in different colours. When viewing the page, people can sort the table by clicking the column headers.

On this page:

- Insert a table
- What you can do with your table in the editor
- Shortcut keys
- Sort the table in view mode
- Sticky table headings in view mode

Related pages:

- Page Layouts, Columns and Sections
- The Editor

Insert a table

To create a table:

1. While editing the page, place your cursor at the point where you want to insert the table.
2. Choose **Table** on the toolbar.
3. A dropdown menu will appear, showing a table with a variable number of rows and columns. Click in a cell to set the number of columns and rows for your table.
4. Add content and more rows and columns as needed. See below for guidelines on what you can do with your table in the editor.

Screenshot: Inserting a table with 3 rows and 3 columns

What you can do with your table in the editor

While editing a page, place your cursor inside a table to see the table toolbar. Here's a summary of the tools:

	Rows
	• Insert rows before or after the current row • Delete the current row • Cut, copy and paste the current row • Mark a row as a header row (shaded with bold text)
	Columns
	• Insert columns before or after the current column • Delete the current column • Cut, copy and paste the current column • Mark a column as a header column (shaded with bold text)
	Cells
	• Merge selected cells • Split selected cells • Change cell colour
	Table
	• Delete table • Add a numbering column to automatically number each row

Shortcut keys

Windows	Action	Mac OS X
Ctrl+Shift+c	Copy the current table row, or the selected rows.	Cmd+Shift+c
Ctrl+Shift+i	Insert a table. (Opens the Insert Table dialog.)	Cmd+Shift+i
Ctrl+Shift+v	Paste the table rows from your clipboard, placing them above the current row.	Cmd+Shift+v
Ctrl+Shift+x	Cut the current table row, or the selected rows.	Cmd+Shift+x
Alt+Up Arrow	Add a row above the current row.	Alt+Up Arrow
Alt+Down Arrow	Add a row below the current row.	Alt+Down Arrow

For more editor keyboard shortcuts, see Keyboard shortcuts.

Sort the table in view mode

When readers view a table on a page, they can sort the table by clicking the sort icons in the header row.

Screenshot: A colourful, sortable table

Sticky table headings in view mode

In some instances table header rows will stick to the top when you're viewing a page, making those *really* lon g tables easier to read. There are a few exceptions though. Table headers aren't sticky when:

- the space is using the documentation theme.
- your table is inside a page layout, or inside another table.
- your table has no header row or there are cells in the top row that aren't marked as headers.
- your table has a header column, instead of a header row, and scrolls horizontally.
- your table contains another table, that has its own header row.

Add, Assign, and View Tasks

Keep track of the things that need to get done with tasks. With the ability to give each task an assignee and due date, and plenty of ways to see your tasks, you can make sure nothing slips through the cracks.

On this page:
- Add a task
- View tasks
- Notes

Add a task

You can add tasks on any page in Confluence. For example, you might add tasks under action items on a meeting notes page, or in a project planning p age – anywhere you need a lightweight task management solution.

To create a task:

1. In the editor, choose the **Task list**

   ☑

   button or use the keyboard shortcut **[ ]**
2. Start typing your task – @mention someone to assign the task to them, and type **//** and choose a due date

The first person you mention in a task is the assignee; you can even assign tasks to yourself.

Note: If you assign a task to someone who doesn't have permission to view the page or space, they won't see the task.

View tasks

There are a number of ways to keep track of tasks assigned to you, or tasks you've created for others.

On a page

The simplest way to see a task is on the page it was originally created on. It's easy to see if a task is complete, who it's assigned to, and when it's due. If a task is nearing or passed its due date, the colour of the date will change (red for overdue, orange for due in the next 7 days).

In your profile

The tasks page in your profile gives you a place to see all the tasks relevant to you. Easily keep track of the status of tasks assigned to you, and tasks you've created and assigned to others.

To view the tasks page, go to **Profile** > **Tasks**. Use the filters to show tasks that were assigned to you or created by you in the last 6 months, and toggle between complete or incomplete tasks.

If you need to see more than just your last 6 months of tasks, use a Task Report.

In a Task Report

If you're looking for a more custom view of tasks, the Task Report blueprint is a great way to track tasks assigned to a specific team or project.

To create a task report:

1. Choose **Create** > **Task Report**
2. Select the type of report:
   - **Assigned to my team** – for tasks assigned to particular people.
   - **In my project** – for tasks that appear in a specific space or page.
   - **Custom** – for a wide range of filtering options, including by date or page label.
3. Follow the prompts to create the report.

This blueprint uses the Task Report Macro. You can also choose to use this macro on an existing page, for example, on a project or team space homepage.

Meeting actions this month
Created and last modified by Rach Admin a minute ago

Description	Due date ⌄	Assignee	Task appears on
☐ @ Leah Admin to schedule a follow up meeting		Leah Admin	2014-03-25 Meeting notes
☐ @ Ewan User to finish creating the new project space by 📅 10 Mar 2014	10 Mar 2014	Ewan User	2014-03-25 Meeting notes
☐ @ Josh User add designs to the project space's file list by 📅 02 Apr 2014	02 Apr 2014	Josh User	2014-03-29 Meeting notes

Notes

- The date picker can be triggered by typing // or by typing a date in the format dd/mm/yyyy or dd-mm-yyyy. Typing other date formats in the editor won't trigger the date picker.
- Personal Tasks (created in the Workbox in older versions of Confluence) don't appear in the Tasks view or Task Report. To migrate any incomplete personal tasks, go to **Workbox** > **Personal Tasks** and follow the prompts.
- The wiki markup based Tasklist Macro has been removed from the macro browser. If you have a Tasklist macro on a page it will continue to work, but you will be unable to add new Tasklists using this macro.

Autocomplete for links, files, macros and mentions

When using the Confluence editor, you can type a trigger character or press a keyboard shortcut to see a list of suggested links, files or macros to add to your page, or to mention another user (and automatically notify them of this).

Summary of autocomplete

What you want to do	Trigger character	Keyboard shortcut	Description
Add a link on your page	[	Ctrl+Shift+K	See a list of suggested pages or other locations to link to from your page. More...
Display an image, video, audio file or document on your page	!	Ctrl+Shift+M	See a list of suggested images, multimedia files and documents to embed in your page. More...

| Add a macro on your page | { | None | See a list of suggestions as you begin typing a macro name. More... |
| Notify another user by email that you have mentioned them on your page | @ | None | See a list of suggested users to mention. More... |

Using autocomplete for links

Type '[', or press Ctrl+Shift+K, to see a list of suggested pages or other locations to link to from your page. You can link to pages, user profiles, images, documents and other file attachments.

To autocomplete a link:

1. Edit the page.
2. Click where you want to insert a link and do one of the following:
   - Type '[' and then the first few characters of the page title, user's name, image name or file name.
   - Type the first few characters of the page title, user's name, image name or file name (or select
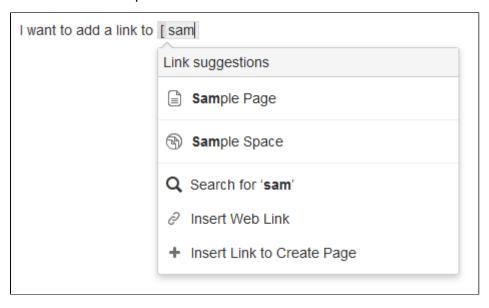
relevant text) and then press Ctrl+Shift+K.
3. Click the relevant link from the list of suggestions.

If the item you need is not in the list, either:

- Choose **Search for 'xxx'** to continue looking for the page within Confluence, or
- Choose **Insert Web Link** to insert a link to an external web page using the link browser.

Screenshot: Autocomplete for a link

I want to add a link to [sam

Link suggestions

📄 **Sam**ple Page

🔄 **Sam**ple Space

🔍 Search for '**sam**'

🔗 Insert Web Link

➕ Insert Link to Create Page

Using autocomplete for images, videos, audio files and documents

You can use the autocomplete as a fast way of embedding images, videos, audio files and documents into your page. Type an exclamation mark or press Ctrl+Shift+M to see a list of suggested images, multimedia files and documents to display on your page. You can use autocomplete to embed the following file types:

- Images – any format that Confluence supports.
- Videos, audio files and all multimedia formats that Confluence supports.
- Office documents supported by the Confluence Office Connector: Word, Excel and PowerPoint.
- PDF files.

ℹ Autocomplete works most efficiently for files that are already attached to the Confluence page.
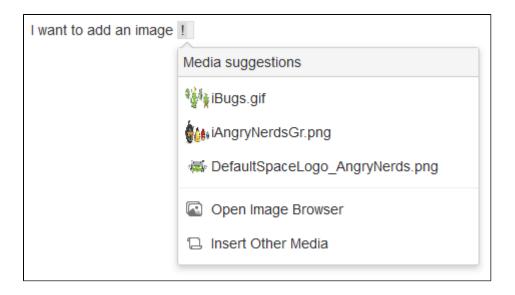
To embed an image, video, audio file or document:

1. Edit the page.
2. Click where you want to insert the image, video, audio file or document and do one of the following:
   - Type '!' and then the first few characters of the image, file or document name.
   - Type the first few characters of the name of the image, file or document (or select relevant text) and then press Ctrl+Shift+M.
3. Choose the relevant file from the list of suggestions.

If the item you need is not in the list, either:

- Choose **Open Image Browser** to find images and documents using the image browser, or
- Choose **Insert Other Media** to embed videos, audio and other multimedia files using the macro browser. Insert the 'Multimedia' macro to display your multimedia file.

Screenshot: Autocomplete for an image or document

I want to add an image !

Media suggestions

iBugs.gif

iAngryNerdsGr.png

DefaultSpaceLogo_AngryNerds.png

Open Image Browser

Insert Other Media

Using autocomplete for macros

Type '{' to see a list of suggested macros to add to your page.

Autocomplete provides access to all available macros in your Confluence site, including any user macros that your administrator has added and made visible to all.
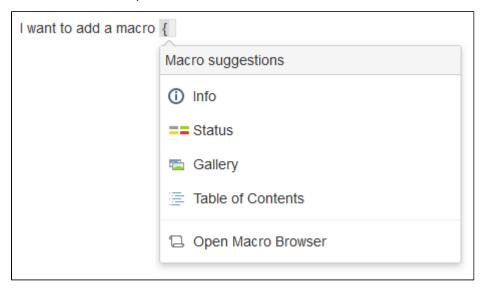
ⓘ You need to know the name of macro. Autocomplete for macros will only match the name of the macro, not the description.

To autocomplete a macro using '{':

1. Edit the page.
2. Click where you want to insert the macro.
3. Type '{' and then the first few characters of the macro name.
4. Choose the relevant macro from the list of suggestions.
5. Configure the macro by completing the form fields as prompted.

If the macro you need is not in the list, choose **Open Macro Browser** in the list of suggestions to continue looking for the macro in the macro browser. See Macros.

Screenshot: Autocomplete for a macro

I want to add a macro {

Macro suggestions

ⓘ Info

Status

Gallery

Table of Contents

Open Macro Browser

Using autocomplete for mentions

You can use autocomplete to automatically notify another Confluence user that you have mentioned them in

a page, blog post, or comment. Type '@' and part of the person's name, to see a list of suggested users.

Note: Use the person's full name. Autocomplete will recognise users' full names only, not their usernames.

Cancelling autocomplete

The autocomplete starts automatically when you press the trigger characters. You may want to close the autocomplete menu or escape from autocomplete once it has started.

There are a few different ways to stop the autocomplete once it has started:

- Press the escape key, 'Esc', on your keyboard.
- Click somewhere else in the editor panel.
- Press an arrow key to move out of the autocomplete area.
- For the *link* autocomplete only: enter a right-hand square bracket, like this: **]**

Enabling and disabling autocomplete

You can turn off the triggering of autocomplete by the '[' and '!' characters. This will prevent the autocomplete from starting automatically when you press one of the trigger characters. You can also turn it back on again.
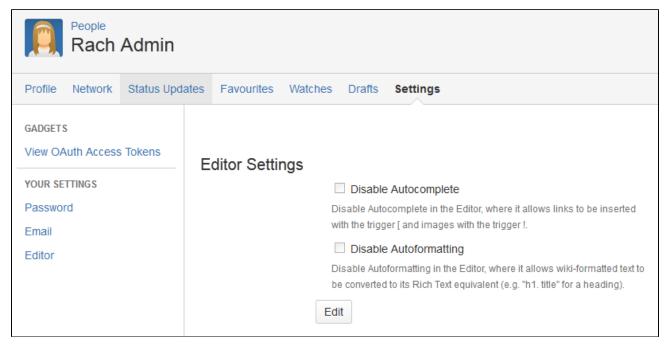
Notes:

- This setting does not affect the keyboard shortcuts for autocomplete (Ctrl+Shift+K and Ctrl+Shift+M). Even if the trigger characters are disabled, you can still use the keyboard shortcuts for autocomplete.
- This setting affects only you. Other people using Confluence can enable or disable the setting on their user profiles independently.
- Note that autocomplete is enabled by default.

To enable or disable the autocomplete trigger characters:

1. Choose **your profile picture** at top right of the screen, then choose **Settings**
2. Choose **Editor** under 'Your Settings' in the left-hand panel.
3. Choose **Edit**.
4. Either:
    - Disable autocompletion by selecting **Disable Autocomplete**.
    - Enable autocompletion by clearing **Disable Autocomplete**.
5. Choose **Submit**.

Screenshot: User settings for the editor

Ignoring autocomplete

You can add macros, links and images by wiki markup alone. Type the macro, including its parameters and the closing curly bracket. Add a link, such as an anchor link, and end it with a square bracket. Insert an image or other embedded object, enclosed between exclamation marks. As soon as you close the macro, link, or embedded image, Confluence will convert it to rich text format and add it to the page.

For more information about mouse-free macros, links and images, choose **Help**

⑦ ▾

> Keyboard Shortcuts from the Confluence header.

Page Layouts, Columns and Sections

The layout of your pages can have a big impact on how they're read, and layouts, used well, allow you to position text, images, macros, charts, and much more, to have the best visual impact.

There are two ways to modify the layout of a Confluence page:

- Use page layouts to add sections and columns
- Use macros to add sections and columns.

Page layouts provide a simple, visual representation of your page layout in the editor, while the macros are more flexible and allow for greater complexity in your layout.
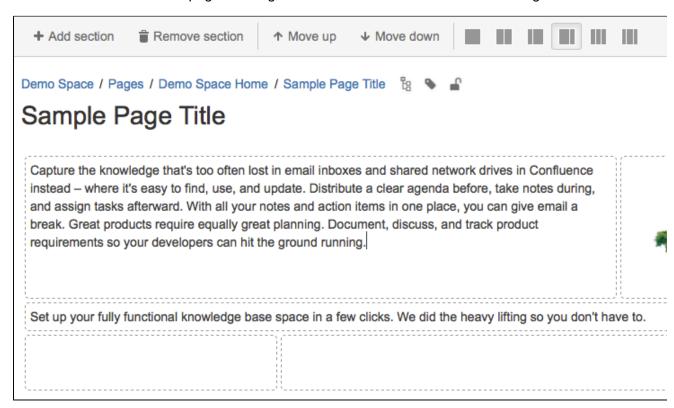
On this page:

- The Section and Column macros

Use page layouts

The page layouts tool allows you to structure your page using horizontal sections and vertical columns. By adding multiple sections with different column configurations you can build quite complex layouts very easily.

Screenshot: Editor view of a page showing three sections with different column configurations.

Start by adding a horizontal section to your page.

To add a section:

1. Choose the **Page Layout**

> ⬚
button in the toolbar

The Page Layout toolbar appears.
2. Choose **Add Section**

The new section appears below your current content, with the boundaries of the section(s) indicated by dotted lines (the dotted lines aren't visible when you view the page).

To change the column layout in a section:

1. Place your cursor in the section you wish to change
2. Choose a layout from the page layout toolbar (for example, two columns or three columns)

Any text, images or macros in your section are not lost when you change the column layout. When you decrease the number of columns, Confluence will move your content to the left. When you increase the number of columns, Confluence will add blank columns to the right of your existing content.

To move a section to another part of the page:

1. Place your cursor in the section you wish to move
2. Choose the **Move up** or **Move down** buttons

The section and all of its content will be moved above or below other sections on the page.

To delete a section:

1. Place your cursor in the section you wish to remove
2. Choose **Remove section**

The section and all of its content will be removed.

Notes about Page Layouts

- **Column width** – The width of the columns are fixed. If you need more than three columns, or columns of a specific width, you should use the Section and Column macros described below.
- **Very wide tables** – The width of each column is set to a percentage of the page width. The icons in the drop-down menu indicate the relative widths for each layout. In most cases, Confluence will adapt the width of the columns to fit the width of the page. If a column includes an item that's too wide for it, you'll see a horizontal scroll bar when viewing the page.

The Section and Column macros

You can use the Section and Column macros to add a set of columns to the page. The Section macro defines an area that will contain the columns. You can have as many sections as you like. Within each section, you can have as many columns as you like.

The Section and Column macros are useful if you want to define a specific percentage or pixel width for each column.
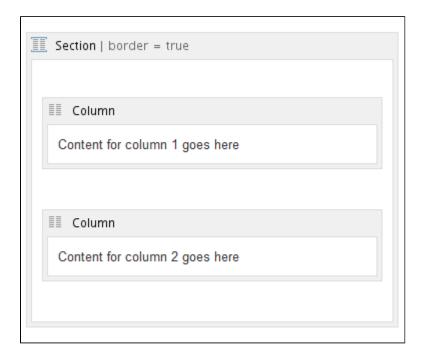
To add a section and some columns to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find the **Section** macro, select it and insert it onto the page
3. Choose **Insert** > **Other Macros** again
4. Find and insert the **Column** macro
5. Add your content to the column

Insert as many columns as you like within the section.

> You should always have at least one column macro within a section macro. Using a section macro without any column macros can negatively affect page loading time.

Screenshot: A section and two columns in the editor

When you see the page in view mode, the above layout is displayed like this:
Content for column 1 goes here
Content for column 2 goes here

Macro parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameters of the Section macro

Parameter	Default	Description
Show Border	false	Select this option to draw a border around the section and columns.
		Note: Without a Column macro , the border will not be displayed correctly.

Parameters of the Column macro

Parameter	Default	Description
Column Width	100% of the page width, divided equally by the number of columns in the section.	Specify the width of the column, in pixels (for example, `400px`) or as a percentage of the available width (for example, `50%`).

All content within your section must be enclosed within a Column macro, otherwise the section layout will not work as expected.

Create Beautiful and Dynamic Pages

Confluence has a number of features that help you build attractive pages to engage your readers and give them the opportunity to interact with up-to-date information. This page summarises those features and provides links to detailed instructions.

Add visual appeal

Pictures, photographs and screenshots. Confluence pages can display images from your Confluence site and from other websites. To put an image into Confluence, you can upload it and attach it to a page or blog post, then display it on any page, blog post or comment. Alternatively, display a remote image using its web address (URL). See Displaying Images.

Galleries. Use the Gallery Macro to display a set of images. When viewing the page, people can click any of the pictures to zoom in and view the images as a slide show.

People. Add a Profile Picture Macro to show a picture of a Confluence user, or a User Profile Macro to show a summary of the person's profile as well as their avatar.
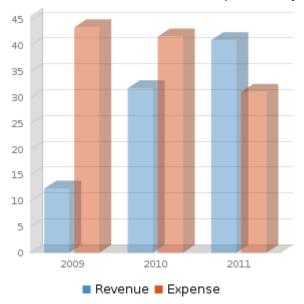
Multimedia. You can display movies, animations and videos, and embed audio files on your Confluence page. For example, Confluence supports Adobe Flash, MP3, MP4, and various other movie formats. See Embedding Multimedia Content.

Social video and image sharing. The Widget macro displays live content from social sites such as YouTube and other video sharing sites, and Flickr for shared photographs. See the guide to the Widget Connector Macro.

Bring numbers to life

The Chart Macro offers a variety of graphs and charts that you can use to illustrate statistics and other numerical data.

Illustration: A 3-dimensional bar chart produced by the Chart macro

Display presentations and documents

Display your Office documents and other presentations directly in Confluence.

Illustration: A PowerPoint slide deck

- Attach your Office documents to a Confluence

page then display them on the page, using the View File Macro. This works for Excel spreadsheets, PowerPoint presentations and Word documents.
- Display PDF files in Confluence too, also with the View File Macro.
- Use the Widget Connector Macro to show slide decks hosted on SlideShare and other online presentation sites.

Pull in content from JIRA applications

Many project teams and customers also use JIRA applications such as JIRA Software or JIRA Service Desk. Rather than copying and pasting issues onto your Confluence page, you can display it directly from the source, thus ensuring that the information shown in Confluence is always up to date.

Link to a feature request in your issue tracker, or display a list of fixed issues – useful for release notes and project planning. See the JIRA Issues Macro.

Tell a story in pictures

A number of Confluence add-ons provide sophisticated tools for creating diagrams and mockups.

For example:

- Balsamiq Mockups for Confluence
- Creately for Confluence
- Gliffy Confluence Plugin
- Graphviz Plugin for Confluence
- Lucidchart for Confluence

Search the Atlassian Marketplace for more add-ons.

Before installing an add-on (also called a plugin) into your Confluence site, please check the add-on's information page to see whether it is supported by Atlassian, by another vendor, or not at all. See our guidelines on add-on support.

Vary the structure of your pages
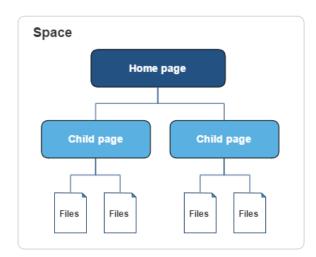
You can build up a custom layout by using the page layout tool to add sections and columns to your page. See the detailed guidelines to Page Layouts, Columns and Sections.

Do you need to display tabular data, which your readers can sort when viewing the page? See Tables.

Use other macros to highlight and format sections of your page:

- Panel
- Info, Tip, Note, and Warning
- Code block
- Noformat

Illustration: A Gliffy diagram

Integrate your content with social media

People share information on various social sites. You can make Confluence a focal point where people collect their shared information and see what is happening in the areas that matter to them.

Illustration: Twitter stream via the Widget macro

Use the Widget Connector macro:

- Show a live stream of tweets from a Twitter user, or tweets matching a Twitter search.
- Display a video from YouTube or other online movie sites.
- Share photographs from Flickr.
- See what else the Widget Connector macro can do.

Show activity streams

Make your Confluence pages dynamic and interactive with:

- An activity stream showing updates and comments on Confluence and other linked applications. See Gadgets.
- An RSS feed from within Confluence or an external site. See Subscribe to RSS Feeds within Confluence.
- A list of recent blog posts from within Confluence. See Blog Posts Macro.

Page Templates

When you add a new page, you don't have to start from scratch. Instead, you can base your new page on a template – a Confluence page with predefined content. Some templates are provided by blueprints or add-ons, and you can even create your own templates.

Some examples of useful templates are:

- The meeting notes template will help you and your team collaborate on notes and follow-up tasks
- The requirements template allows you to capture your software/hardware product requirements, and create related JIRA issues from the page

On this page:

- Global templates and space templates
- Create a template
- Use a template
- Templates provided by blueprints
- Promote templates in the Create dialog
- System templates

Related pages:

- Create a Template
- Administering Site Templates
- Importing Templates
- Create a Page from a Template

Global templates and space templates

In Confluence, there are two categories of page templates:
- **Space templates:** These page templates are available in a specific space only. If you have space administrator permission, you can define templates via the space administration screen.
- **Global templates:** These page templates are available in every space on your site. If you have Confluence Administrator permission, you can define global templates via the Confluence Administration Console.

If you're a system administrator, you can also download predefined templates.

Create a template

You can write your template using the Confluence editor. You can also add special variables to the page, if you want to include fields that the author will complete when adding the page. See Create a Template for more information.

Use a template

Page templates are used only when adding a page. It is not possible to apply a template to an already-existing page. Once a page has been added using a template, the template is no longer linked to the page. All further editing is performed as if the template was never used. Some plugins provide enhanced template functionality. You can search the Atlassian Marketplace for template add-ons. See Create a Page from a Template for more information.
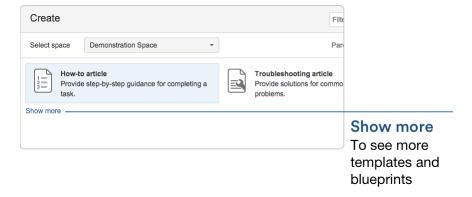
Templates provided by blueprints

A blueprint is a page template with added functionality to help you create, manage and organise content in Confluence, and there's a collection of predefined ones that ship with Confluence. You can also download additional blueprints from the Atlassian Marketplace. You can customise the blueprint templates to suit your individual needs, disable particular blueprints or even develop your own blueprints.

Promote templates in the Create dialog

If you're a space administrator, you can choose to promote specific templates and blueprints in the Create dialog. Promoting items can help ensure consistency in a space by encouraging users to create particular types of content, instead of blank pages.

The promoted templates or blueprints will appear at the top, with all other content types, including Blank Page and Blog Post collapsed under them. To view the other types of content available choose the **Show more** link.

Show more
To see more templates and blueprints

To promote a template or blueprint:

1. Go to **Space Tools** > **Content Tools**
   If your space is using the Documentation theme choose **Browse** > **Space Admin** > **Templates**
2. Choose **Promote** next to the templates or blueprints you want to appear in the Create dialog

Remember, by promoting a blueprint or template you'll be hiding all other items, including blank page and blog post, under the **Show more** link.

If you use the **Show more** link in the create dialog more than three times in a single space, the dialog will show you all templates by default from then on.

System templates

Confluence also provides 'system templates' containing content like the site welcome message and default space content. See Administering Site Templates.

Create a Template
In Confluence, there are two categories of page templates:
  • **Space templates:** These page templates are available in a specific

space only. If you have space administrator permission, you can define templates via the space administration screen.

- **Global templates:** These page templates are available in every space on your site. If you have Confluence Administrator permission, you can define global templates via the Confluence Administration Console.

Add a space template

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar
2. Choose **Templates** > **Create new template**.

If your space uses the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **Templates** from the space administration options.
3. Choose **Create new template**.

Add a global template

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Global Templates and Blueprints** in the left-hand panel
3. Choose **Add New Global Template**

On this page:

- Add a space template
- Add a global template
- The template editor
  - Template variables
  - Labels
  - Images and other attachments
  - Instructional text
- Add a description to your template
- Edit or delete a template
- Notes

Related pages:
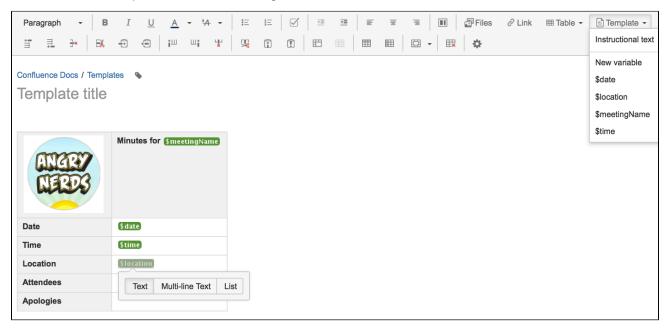
- Create a Page from a Template
- Page Templates
- Add, Remove and Search for Labels
- Macros
- The Editor

Check out the article How to document releases and share release notes to see how helpful creating templates in Confluence can be.

The template editor

When you create or edit a template, you'll be using the editor in much the same way as when you edit a page or blog post. In addition you can add variables, which will produce a form for data collection when anyone adds a page based on the template.

Screenshot: The template editor with an image, table, text, and variables

Screenshot: The form displayed when you create a page based on the template

Confluence Docs / Pages / Confluence Docs Home / Confluence top 50 pages

Page Template Wizard

Step 2: Fill in template variables

The information you supply will become part of the content in your new page.

ANGRY NERDS	**Minutes for** [_____] (meetingName)
Date	[_____] (date)
Time	[_____] (time)
Location	Meeting room 1 ⌄ (location)
Attendees	
Apologies	

Minutes

- *Add your first meeting minute, here.*

[<< Back] [Next >>]

Template variables

When you add variables to your template, they will act as form fields. When you create a page based on a template, you'll see a text entry box for each field. Enter data into each field, and it'll be added to the page.
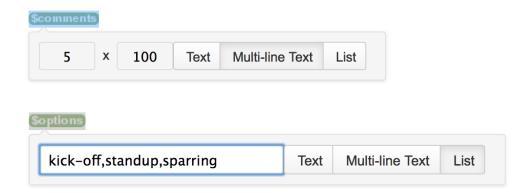
You can add the same variable more than once in the same template, which is useful if you need the same information in more than one place on the page.

To insert a variable into a template:

1. Choose **Template** > **New Variable** from the editor toolbar (or choose an existing variable to add it to the page)
2. Enter a name for the variable
3. Press **Enter** (by default this will create a single-line text input field)

To change the variable type, click the variable placeholder and the variable's property panel will appear. Choose one of the variable types: **Text**, **Multi-line Text**, or **List**.
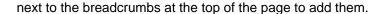
You can change the number of lines and width in characters of a **Multi-line Text** field. If you choose **List**, enter each of the items in your list, separated by commas.

Hint: Type **$** and the variable name, then press **Enter** to add a new variable or to select an existing variable from a list of suggestions. The suggestions dialogue shows variables already defined in this template.

Labels

If you'd like all pages created using this template to have one or more labels, choose the labels icon

🏷️

next to the breadcrumbs at the top of the page to add them.

Images and other attachments

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.

Instructional text

Instructional text is placeholder content in a template, and is only visible while you're editing the page. Use it to give guidance to whoever is creating a page from the template.

To insert instructional text:

1. Choose **Template** > **Instructional Text** in the toolbar
2. Type in your instructional text (for example, *Insert an image of the interface here.*)

Instructional text appears in italics with a shaded background, to distinguish it from normal paragraph text.

You can also change the placeholder type from **Text** to either:

- **User mention** – Opens the user mention dialogue.
- **JIRA Macro** – Opens a dialogue that allows you to create a new JIRA issue, or search for one or more JIRA issues to include on the page.

Add a description to your template

The template description displays in the 'Create' dialog, and is useful for explaining the purpose of your template to other users.

To add a description to a template:

- Go to the space or global templates page (as described above)
- Choose the **Edit** icon in the 'Description' column
- Enter your description and choose **Save**

User Created Templates				Create New Template
Name	**Description**	**Last Updated By**	**Updated**	**Operations**
Meeting Minutes Template	My new meeting minutes template. 🖉	Gil User	Aug 27, 2014	Promote · Edit · Delete

Choose the pencil to edit your template's description

Edit or delete a template

If you need to change anything about your template, or want to delete it, navigate to either your space or global template (as described above) and choose either **Edit** or **Delete**.

Notes

- Page templates are used only when adding a page. It is not possible to apply a template to an already-existing page. Once a page has been added using a template, the template is no longer linked to the page. All further editing is performed as if the template was never used. Some plugins provide enhanced template functionality. You can search the Atlassian Marketplace for template add-ons.
- When you use a Table of Contents macro in a template, you'll see an error when you preview the template, but the Table of Contents macro works on the pages that people create from the template.
- The editor for templates is available only in **Confluence 4.3 and later**. Please refer to the earlier documentation for a description of the wiki markup editor templates.
- Confluence also provides 'system templates' containing content like the site welcome message and default space content. See Administering Site Templates.

Create a Page from a Template

You can create a page based on a global template (available to all spaces) or a space template (available only to that space).

> If you want to quickly create a blank page, hit the **Create** button in the header; if you want to create a page from a template, hit the **Create from template** button.
>
> Create blank page | **Create** | ••• | Create from template

On this page:
- Information copied from the template to the page
- Form fields displayed by the template
- Using a template to create a page
- Notes

Related pages:
- Create a Template
- The Editor
- Add, Remove and Search for Labels

Information copied from the template to the page

When you create a page based on a template, Confluence will copy the following content and information from the template to the new page:

- Labels
- Text and styles
- Layouts and formatting
- Macros
- Embedded images and other files. Note that you cannot attach an image or other file to a template. But if the template displays an image or file from another page, the new page will display that image or file too.

Form fields displayed by the template

If the template author included variables in the template, Confluence will display a form prompting you to supply values for the variables when you add the page.

Using a template to create a page

To create a page based on a template:

1. Choose **Create from template**

   [•••]
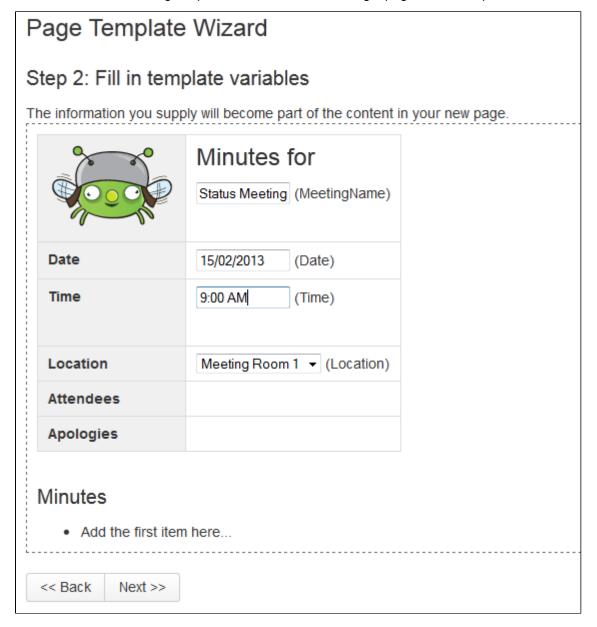
   in the Confluence header
2. Select a **space** and the **template** you want to use and choose **Next**
   If the template contains variables, you'll see a form allowing you to add values for the form variables.
3. Type the relevant information into the form fields, and choose **Next**
   Now you'll see a new page based on the template. If you added information in the form fields, the page content will include that information.
4. Name your page, add content or make any other changes required and hit **Save**

Screenshot: Form showing template variables when creating a page from a template

Page Template Wizard

Step 2: Fill in template variables

The information you supply will become part of the content in your new page.

	Minutes for
	Status Meeting (MeetingName)
Date	15/02/2013 (Date)
Time	9:00 AM (Time)
Location	Meeting Room 1 ▼ (Location)
Attendees	
Apologies	

Minutes

- Add the first item here...

[<< Back] [Next >>]

Notes

Page templates are used only when adding a page. It is not possible to apply a template to an already-existing page. Once a page has been added using a template, the template is no longer linked to the page. All further editing is performed as if the template was never used. Some plugins provide enhanced

template functionality. You can search the Atlassian Marketplace for template add-ons.

Blueprints

What's a blueprint?

A blueprint is a set of page templates with added functionality to help you create, manage and organise content in Confluence more easily.

Create meeting notes, shared file lists and requirements documentation out of the box, and Customise the blueprint templates to suit your individual needs. You can even develop your own blueprints.

Create content using a blueprint

1. Choose **Create from template**

   ![...]

   in the Confluence header
2. Select a blueprint from the create dialog
3. Hit **Create**

The editor will open, and, depending on the blueprint selected, a prompt to enter information or the page will appear. You can now follow the instructions built in to the blueprint to add content.

On this page:
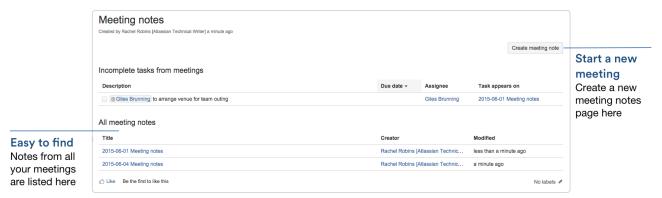
- What's a blueprint?
- Create content using a blueprint
- Customise blueprint templates
- Promote blueprints in the Create dialog
- Add more blueprints
- Disable a blueprint
- Full list of blueprints

Related pages:

- Page Templates
- Request Add-ons

The first time a blueprint is used in a space, Confluence creates an index page and adds a shortcut to your sidebar (if you're using the default theme). The index displays a list of pages made with the blueprint, and information selected information from your blueprint pages. For example, the meeting notes index displays a list of all meeting notes pages in the space, who created them, and when they were last modified.

Screenshot: Index page for the Meeting Notes blueprint

If you're using the Documentation theme, your blueprint index page(s) will appear in the sidebar as children of the homepage.

Customise blueprint templates

Blueprints are made up of templates that can be customised for an individual space or the whole site. This means you can adapt the content of the blueprint pages to suit your specific needs. For example, you might update the Meeting Notes blueprint templates to include a heading for apologies.

If you have space administrator permissions, you can customise blueprint templates for the spaces you are an administrator of. You must be a Confluence Administrator to customise blueprint templates for a whole

site. See Administering Site Templates for more information.

To customise a blueprint template for a space:

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar
2. Choose **Edit** beside the blueprint template you wish to edit
3. Make your changes to the template and choose **Save**

Editing a blueprint template is very similar to editing a page template, except:

- Be careful not to remove any macros that the blueprint page or index page may use to store and display information
- You can't remove a blueprint template or change the template name

To reset a blueprint template back to the default:

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar
2. Choose **Reset to default** beside the blueprint template you wish to reset
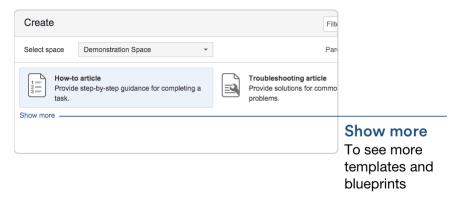
See Working With Templates and Administering Site Templates for more information on templates.

As with user created space and site templates, editing a blueprint template will not change existing pages, but any new blueprint pages will be based on the updated template.

Promote blueprints in the Create dialog

If you're a space administrator, you can choose to promote specific templates and blueprints in the Create dialog. Promoting items can help ensure consistency in a space by encouraging users to create particular types of content, instead of blank pages.

The promoted templates or blueprints will appear at the top, with all other content types, including Blank Page and Blog Post collapsed under them. To view the other types of content available choose the **Show more** link.

Show more
To see more templates and blueprints

To promote a template or blueprint:

1. Go to **Space Tools** > **Content Tools**
   If your space is using the Documentation theme choose **Browse** > **Space Admin** > **Templates**
2. Choose **Promote** next to the templates or blueprints you want to appear in the Create dialog

Remember, by promoting a blueprint or template you'll be hiding all other items, including blank page and blog post, under the **Show more** link.

If you use the **Show more** link in the create dialog more than three times in a single space, the dialog will show you all templates by default from then on.

Add more blueprints

You can find more blueprints for Confluence in the Atlassian Marketplace. Blueprints are managed using add-ons (also known as plugins).

See Request Add-ons for information on how you can search for new blueprint add-ons and send a request to your System Administrator.

If you are a System Administrator, see Managing Add-ons or Plugins for information on how to install new blueprint add-ons.

You can also develop your own blueprints. See our developer documentation on Writing a Blueprint.

Disable a blueprint

You may want to disable particular blueprints. For example, you may not want to see the Product Requirements blueprint in the create dialog in an HR or Social space. If you are a Confluence Administrator you can also disable particular page and space blueprints for the whole site.

To disable a blueprint in a space:

- Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar
- Choose **Disable** next to the blueprint you wish to disable in that space

You can re-enable the blueprint at any time.

To disable a blueprint across a whole site:

- Choose the **cog icon**

  , then choose **General Configuration** under Confluence Administration (You need Confluence Administrator permissions to do this)
- Choose **Global Templates and Blueprints**
- Choose **Disable** next to the page or space blueprint you wish to disable

The blueprint will not appear in the 'Create' or 'Create Space' dialogs.

Full list of blueprints
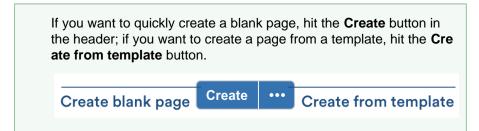
Here's the full list of blueprints bundled with Confluence.

Page blueprints	Space blueprints
- Meeting notes - File list - Decision - How-to article - Troubleshooting article - JIRA report - Product requirements - Retrospective - Share a link - Task report	- Documentation space - Team space - Knowledge base space

Decisions Blueprint

The Decisions blueprint helps you make decisions and record the outcomes with your team.

The first time you use the Decisions blueprint in a space, Confluence will create an index page and add a shortcut on your space sidebar (if you're using the default theme). The index acts as your Decision Register and lists all the decisions in that space.

> If you want to quickly create a blank page, hit the **Create** button in the header; if you want to create a page from a template, hit the **Create from template** button.
>
>

Related pages:

- Blueprints
- File List Blueprint
- Meeting Notes Blueprint
- Product Requirements Blueprint

To create a decision page:

1. Choose **Create from template**
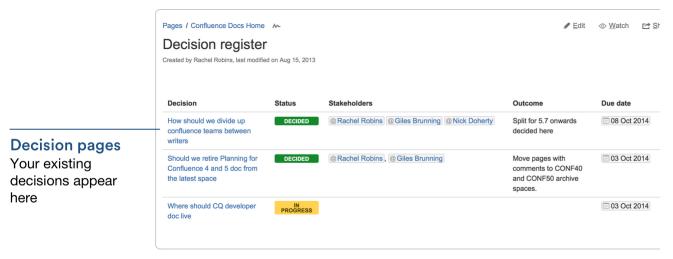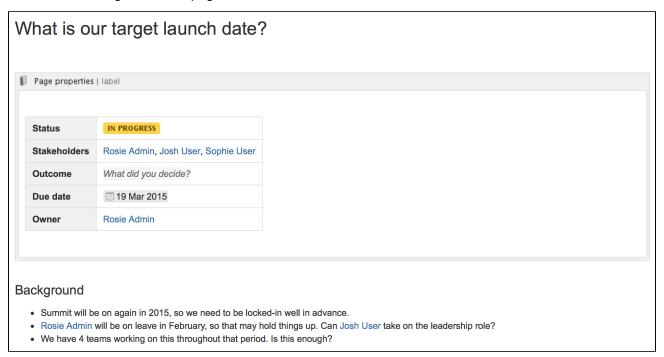
   in the Confluence header
2. Select **Decision** and hit **Next**
3. Enter information about the decision and relevant stakeholders (the blueprint will prompt you) and hit **Create**

Once you save your first decision page, Confluence will create a decision register page for the space you're in, and add a shortcut to it in the space's sidebar.

Screenshot: Decision Register showing a series of Decision pages

Screenshot: Editing a Decision page

The Decisions blueprint uses these Confluence features:

* Page Properties and Page Properties Report macro - content that you enter within the page properties macro can appear on the index page.

- Mentions - add a user as a stakeholder, owner or @mention them on the page and they will be notified in their workbox.

> For an example of the Decisions Blueprint, and some other great page elements, check out: How to make better decisions as a development team.

Customising this blueprint

You can customise the templates that are used by the Decisions blueprint - see Customising the blueprint templates. For example, you might choose to edit the decisions **index page** in a space to change the columns displayed by the Page Properties Report macro.
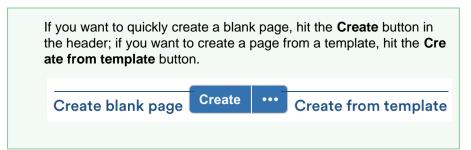
You can also edit the **page template** to add headings or instructional text to the background section, or even add rows to the Page Properties macro. For example, a row for the date the decision was made.

See Instructional text to find out more about using instructional text in templates.

File List Blueprint

The File List blueprint helps you to create lists of files to share with your team. Great for organising documents, images and presentations.

The first time you use the File List blueprint in a space, Confluence will create an index page and add a shortcut to your space sidebar (if you're using the default theme). The index page lists the latest File List pages in that space. You can have as many File List pages as you need.

Related pages:

- Blueprints
- Meeting Notes Blueprint
- Product Requirements Blueprint

> If you want to quickly create a blank page, hit the **Create** button in the header; if you want to create a page from a template, hit the **Create from template** button.

To create a file list:

1. Choose **Create from template**
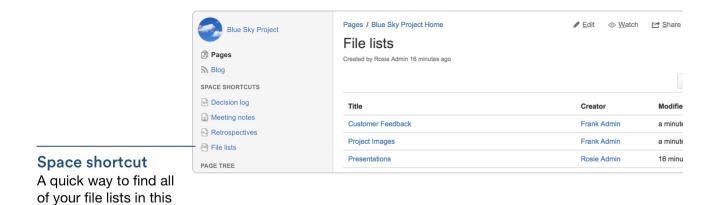
   in the Confluence header
2. Select **File List** and hit **Next**
3. Enter the details for your file list and hit **Create**
4. Drag files from your desktop or choose **browse for files** to search for files on your computer

Attachments appear on the page, and you can expand each attachment to preview the file and/or view its details.
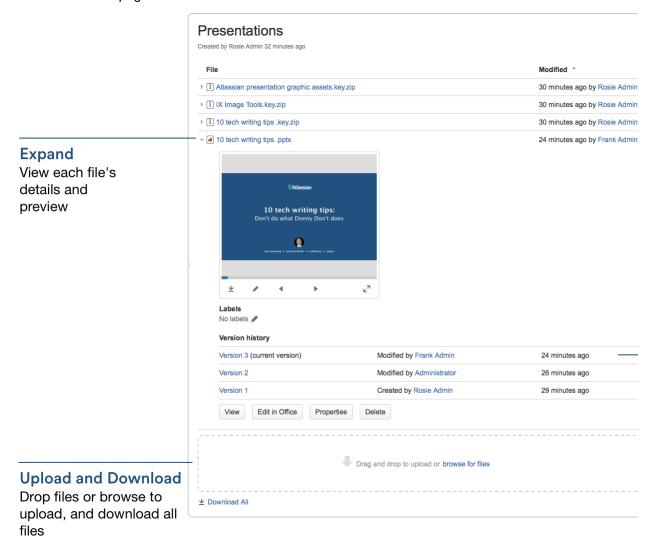
In this example, three file list pages have been created to store project related presentations, images and customer feedback. Confluence looks after the versioning of the files, so there's no need to use the document file name to mark version numbers.

Screenshot: Index page showing File List pages

Space shortcut

A quick way to find all of your file lists in this space

Screenshot: A File List page

Expand

View each file's details and preview

Upload and Download

Drop files or browse to upload, and download all files

Customising this blueprint

You can customise the templates that are used by the File List blueprint - see Customising blueprint templates.

The File List blueprint template uses the attachments macro. You can customise the macro to change the sort order or hide features such as version history and the upload attachment fields.

You can also edit the Content Report Table macro used on the Index page to specify the number of pages you want to display.

Meeting Notes Blueprint

The Meeting Notes blueprint helps you to plan your meetings and share notes and actions with your team.

The first time you use the Meeting Notes blueprint in a space, Confluence will create an index page and add a shortcut on your space sidebar (if you are using the default theme). The index page lists the latest Meeting Notes pages in that space.

Related pages:

- Blueprints
- File List Blueprint
- Product Requirements Blueprint

If you want to quickly create a blank page, hit the **Create** button in the header; if you want to create a page from a template, hit the **Create from template** button.

Create blank page **Create** ••• Create from template

To create a meeting notes page:

1. Choose **Create from template**

   •••
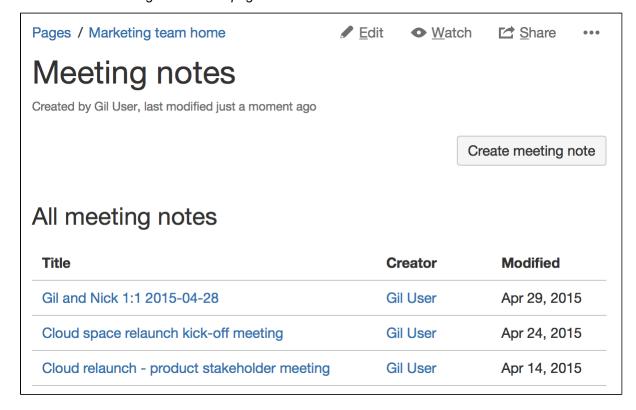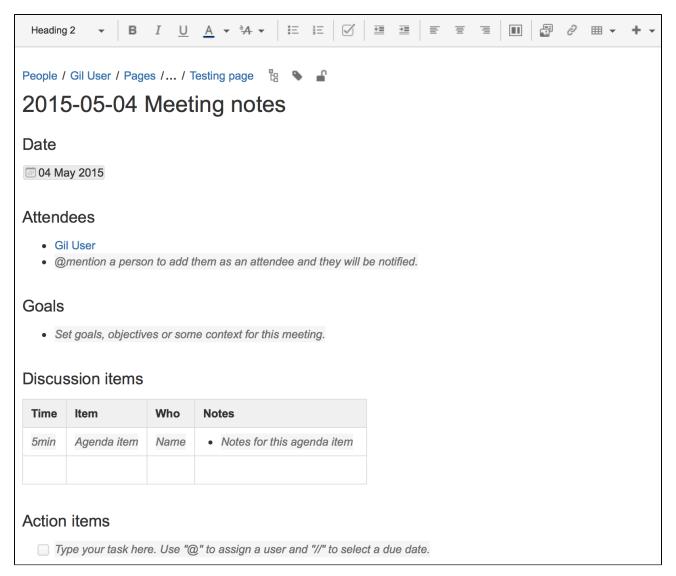
   in the Confluence header
2. Select **Meeting Notes** and hit **Next**
3. Enter the information required by the template and hit **Create**
4. **Save** your page and get ready to attend your meeting

You can edit the page during or after your meeting, and enter your notes, action items and @mention users to assign tasks to them.

Screenshot: A Meeting Notes index page

Pages / Marketing team home   ✎ Edit   👁 Watch   ⬈ Share   •••

Meeting notes

Created by Gil User, last modified just a moment ago

Create meeting note

All meeting notes

Title	Creator	Modified
Gil and Nick 1:1 2015-04-28	Gil User	Apr 29, 2015
Cloud space relaunch kick-off meeting	Gil User	Apr 24, 2015
Cloud relaunch - product stakeholder meeting	Gil User	Apr 14, 2015

Screenshot: A blank Meeting Notes page showing instructional text

---

| Heading 2 ▾ | **B** | *I* | U̲ | A ▾ | ³⁄A ▾ | ☰ | ≣ | ☑ | ⇥ | ⇤ | ≡ | ≡ | ≡ | ▣ | ⧉ | ⊘ | ⊞ ▾ | + ▾ |

People / Gil User / Pages /... / Testing page 🔖 🏷 🔓

2015-05-04 Meeting notes

Date

📅 04 May 2015

Attendees

- Gil User
- *@mention a person to add them as an attendee and they will be notified.*

Goals

- *Set goals, objectives or some context for this meeting.*

Discussion items

Time	Item	Who	Notes
5min	*Agenda item*	*Name*	• *Notes for this agenda item*

Action items

- ☐ *Type your task here. Use "@" to assign a user and "//" to select a due date.*

The Meeting Notes blueprint uses some cool Confluence features:

- Instructional text - this handy text prompts you to enter information and disappears when you start typing or view the page.
- Mentions - @mention a user on the page and they will be notified in their workbox.
- Task lists - @mention a user in a task to assign it to them – the task will appear as a personal task in their workbox. You can also add a due date by typing **//**, then choosing a date from the calendar.

Customising this blueprint

You can customise the templates that are used by the Meeting Notes blueprint – see Customising the blueprint templates.

You might choose to edit the headings or add additional headings, or change the instructional text that prompts users to enter information to suit your context. To find out more about using instructional text in a template, see Instructional text.

You can also edit the Content Report Table macro used on the Index page to specify the number of pages you want to display.

Product Requirements Blueprint

The Product Requirements blueprint helps you to define, scope and track requirements for your product or feature.

The first time you use the Product Requirements blueprint in a space, Confluence will create an index page and add a shortcut on your space sidebar (shortcut only available in the default theme). The index lists all the Product Requirements pages in that space, and displays a summary of the

information on each page (such as status and owner). You can have as many Product Requirements pages as you need.

Related pages:

- Blueprints
- File List Blueprint
- Meeting Notes Blueprint

If you want to quickly create a blank page, hit the **Create** button in the header; if you want to create a page from a template, hit the **Create from template** button.

Create blank page | **Create** ••• | Create from template

To create a requirements page:

1. Choose **Create from template**

   •••
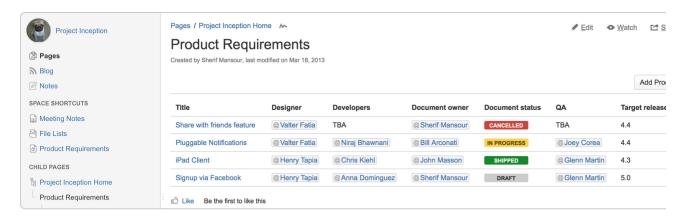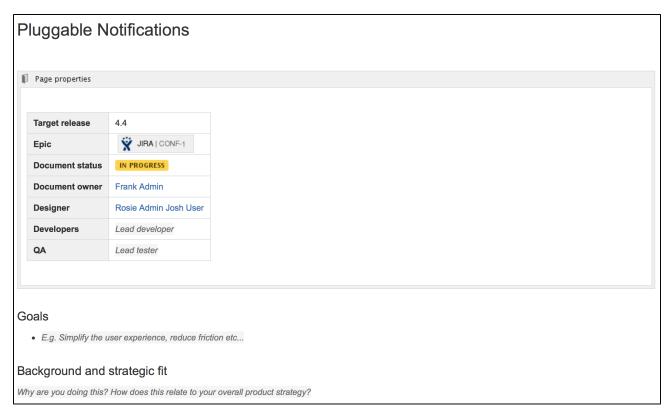
   in the Confluence header
2. Select **Product Requirements** and hit **Next**
3. Enter information about your product or feature (the instructional text will prompt you) and hit **Create**

You can @mention team members to bring them into the conversation about the page.

Screenshot: Index showing a series of Product Requirements pages and summary information

Project Inception

- Pages
- Blog
- Notes

SPACE SHORTCUTS
- Meeting Notes
- File Lists
- Product Requirements

CHILD PAGES
- Project Inception Home
  - Product Requirements

Pages / Project Inception Home

Product Requirements

Created by Sherif Mansour, last modified on Mar 18, 2013

✏ Edit   👁 Watch   ↗ S

Add Pro

Title	Designer	Developers	Document owner	Document status	QA	Target release
Share with friends feature	@ Valter Fatia	TBA	@ Sherif Mansour	CANCELLED	TBA	4.4
Pluggable Notifications	@ Valter Fatia	@ Niraj Bhawnani	@ Bill Arconati	IN PROGRESS	@ Joey Corea	4.4
iPad Client	@ Henry Tapia	@ Chris Kiehl	@ John Masson	SHIPPED	@ Glenn Martin	4.3
Signup via Facebook	@ Henry Tapia	@ Anna Dominguez	@ Sherif Mansour	DRAFT	@ Glenn Martin	5.0

👍 Like   Be the first to like this

Screenshot: Editing a Product Requirements page

Pluggable Notifications

📄 Page properties

Target release	4.4
Epic	🦅 JIRA	CONF-1
Document status	**IN PROGRESS**
Document owner	Frank Admin
Designer	Rosie Admin Josh User
Developers	*Lead developer*
QA	*Lead tester*

Goals

- *E.g. Simplify the user experience, reduce friction etc...*

Background and strategic fit

Why are you doing this? How does this relate to your overall product strategy?

The Product Requirements blueprint uses these Confluence features:

- Page Properties and Page Properties Report macro - content that you enter within the page properties macro can appear on the index page.
- Instructional text - this handy text prompts you to enter information or create a JIRA issue and disappears when you start typing or view the page.
- Mentions - @mention a user on the page and they will be notified in their workbox.

Customising this blueprint

As no two products or projects are alike, you can customise the templates that are used by the Product Requirements blueprint - see Customising the blueprint templates.

You might choose to edit the **index** page in a space to change the columns to be displayed by the Page Properties Report macro.

You might choose to edit the **page template** to:

- edit the headings or add additional headings
- change the instructional text that prompts users to enter information to suit your context
- add or remove rows within the Page Properties macro.

See Instructional text to find out more about using instructional text in templates.

Shared Links Blueprint

The Shared Links blueprint helps you take content from the web and share it with your team. You can use Shared Links to share and collaborate on web content, or to create a centralised repository of useful links.

The first time you use the Shared Links blueprint in a space, Confluence will create an index page and add a shortcut on your space sidebar (if you're using the default theme). The index lists all the shared links in that space.

> If you want to quickly create a blank page, hit the **Create** button in the header; if you want to create a page from a template, hit the **Create from template** button.

To create a shared links page:

1. Choose **Create from template**

   in the Confluence header
2. Select **Share a link** and hit **Next**
3. Enter the URL of the web content you want to share, then hit **Create**

You can also:

- Include topics to help categorise your links – these are added as labels to your page.
- Share the link immediately with another user or group – users will receive a notification.
- Add a comment to start the discussion.

> To make sharing links even faster, you can add a **Share on Confluence** button to your browser's toolbar. Click this button and the webpage you're currently viewing will be added as a shared link!

To add the Share on Confluence button to your browser:

1. Choose **Create from template**
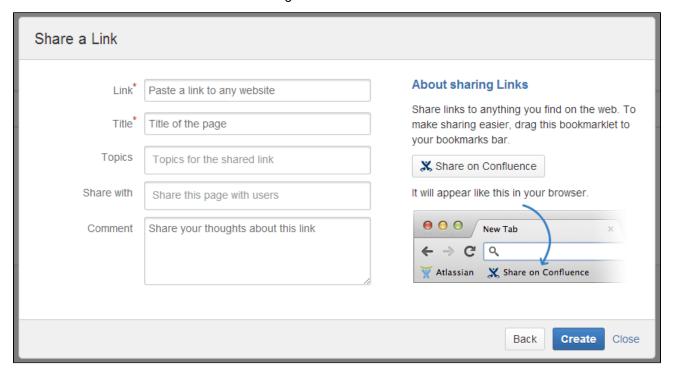
   , then select **Share a link**
2. Drag the Share on Confluence button to your browser toolbar

Now, when you want to share a link in Confluence, choose the **Share on Confluence** button in your browser and follow the prompts.
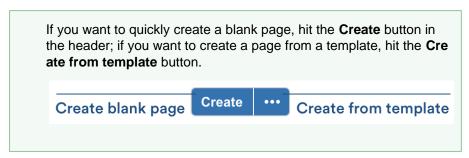
Screenshot: Share a link from the Create dialog.

JIRA Report Blueprint

The JIRA Report blueprint helps you create easy to read reports to communicate the progress of your JIRA Software projects and releases. You can choose from a Change Log report that generates a list of JIRA issues or a Status Report that includes charts to visually communicate your progress.

The first time you use the JIRA Reports blueprint in a space, Confluence will create an index page and add a shortcut on your space sidebar (if you're using the default theme).

> If you want to quickly create a blank page, hit the **Create** button in the header; if you want to create a page from a template, hit the **Create from template** button.
>
> Create blank page | **Create** | ••• | Create from template

To use the JIRA Report Blueprint your Confluence and JIRA application (such as JIRA Software) must be connected via Application Links .

Create a Change Log

The Change Log report displays a list of issues from your JIRA application. This list can be static or dynamic, automatically updating as the status of your issues change in JIRA.

To create a static change Log:

1. Choose **Create from template**
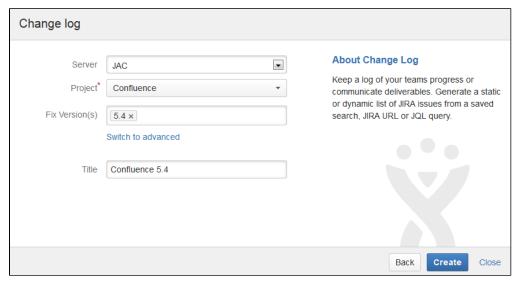
   •••

   in the Confluence header
2. Select **JIRA report** and hit **Next**
3. Select a **Change log** and hit **Next**
4. Enter the information required for the change log and hit **Create**

A report page will be created with sample text and a list of all issues for the project and fix versions selected, organised by issue type. This list of issues is static; it won't be updated when the issues are updated, and is visible to users who don't have JIRA access or permissions to view that project.

Screenshot: Creating a Change Log in simple mode.

Screenshot: Static list of JIRA Issues displaying in the Change Log.

To create a dynamic change log:

1. Choose **Create from template**

   in the Confluence header
2. Select **JIRA report** and hit **Next**
3. Select a **Change log** and hit **Next**
4. Choose **Switch to advanced**
5. Enter a JQL query or paste in the URL of a JIRA search (find out about using JQL in the JIRA Documentation)
6. Hit **Create**

A report page will be created with sample text and a JIRA issues macro that's configured to show your issues. The macro is dynamic and will update when the issues are updated. For more information on changing the information displayed, refer to the JIRA Issues macro.

Screenshot: Dynamic list of JIRA Issues displaying in the Change Log.

Create a Status Report

The Status Report displays the progress of a JIRA project and fix version in pie charts by status, priority, component and issue type. The Status Report uses the JIRA Chart macro, and is dynamic.

To create a status report:

1. Choose **Create from template**
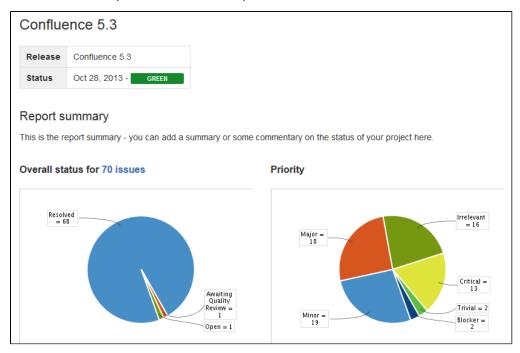
   in the Confluence header
2. Select **JIRA report** and hit **Next**
3. Select a **Status report** and hit **Next**
4. Enter the information required for the report and hit **Create**

A report page will be created with sample text and a series of pie charts, using the JIRA Chart macro. The macro is dynamic and will update when the issues in JIRA are updated. For more information refer to the JIRA Chart macro.

As with the Change Log, you can switch to Advanced mode and use JQL or paste in a JIRA URL to search for issues to display in the report.

Screenshot: Excerpt from the Status Report.

Customising this blueprint

You can customise the templates used by this blueprint. The Change Log uses the **Snapshot JIRA Report Template** (for static list of issues) and the **Dynamic JIRA Report Template**, and the Status Report uses the **Status Report Template**. See Customising the blueprint templates. Variables represent the JIRA Issues and JIRA Chart Macros. While these can't be edited, they can be moved around the page or deleted if you don't want every chart to be included.

You can also choose to edit the **page template** to modify the format of the page, change some headings, or modify the instructional text. To See Instructional text to find out more about using instructional text in templates.

Retrospective Blueprint

Retrospective pages help you track team successes and opportunities after projects or at the end of a sprint. Use this blueprint to document what went well, what needed improvement, and assign actions for the future.

The first time you create a retrospective page in a space, Confluence will automatically create an 'index' page, which will list all retrospectives in the space, and add a shortcut to it in the space sidebar.

Related pages:

- Blueprints
- File List Blueprint
- Meeting Notes Blueprint

If you want to quickly create a blank page, hit the **Create** button in the header; if you want to create a page from a template, hit the **Create from template** button.
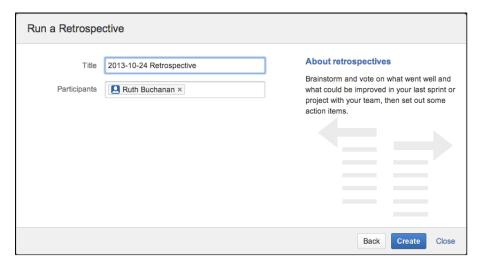
To create a retrospective page:

1. Choose **Create from template**

in the Confluence header

2. Select **Retrospective** and hit **Next**
3. Add participants, change the title if you want to and click **Create**

Screenshot: The 'retrospective' template page

The Retrospective blueprint uses the following Confluence features:

- Page Properties and the Page Properties Report macro make content listed within the macro visible on the index page.
- Instructional text prompts you to enter information and disappears when you start typing or view the page.
- Mention a user on the page to notify them in their workbox.

> Check out how the retrospectives blueprint can be used in the article Create sprint retrospective and demo pages (like a BOSS).

Customise this blueprint

Every team conducts retrospective meetings differently, so you can customise the Retrospective blueprint template to match your team's culture and practices. You can:

- Edit headings and pre-populated text
- Add instructional text to capture specific information
- Add additional sections and content

See Customising blueprint templates for instructions.

Create a Blueprint-Style Report

Using a combination of templates and macros you can make a wide range of reports for managing anything from customer interviews, product requirements to IT service catalogues and more. In this tutorial we'll guide you through the process of creating a blueprint-style report.

In this example, we'll create a multi-team status report. Here's the scenario we'll use for this tutorial.

> The Design, Development and QA teams working on the Blue Sky Project need to produce a short status update page each week, containing the focus area for the week, contact person, risks and overall status for each team. They like the way the Product Requirements blueprint works and want to be able to manage their status updates in a similar way.
>
> What do each of the players want out of this report?
>
> - **Project Lead** – Wants an at-a-glance report that shows only the status for each team.
> - **Team Leads** – Want a summary report, including the focus areas and risk, just for their team.
> - **All team members** – Want it to be easy to create the new page each week.
> - **Management Team** – Want to see all the details for a week on one page, and don't want to have to look at a different page for each team.

With this scenario in mind, this tutorial will guide you through how to:

1. Create a status update template containing a separate page properties macro for each team's section of the report.
2. Create a high level status report, showing just the status of all teams.
3. Create a summary report for each team.
4. Create your first status update page.

> You'll need Space Administrator permissions to complete some of the steps in this tutorial.

Part 1: Create a status update template
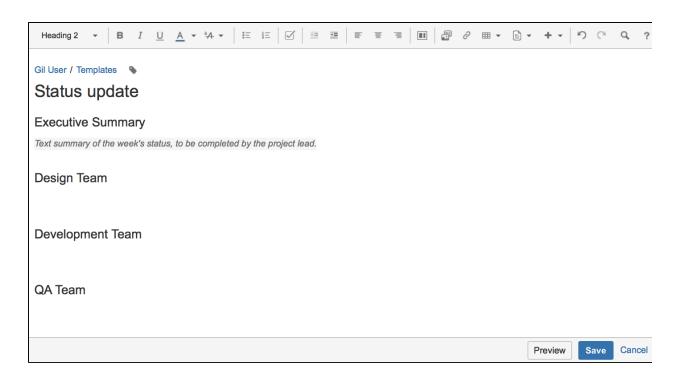
First we'll create a page template and add the Page Properties macros.

1. Go to **Space Admin** > **Content Tools** > **Templates**
2. Choose **Create Template**
3. Give the template a name (in this example the template will be called 'Status Update')
4. Add the skeleton of your status report to the page
5. Choose the label icon

   🏷

   at the top of the page to add a label to the template (in this example, we'll add the label: 'status-update')
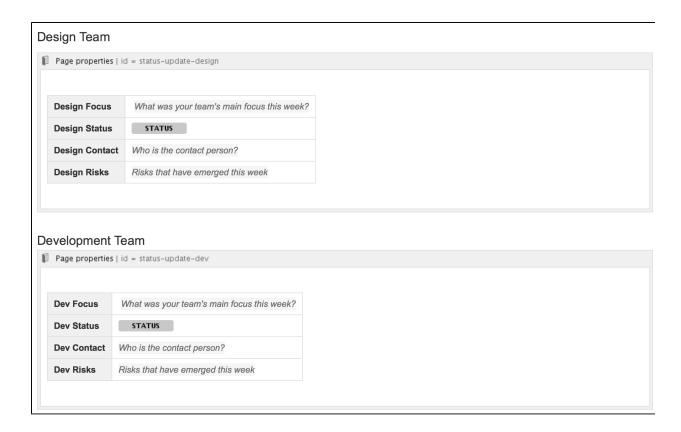
   Screenshot: Adding teams to our status update template

Now we'll add a Page Properties macro to record the status of the Design team.

6. Choose **Insert** > **Other Macros** > **Page Properties** to add the Page Properties macro to the page
7. In the macro body create a two column table and remove the heading row
8. In the left column enter the column headings for your report (these are known as metadata 'keys')
   In this example we'll add 'Design Focus', 'Design Status', 'Design Contact' and 'Design Risks').
9. In the right column, leave the cells blank, or enter some instructional text to prompt your users (Choose **Template** > **Instructional Text**)
   We've also added a status macro.
10. **Edit** the Page Properties macro and enter a **Page Properties ID** for this macro (in this example we'll use 'status-update-design'. This will allow us to report on the status of just the Design team later on)
    Repeat this process for the Development and QA teams, remembering to specify a different ID for each macro (we used 'status-update-dev' and 'status-update-qa').
11. Finally, add any other headings, instructional text or content to your template and **Save**.
    You can enter a **Description** for your template - this appears in the Create dialog.

Screenshot: Our status update template
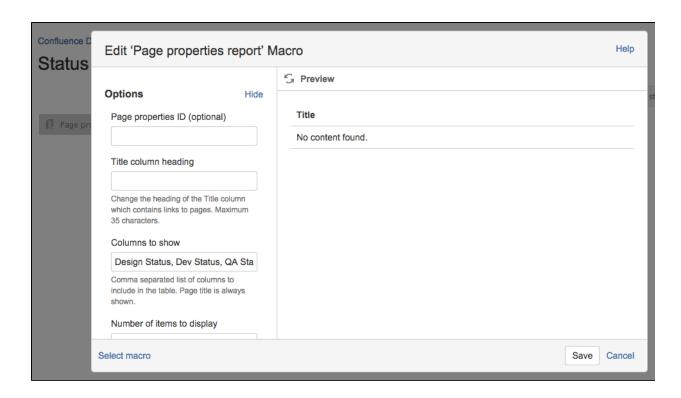
Part 2: Create a report showing the high level status of each team

Next we'll create an index page, just like you see in many blueprints.

1. In your space create a new blank page (this will be our 'Status Report - all teams' page, showing just the status of each team)
2. Choose **Insert** > **Other Macros** > **Page Properties Report** to add the Page Properties Report macro to the page
3. Enter the **Label** to report on (in this example, it'll be the 'status-update' label we added to the template page)
4. Leave the **Display options** > **Page Properties ID** field blank (we want to report on all the macros on the page)
5. In the **Columns to Show** field, list the 'keys' from each macro that you want to include in the report (in this example, we only want to show the values of 'Design Status', 'Dev Status', 'QA Status')
6. Choose **Save** to add the macro to the page

   Screenshot: The page properties report macro on the 'Status Report - all teams' page

Now we'll add a button to the page to allow team leads to easily create new status update pages from the template we created earlier.
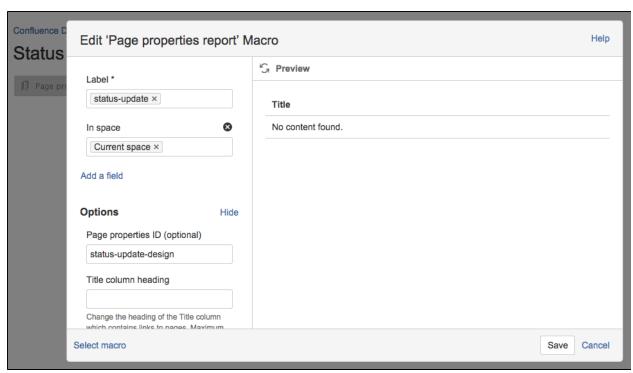
7. Choose **Insert** > **Other Macros** > **Create from Template** to add the Create from Template macro to the page
8. Enter the text for the button (in this example we'll call the button 'New Status Update Page')
9. Select the template from the **Template Name** drop down (in this example our template was called 'Status Update')
10. Specify the title of any pages to be created (This is a great way to keep your titles consistent. In this example we'll call the page 'Status update week ending @currentDate', which will append the current date when the page is created, as in the meeting notes blueprint)
11. Choose **Insert**
12. Add any other content, links or images to the page and **Save**
13. Choose **Space Tools** > **Configure Sidebar** > **Add Link** to add a shortcut to the page on the sidebar

Part 3: Create a separate report for each team

Now we'll create some index pages that show a more detailed summary for each team, starting with the Design team.

1. Create a new blank page – this will be the 'Design Status Report' index page, showing just information for that team.
2. Choose **Insert** > **Other Macros** > **Page Properties Report** to add the Page Properties Report macro to the page
3. Enter the **Label** (the page label is once again 'status-update', the label we added to the template)
4. Expand the **Display options** and enter the **Page Properties ID** that was specified in the Page Properties macro in the template (in this example it was 'status-update-design') – this allows us to report on just information in that macro.
5. Leave all of the other fields blank (we want to show all columns from this Page Properties macro)
6. Choose **Save** to add the macro to the page

Screenshot: The page properties report macro on the 'Design Status Report' index page

7. Add any other content, links or images to the page and **Save**
8. Choose **Space Tools** > **Configure Sidebar** > **Add Link** to add a shortcut to the page on the sidebar
9. Create a new page and repeat this process for each team
   Remember to specify a different Page Properties ID each time (in this example 'status-report-dev' and 'status-report-qa').

If your Design, Dev and QA teams have their own team spaces, this summary report could even be created in their team spaces. Just be sure to specify the space where the Status Updates pages are created in the **Restrict to spaces** field, to make sure the macro can find the pages to report on.

Part 4: Create your first status update page
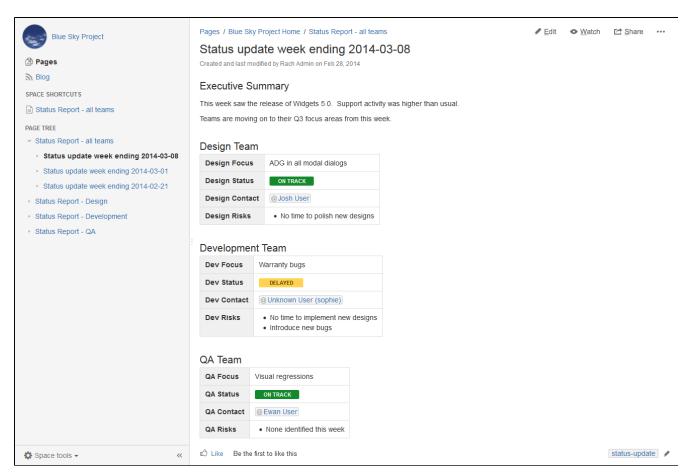
That's it! **Create from template**
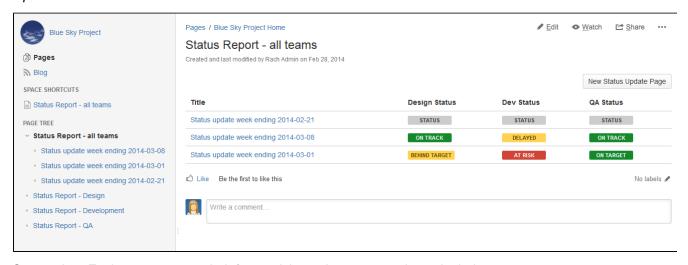
in the Confluence header, then select **Status Update**, or use the **Create a new status update** button to make your first status update page. Just like a blueprint, but 100% made by you.

Here's how our finished pages look.
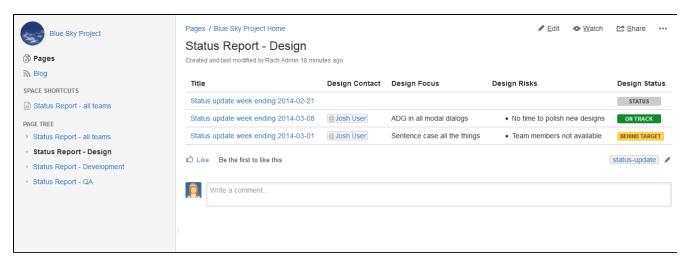
Screenshot: Team Leads and the management team still have a single page for the weekly status update

Screenshot: The Project Lead can see the status of each team, each week, at a glance in the All Teams status report.

Screenshot: Each team can see their focus, risks and status at a glance in their status report.

Title	Design Contact	Design Focus	Design Risks	Design Status
Status update week ending 2014-02-21				STATUS
Status update week ending 2014-03-08	@Josh User	ADG in all modal dialogs	• No time to polish new designs	ON TRACK
Status update week ending 2014-03-01	@Josh User	Sentence case all the things	• Team members not available	BEHIND TARGET

Remember, these concepts don't just apply to status updates – you can use them for any purpose at all.

Import Content Into Confluence

There are a number of ways you can get existing content, such as text, images and other content into Confluence.

On this page:
- Import content from other Confluence sites
- Import content from a Microsoft Word document
- Import web content
- Import other content

Import content from other Confluence sites

To import content from another Confluence site you can:

- Import a backup of the entire Confluence site
- Import an XML export of an individual space. Page history, attachments, and page content will be preserved.

See Restoring a Site and Restoring a Space for more information.

Import content from a Microsoft Word document

The Office Connector allows you to create pages by importing Word documents. The document content is copied onto one or more Confluence pages. See Import a Word Document into Confluence.

Import web content

To displaying web content on a page:

- Use the Widget Connector Macro to display videos, slide shows, twitter chats, documents and more, sourced from other web sites and displayed on your Confluence page.

- Embed an external web page into Confluence with the HTML Include macro.
- Use HTML code in a page with the HTML macro.

Import other content

Importing non-wiki markup into Confluence requires a conversion process:

- Text with basic formatting can be pasted directly into the editor. This includes simple Word documents or web pages.
- Confluence pages saved to disk can be imported from disk.
- Files can be uploaded in bulk using the Confluence WebDav Plugin
- Full featured customisation is available using the Confluence APIs.

Import a Word Document into Confluence

The Office Connector allows you to import Word documents and create one or more Confluence pages from the content.

You can create a single page, or divide the contents up into multiple pages, based on the headings in your document.

This is useful if you have a lot of content stored in existing documents, or if you are migrating from another system or platform that allows you to export to Word format.

⚠ The Office Connector is only available for some browsers, operating systems and applications, so you'll need to check the Office Connector Prerequisites before you start.

On this page:
- Import a Word document
- Import options

Related pages:

- Export Content to Word, PDF, HTML and XML

Import a Word document

1. Create a page in Confluence or go to an existing page (you want to view the page, not edit it).
2. Choose
   •••
   > **Import Word Document**
3. Choose **Browse** and locate the Word document you want to import then choose **Next**.
   The import document options appear.
4. Enter a **title** for the new page (useful if you do not want to use the file name as your page title).
5. Choose **where** you want to import the file (as a brand new page, or overwriting an existing page with the same title).
6. Choose how to handle **title conflicts** (rename the new pages or replace existing pages).
7. Choose whether to create a single page or multiple pages based on the heading styles in the file (this option is only available if the file contains heading styles).
8. Click **Import**.

When the upload has finished, pages will be created with the content of the Word documents. You can then view and edit this page as normal. There's no connection between the original Word document and this page.

Import options

There are a number of options when importing a Word document that control how pages are created, whether the import should overwrite existing pages in the space, and how it should handle page name conflicts.

Option	Description
Root page title	This is the title of the page that will be created or updated by the import.
Where to import	Controls whether the document is imported into the current page (the page you were viewing when you selected Tools > Import) or created as a new page. Choose from • **Import as a new page in the current space** - a new page will be created as a child of the space home page. • **Replace <page name>** - content will be imported into the current page. The title of this page will change to the title you specified in the Root page title field. • **Delete existing children of <pagename>** - any existing children of the current page will be removed when the content of the page is replaced.

Title conflicts	Controls how page name conflicts (a page with the same title already exists in the space) are handled. • **Rename imported pages if page name already exists** - new pages get a new name (a number added to the end of the page title). Existing pages will be unchanged. • **Replace existing pages with imported pages of the same title** - overwrite the content of existing pages. The change will be shown in the Page History for the page. • **Remove existing pages with the same title as imported pages** - remove original pages and then create new pages. The change is not shown in the Page History for the page.
Split by heading	If the document contains Word heading styles you can choose to create multiple pages based on the heading. Options are: • **Don't split** - creates a single page. • **Level Headings** - creates multiple pages in a hierarchy based on the heading levels in the document. A preview of the pages that will be created appears under **Document Outline**.
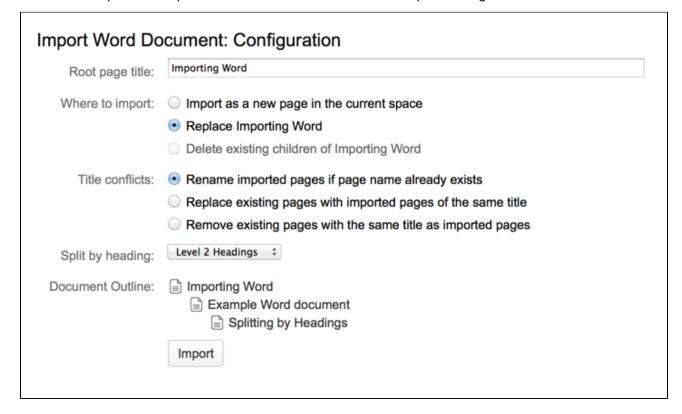
Screenshot: Import Word options for a document that contains multiple heading levels.

Orphaned Pages

An orphaned page is a page without any incoming links and is located at what's called the 'root of the space', meaning it sits alongside the space's home page. This means that, unless you know the page exists, you're not likely to come across it during the natural course of navigation.

There may be a legitimate reason you want to have a page at the root of a space. For example, you may be effectively hiding it in the navigation. If you don't want orphaned pages though, you can easily view all orphaned pages in a space so you can tidy delete the pages or reorganise them so they're no longer orphaned.

Related pages:

- Undefined Page Links
- Create and Edit Pages

To view the orphaned pages in a space:

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar
2. Choose **Orphaned Pages**

If your space uses the Documentation theme:

1. Choose **Browse** > **Space Operations** at the top of the screen
2. Choose **Orphaned Pages** in the space operations options

While in the orphaned pages view, you can do any of the following:

- Delete an orphaned page by choosing the 'trash can' icon next to the page name.
- Edit a page by choosing the 'pencil' icon next to the page name.
- Give an orphaned page a parent — see Move and Reorder Pages.

Screenshot: Managing orphaned pages

Publishing	Templates	Reorder Pages	Orphaned Pages	Undefined Pages	Attachments	Trash	Export	RSS Feeds

This report shows all the 'orphaned' pages which are not linked to by other pages.

Title (Space)	Creator	Created	Last Updated By	Updated	Operations
Appendix A - Extending JIRA (Atlassian OnDemand Documentation)	Atlassian Documentation Bot	Sep 25, 2012	Atlassian Documentation Bot	Jul 10, 2014	✏ · 🗑
_InclusionsLibrary (Atlassian OnDemand Documentation)	Andrew Lui [Atlassian Technical Writer]	Oct 22, 2008	Andrew Lui [Atlassian Technical Writer]	Oct 22, 2008	✏ · 🗑
Enterprise Hosted Service Deactivation (Atlassian OnDemand Documentation)	Christine Burwinkle [Atlassian Technical Writer]	Apr 04, 2013	Andrew Lui [Atlassian Technical Writer]	Jul 09, 2014	✏ · 🗑
_Redirects (Atlassian OnDemand Documentation)	Christine Burwinkle [Atlassian Technical Writer]	Feb 07, 2013	Andrew Lui [Atlassian Technical Writer]	Sep 24, 2013	✏ · 🗑

Undefined Page Links

You can add links to pages that don't yet exist in Confluence, but you intend to create later. Known as links to 'undefined pages', they allow you to create a link which, when clicked, will create a page with the name you specify in the link.

Related pages:

- Orphaned Pages
- Create and Edit Pages
- Links

Create an undefined page link

1. Choose **Insert** > **Link** or press **Ctrl+K** on your keyboard
2. Choose **Advanced**
3. Enter the name of the page to be created in the **Link** field

A link to an undefined page is shown in dark red while in the editor. When anyone clicks the link, Confluence will create a new page with the name you typed in the **Link** field.

View undefined pages in a space

The **Undefined Pages** view shows you all undefined pages in your space. The undefined page links are badged with a
✚
icon to remind you that those pages are yet to be created.

To view a list of the undefined links in a space:

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar
2. Choose **Undefined Pages**

If your space uses the Documentation theme:

1. Choose **Browse** > **Space Operations** at the top of the screen
2. Choose **Undefined Pages** in the space operations options

You can choose the link for an undefined page to create the page and add content to it.

View Page Information

The Page Information view for a page shows you useful information about the page.

To see the information about the page:

1. View the page.
2. Choose
   ●●●
   > **Page Information**
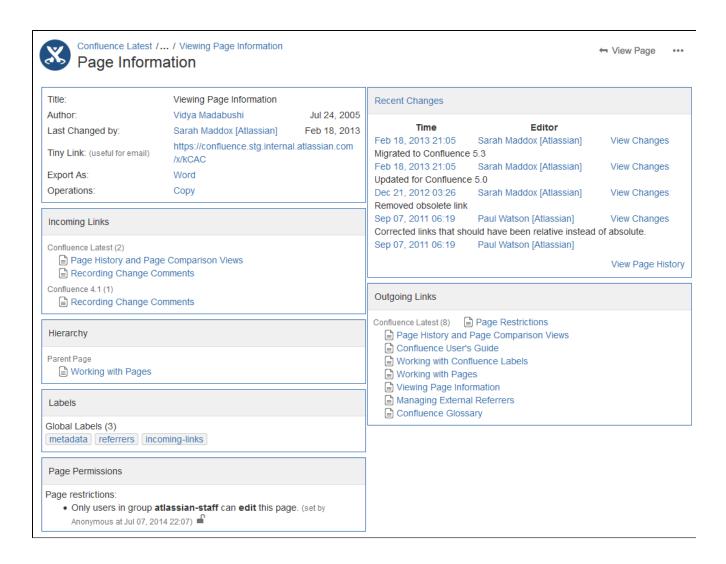
You will see the following information:

- **Page details:** Title, author, date of creation, date of last modification and the tiny link (permalink) of the page.
- **Page hierarchy:** Parent-child relationships of the page.
- **Incoming links:** Lists other pages in your Confluence Site that have links to this page, or reference this page in an Include Page or Excerpt Include macro.
- **Labels:** Any labels (tags) that have been applied to this page. See Add, Remove and Search for Labels.
- **Page Permissions:** Displays page-level security restrictions that apply to the page (if present). See Page Restrictions.
- **Hot Referrers:** The external website pages which send the most viewers to the page. See Trackback and External Referrers.
- **Recent Changes:** Links to the five most recent versions of the page along with the name of the editor and the date of modification. See Page History and Page Comparison Views. Choose **View page history** to see the page history view, all the versions of the page in reverse chronological order and allows you to compare versions or to restore a previous version.
- **Outgoing links:** A summary of the links contained on this page, pointing to other pages on the Confluence site or to external websites.

Note: if there is no information to report (for example the page has no restrictions or no incoming links), that section of the Page Information won't appear.

Screenshot: Page information for this page

Page History and Page Comparison Views

Confluence tracks the history of changes to each page by creating a new version of the page each time it's modified. You can view the changes between different versions, and roll back to a previous version if you need to.

Access the page history

To view the history of a page:

1. Go to the page and choose
   •••
   > **Page History**
2. Choose a version number to view the content of that version

Screenshot: Page history

View an older version

When you select a previous version of the page, you'll see a header like this at the top of the page:

ⓘ If you want to send this page version to someone, copy and paste the URL from your browser. The link will look something like this: `http://confluence.atlassian.com/pages/viewpage.action?pageId=12345.`

When you're viewing a specific version of the page, the following functions are available:

Function	Description
current version	View the latest version of the page.
Compare with Current	Compare the differences between the version of the page you are viewing and the current version.

Restore this Version	Roll back the content of the page to the previous version that you are viewing.
View Page History	Return to the list of page versions.
<< Previous and **Next >>**	View the previous or next version of the page.

Restore a previous version

1. Go to the page and choose
   •••
   > **Page History**
2. Choose **Restore this version** beside the version you want to restore (or at the top of the page if you've opened the version)
3. Change the default change comment if necessary, and choose **OK**

ℹ️ All page history is retained; restoring an older version creates a copy of that version. For example, if you restore version 39, Confluence will create a copy of version 39 and the copy will become the new, current version.

Delete a specific version

Choose **Delete** next to a version in the page history, to remove that version.

View the changes made

Using the page history view or the page information view, you can see the recent changes made to a page.

To view recent changes made to a page:

1. Choose
   •••
   > **Page Information**
   In the section titled 'Recent Changes' you'll see the most recent versions of the page, along with the date of their modification and the name of the modifying author.
2. Choose **View Changes** beside the required version
   The page comparison view is displayed, showing the differences between the selected and previous versions.

Compare two versions

1. Go to the page and choose
   •••
   > **Page History**
2. Choose the versions you want to compare by selecting the check boxes beside them
3. Choose **Compare selected versions**
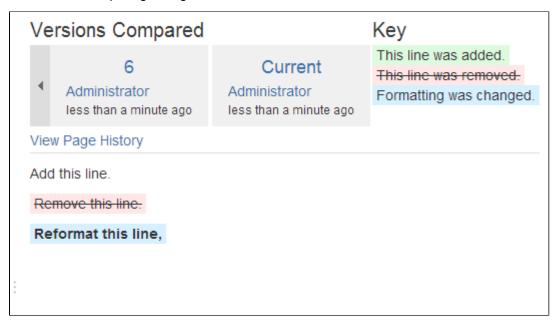
You'll see the page comparison view showing the differences between the selected versions. Changes are highlighted as follows:

Highlighted colour	Meaning
Green	Added content
Red	Deleted content

Blue	Changed formatting

Screenshot: Comparing changes

More about the comparison view

When you view a page comparison, all large sections of unchanged text are hidden and reduced to an ellipsis (. . .).

You can view page changes between versions which are adjacent to your current page comparison view. Click the link containing:

- **<<** to view the page comparison with the earlier adjacent version
- **>>** to view the page comparison with the more recent adjacent version

For example, if your page comparison view is between v. 30 and v. 34 of a page, you can view changes between:

- v. 29 and v. 30 by clicking **<< Changes from 29 to 30**
- v. 34 and v. 35 by clicking **Changes from 34 to 35 >>**

Confluence Markup

This section describes two types of markup found in Confluence:

Related pages:

- Macros
- The Editor

- **Confluence storage format.** Confluence stores the content of pages and blog posts in an XHTML-based format. Advanced users can view the storage format of a page and even edit it, provided their Confluence site is configured to allow that. See Confluence Storage Format for a description of the core elements of the storage format.
- **Wiki markup.** Confluence allows data entry via a shorthand code called wiki markup. Some parts of the Confluence administration interface also accept wiki markup for defining content. For a description of the wiki markup syntax, see Confluence Wiki Markup.

Storage format and wikimarkup code examples for macros can be found in the documentation for each macro.

Confluence Storage Format

This page describes the XHTML-based format that Confluence uses to store the content of pages, page templates, blueprints, blog posts and comments. This information is intended for advanced users who need to interpret and

edit the underlying markup of a Confluence page.

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.

You can view the Confluence storage format for a given page by choosing ••• > **View Storage Format**. This option is only available if one of the following is true:

- You are a Confluence administrator.
- Your Confluence site has the Confluence Source Editor plugin installed and you have permission to use the source editor.

- If you would like to edit the storage format for a page, your Confluence system administrator will need to install the Confluence Source Editor plugin.
- Clarification of terminology: If you choose ••• > **View Source**, you'll see the format used within the editor panel, not the storage format of the page.

On this page:

- Macros
- Headings
- Text effects
- Text breaks
- Lists
- Links
- Images
- Tables
- Page layouts
- Resource identifiers
- Template variables
- Instructional Text

Macros

Storage format and wiki markup examples have been included in the documentation for each macro.

Headings

Format type	In Confluence 3.5 and earlier	In Confluence 4.0 and later	What you will get	
Heading 1	h1. Heading 1	`<h1>Heading 1</h1>`	**Heading 1**	Underlined in the Documentation Theme
Heading 2	h2. Heading 2	`<h2>Heading 2</h2>`	**Heading 2**	Underlined in the Documentation Theme
Heading 3	h3. Heading 3	`<h3>Heading 3</h3>`	**Heading 3**	
Headings 4 to 6 are also available and follow the same pattern				

Text effects

Format type	In Confluence 3.5 and earlier	In Confluence 4.0 and later	What y
strong/bold	`*strong*`	`<strong>strong text</strong>`	**strong
emphasis	`_emphasis_`	`<em>Italics Text</em>`	*emphas*
strikethrough	`-strikethrough-`	`<span style="text-decoration: line-through;">strikethrough</span>`	~~striketh~~
underline	`+underline+`	`<u>underline</u>`	<u>underlin</u>
superscript	`^superscript^`	`<sup>superscript</sup>`	superscri
subscript	`~subscript~`	`<sub>subscript</sub>`	subscript
monospace	`{{monospaced}}`	`<code>monospaced</code>`	`monos`
preformatted	n/a	`<pre>preformatted text</pre>`	`prefor text`

block quotes	`bq. block quote`
	
	or
	
	`{quote}`
	
	`          block`
	
	`          quote`
	
	`{quote}`
		`<blockquote>`
		`<p>block quote</p></blockquote>`	

text colour	`{color:red}red text{color}`	`<span style="color: rgb(255,0,0);">red text</span>`	red text
small	n/a	`<small>small text</small>`	small te
big	n/a	`<big>big text</big>`	big tex
center-align	n/a	`<p style="text-align: center;">centered text</p>`	cent
right-align	n/a	`<p style="text-align: right;">right aligned text</p>`	right

Text breaks

Format type	In Confluence 3.5 and earlier	In Confluence 4.0 and later	What you will get

New paragraph	`Paragraph 1`
	
	`          (empty line)`
	
	`          Paragraph 2`
		`<p>Paragraph 1</p>`
		
		`<p>Paragraph 2</p>`
			Paragraph 1
			
			Paragraph 2

Line break	`Line 1 \\ Line 2`	`Line 1   Line 2` Note: Created in the editor using `Shift + Return/Enter`	Line 1 Line 2
Horizontal rule	`----`	`<hr />`	———
— symbol	`---`	`—`	—
– symbol	`--`	`–`	–

Lists

Format type	In Confluence 3.5 and earler	In Confluence 4.0 and later	What y will ge

Unordered list – round bullets	`* Round bullet list item`	`<ul>`
		
		`<li>round bullet list item</li>`
		`</ul>`
			• Rou bull list

Ordered list (numbered list)	`# Ordered list item`	`<ol>`
		
		`<li>numbered list item</li>`
		`</ol>`
			1. Ord list

Task Lists	`[] Task list item`	`<ac:task-list>`
		`    <ac:task>`
		`        <ac:task-status>incomplete</ac:task-status>`
		`        <ac:task-body>task list item</ac:task-body>`
		`    </ac:task>`
		`</ac:task-list>`
			☐ tas list ite

Links

Format type	In Confluence 3.5 and earlier	In Confluence 4.0 and later

Link to another Confluence page	`[Link to another Confluence page|Page Title]`	```<ac:link> <ri:page ri:content-title="Page <ac:plain-text-link-body> <![CDATA[Link to another Conflu </ac:plain-text-link-body> </ac:link>```
Link to an attachment	`[Link to an attachment^atlassian_logo.gif]`	```<ac:link> <ri:attachment ri:filename="atla <ac:plain-text-link-body> <![CDATA[Link to a Confluence A </ac:plain-text-link-body> </ac:link>```
Link to an external site	`[Atlassian|http://www.atlassian.com/]`	```<a href="http://www.atlassian.cc```
Anchor link (same page)	`[Anchor Link|#anchor]`	```<ac:link ac:anchor="anchor"> <ac:plain-text-link-body> <![CDATA[Anchor Link]]> </ac:plain-text-link-body> </ac:link>```
Anchor link (another page)	`[Anchor Link|pagetitle#anchor]`	```<ac:link ac:anchor="anchor"> <ri:page ri:content-title="pag <ac:plain-text-link-body> <![CDATA[Anchor Link]]> </ac:plain-text-link-body> </ac:link>```
Link with an embedded image for the body	`[!google.png!|pagetitle#anchor]`	```<ac:link ac:anchor="Anchor Link" <ac:link-body> <ac:image><ri:url ri:value="http://confluence.atla /></ac:image> </ac:link-body> </ac:link>``` For rich content like *images*, you nee

A note about link bodies

All links received from the editor will be stored as plain text by default, unless they are detected to contain the limited set of mark up that we allow in link bodies. Here are some examples of markup we support in link

bodies.

> **An example of different link bodies**
>
> ```
> <ac:link>
>   <!-- Any resource identifier -->
>   <ri:page ri:content-title="Home" ri:space-key="SANDBOX" />
>   <ac:link-body>Some <strong>Rich</strong> Text</ac:link-body>
> </ac:link>
> <ac:link>
>   <ri:page ri:content-title="Plugin developer tutorial stuff"
> ri:space-key="TECHWRITING" />
>   <ac:plain-text-link-body><![CDATA[A plain <text> link
> body]]></ac:plain-text-link-body>
> </ac:link>
> <ac:link>
>   <ri:page ri:content-title="Plugin developer tutorial stuff"
> ri:space-key="TECHWRITING" />
>   <!-- A link body isn't necessary. Auto-generated from the resource identifier
> for display. -->
> </ac:link>
> ```

The markup tags permitted within the <ac:link-body> are <b>, <strong>, <em>, <i>, <code>, <tt>, <sub>, <sup>, <br> and <span>.

Images

Format type	In Confluence 3.5 and earlier	In
Attached image	`!atlassian_logo.gif!`	
External image	`!http://confluence.atlassian.com/images/logo/confluence_48_trans.png!`	

Supported image attributes (some of these attributes mirror the equivalent HTML 4 IMG element):

Name	Description
ac:align	image alignment
ac:border	Set to "true" to set a border
ac:class	css class attribute.
ac:title	image tool tip.

ac:style	css style
ac:thumbnail	Set to "true" to designate this image as a thumbnail.
ac:alt	alt text
ac:height	image height
ac:width	image width
ac:vspace	the white space on the top and bottom of an image
ac:hspace	the white space on the left and right of an image

Tables

Format type	In Confluence 3.5 and earlier	In Confluence 4.0 and later	What you will get

Two column, two row (top header row)	`||Table Heading Cell 1||Table Heading Cell 2||`
	
	`                |Normal Cell 1|Normal Cell 2|`
		`<table>`
		`<tbody>`
		`<tr>`
		`<th>Table Heading Cell 1</th>`
		`<th>Table Heading Cell 2</th>`
		`   </tr>`
		`<tr>`
		`<td>Normal Cell 1</td>`
		`<td>Normal Cell 2</td>`
		`   </tr>`
		`</tbody>`
		`</table>`
			Table Heading Cell 1 / **Table Heading Cell 2** / Normal Cell 1 / Normal Cell 2

Two column, three rows, 2nd and third with merged cells in first row	N/A	```xml
		<table>
		
		<tbody>
		
		<tr>
		
		<th>Table Heading Cell 1</th>
		
		<th>Table Heading Cell 2</th>
		  </tr>
		
		<tr>
		
		<td rowspan="2">Merged Cell</td>
		
		<td>Normal Cell 1</td>
		  </tr>
		
		<tr>
		
		<td colspan="1">Normal Cell 2</td>
		  </tr>
		</tbody>
		</table>
		``	

Table Heading Cell 1	**Table Heading Cell 2**
Merged Cell	Normal Cell 1
	Normal Cell 2

Page layouts

Confluence supports page layouts directly, as an alternative to macro-based layouts (using, for example, the section and column macros). This section documents the storage format XML created when these layouts are used in a page.

Notes:

- Page layouts were originally introduced in Confluence 4.2, and are not available in earlier versions of Confluence. If you are using Confluence 4.2 - 5.1 you should refer to the documentation for your version of Confluence:
    - Confluence Storage Format (Confluence 4.2)
    - Confluence Storage Format (Confluence 4.3)
    - Confluence Storage Format (Confluence 5.0)
    - Confluence Storage Format (Confluence 5.1)
- Pages with a layout created in the old format will be converted to 5.2 format.
- Confluence 5.2 provides more flexible layouts with a more concise storage format.

Element name	In Confluence 5.2 and later	Attributes
`ac:layout`	Indicates that the page has a layout. It should be the top level element in the page.	None
`ac:layout-section`	Represents a *row* in the layout. It must be directly within the `ac:layout` tag. The type of the section indicates the appropriate number of cells and their relative widths.	`ac:type`

`ac:layout-cell`	Represents a *column* in a layout. It must be directly within the `ac:layout-section` tag. There should be an appropriate number of cells within the layout-section to match the `ac:type`.	None

The recognised values of `ac:type` for `ac:layout-section` are:

`ac:type`	Expected number of cells	Description
`single`	1	One cell occupies the entire section.
`two_equal`	2	Two cells of equal width.
`two_left_sidebar`	2	A narrow (~30%) cell followed by a wide cell.
`two_right_sidebar`	2	A wide cell followed by a narrow (~30%) cell.
`three_equal`	3	Three cells of equal width.
`three_with_sidebars`	3	A narrow (~20%) cell at each end with a wide cell in the middle.

The following example shows one of the more complicated layouts from the old format built in the new. The word `{content}` indicates where further XHTML or Confluence storage format block content would be entered, such as `<p>` or `<table>` tags.

```
<ac:layout>
  <ac:layout-section ac:type="single">
     <ac:layout-cell>
        {content}
     </ac:layout-cell>
  </ac:layout-section>
 <ac:layout-section ac:type="three_with_sidebars">
     <ac:layout-cell>
       {content}
     </ac:layout-cell>
     <ac:layout-cell>
       {content}
     </ac:layout-cell>
     <ac:layout-cell>
       {content}
     </ac:layout-cell>
   </ac:layout-section>
   <ac:layout-section ac:type="single">
     <ac:layout-cell>
        {content}
     </ac:layout-cell>
   </ac:layout-section>
 </ac:layout>
```

Emoticons

Format type	In Confluence 3.5 and earlier	In Confluence 4.0 and later	What you will get
Emoticons	:)	`<ac:emoticon ac:name="smile" />`	

	:(	`<ac:emoticon ac:name="sad" />`	😦
	:P	`<ac:emoticon ac:name="cheeky" />`	
	:D	`<ac:emoticon ac:name="laugh" />`	
	;)	`<ac:emoticon ac:name="wink" />`	
	(y)	`<ac:emoticon ac:name="thumbs-up" />`	
	(n)	`<ac:emoticon ac:name="thumbs-down" />`	
	(i)	`<ac:emoticon ac:name="information" />`	
	(/)	`<ac:emoticon ac:name="tick" />`	
	(x)	`<ac:emoticon ac:name="cross" />`	
	(!)	`<ac:emoticon ac:name="warning" />`	

Resource identifiers

Resource identifiers are used to describe "links" or "references" to resources in the storage format. Examples of resources include pages, blog posts, comments, shortcuts, images and so forth.

Resource	Resource identifier format

Page	
	```<ri:page ri:space-key="FOO" ri:content-title="Test Page"/>```

Notes:

- `ri:space-key`: (optional) denotes the space key. This can be omitted to create a relative reference.
- `ri:content-title`: (required) denotes the title of the page.

Blog Post	
	```<ri:blog-post ri:space-key="FOO" ri:content-title="First Post" ri:posting-day="2012/01/30" />```

Notes:

- `ri:space-key`: (optional) denotes the space key. This can be omitted to create a relative reference.
- `ri:content-title`: (required) denotes the title of the page.
- `ri:posting-day`: (required) denotes the posting day. The format is YYYY/MM/DD.

Attachment	
	```<ri:attachment ri:filename> ... resource identifier for the container of the attachment ... </ri:attachment>```

Notes:

- `ri:filename`: (required) denotes the name of the attachment.
- the body of the `ri:attachment` element should be a resource identifier denoting the container of the attachment. This can be omitted to create a relative attachment reference (similar to [foo.png] in wiki markup).

Examples:

Relative Attachment Reference
```
<ri:attachment ri:filename="happy.gif" />
```

Absolute Attachment Reference
```
<ri:attachment ri:filename="happy.gif">
 <ri:page ri:space-key="TST" ri:content-title="Test Page"/>
</ri:attachment>
```

URL	
	```<ri:url ri:value="http://example.org/sample.gif"/>```

Notes:

- `ri:value`: (required) denotes the actual URL value.

Shortcut	`<ri:shortcut ri:key="jira" ri:parameter="ABC-123">` Notes: • `ri:key`: (required) represents the key of the Confluence shortcut. • `ri:parameter`: (required) represents the parameter to pass into the Confluence shortcut. • The example above is equivalent to `[ABC-123@jira]` in wiki markup.
User	`<ri:user ri:userkey="2c9680f7405147ee0140514c26120003"/>` Notes: • `ri:userkey`: (required) denotes the unique identifier of the user.
Space	`<ri:space ri:space-key="TST"/>` Notes: • `ri:space-key`: (required) denotes the key of the space.
Content Entity	`<ri:content-entity ri:content-id="123"/>` Notes: • `ri:content-id`: (required) denotes the id of the content.

Template variables

This screenshot shows a simple template:

The template contains the following variables:

Variable name	Type	Values
$MyText	Single-line text	
$MyMulti	Multi-line text	Size: `5 x 100`
$MyList	List	List items: `Apples,Pears,Peaches`

The XML export produces the following code for the template:

```
<at:declarations>
  <at:string at:name="MyText" />
  <at:textarea at:columns="100" at:name="MyMulti" at:rows="5" />
  <at:list at:name="MyList">
    <at:option at:value="Apples" />
    <at:option at:value="Pears" />
    <at:option at:value="Peaches" />
  </at:list>
</at:declarations>


<p>This is Sarah's template</p>

<p>A single-line text variable: <at:var at:name="MyText" /></p>

<p>A multi-line text variable: <at:var at:name="MyMulti" /></p>

<p>A selection list: <at:var at:name="MyList" /></p>

<p>End of page.</p>
```

Instructional Text

Instructional text allows you to include information on how to fill out a template for an end-user (the person using creating a page from the template). Instructional text will:

- automatically clear all *instructional text* as the user types in a specific text block, and
- automatically trigger a @mention prompt for user selection (for 'mention' type instructional text).

Screenshot: Example of instructional text.

```
<ul>

<li><ac:placeholder>This is an example of instruction text that will get replaced
when a user selects the text and begins typing.</ac:placeholder></li>
</ul>
<ac:task-list>
    <ac:task>
        <ac:task-status>incomplete</ac:task-status>
        <ac:task-body><ac:placeholder ac:type="mention">@mention example. This
placeholder will automatically search for a user to mention in the page when the
user begins typing.</ac:placeholder></ac:task-body>
    </ac:task>
</ac:task-list>
```

Confluence Wiki Markup

This page describes the wiki markup used on some administration screens in Confluence.

Wiki markup is useful when you want to do one of the following:
- Configure the Documentation theme.
- Type wiki markup directly into the editor. Confluence will convert it to the rich text editor format as you type.
- Create links using the **Advanced** tab of the Links Browser.
- Insert a block of wiki markup (or markdown) into the Confluence editor. (Choose **Insert** > **Markup**.)

Note: You cannot edit content in wiki markup. Confluence does not store page content in wiki markup. Although you can enter wiki markup into the editor, Confluence will convert it to the rich text editor format immediately. You will not be able to edit the wiki markup after initial entry.

Can I type wiki markup into the editor?

On this page:

- Can I type wiki markup into the editor?
- Can I insert markdown?
- Headings
- Lists
- Tables
- Text Effects
- Text Breaks
- Links
- Images
- Page Layouts
- Macros

Yes. You can type wiki markup directly into the editor, and Confluence will convert it as you type. (You cannot edit the wiki markup after conversion.) **See it in action in this video**:

Read this blog post for more examples: 12 Things You Didn't Know About the Confluence Editor.

Can I insert markdown?

Confluence supports inserting content in markdown. This is often used in ReadMe files. See Markdown syntax guide for some examples of markdown syntax.

To insert markdown in the editor:

1. Choose **Insert** > **Markup**
2. Select **Markdown**
3. Type or paste your text - the preview will show you how it will appear on your page
4. Choose **Insert**.

As with wiki markup, Confluence will convert your markdown to the rich text editor format. You will not be able to edit your content using markdown.

Headings

To format a line as a heading, type "hn." at the start of your line, where n can be a number from 1 to 6.

What you need to type	What you will get
h1. Biggest heading	**Biggest heading**
h3. Big heading	***Big heading***
h5. Small heading	***Small heading***

Lists

Wiki markup allows you to create bulleted or numbered lists, and is flexible enough to allow a combination of the two list types.

ℹ️ If you need to separate the text within lists using line breaks, make sure you do so using a double slash (//). Empty lines may disrupt the list.

Simple lists

Use the hyphen (-) to create simple lists with square bullets. Make sure there's a space between the hyphen and your text.

What you need to type	What you will get
| ```
- some
- bullet
- points
``` | ▪ some<br>▪ bullet<br>▪ points |

Bulleted lists

Use the asterisk (*) to create bullets. For each subsequent level, add an extra asterisk.
Make sure there is a space between the asterisk and your text.

What you need to type	What you will get
| ```
* some
* bullet
** indented
** bullets
* points
``` | • some<br>• bullet<br>  • indented<br>  • bullets<br>• points |

Numbered lists

Use the hash (#) to create numbered lists.
Make sure there is a space between the hash and your text.

What you need to type	What you will get

| ```
# a
# numbered
# list
``` | 1. a<br>2. numbered<br>3. list |

A second level of hashes will produce a sub-list, such as the **alphabetical** sub-list shown below.

What you need to type	What you will get
| ```
# Here's a sentence.
## This is a sub-list point.
## And a second sub-list point.
# Here's another sentence.
``` | 1. Here's a sentence.<br>   a. This is a sub-list point.<br>   b. And a second sub-list point.<br>2. Here's another sentence. |

You can use a third level of hashes to produce a sub-sub-list.

What you need to type	What you will get
| ```
# Here's a sentence.
## This is a sub-list point.
### Third list level.
### Another point at the third level.
## And a second sub-list point.
# Here's another sentence.
``` | 1. Here's a sentence.<br>   a. This is a sub-list point.<br>      i. Third list level.<br>      ii. Another point at the third level.<br>   b. And a second sub-list point.<br>2. Here's another sentence. |

Note: In numbered lists as described above, the format of the 'number' displayed at each list level may be different, depending upon your browser and the style sheets installed on your Confluence instance. So in some cases, you may see letters (A, B, C, etc; or a, b, c, etc) or Roman numerals (i, ii, iii, etc) at different list levels.

Mixed lists

What you need to type	What you will get
| ```
# Here
#* is
#* an
# example
#* of
#* a
# mixed
# list
``` | 1. Here<br>   • is<br>   • an<br>2. example<br>   • of<br>   • a<br>3. mixed<br>4. list |

Tables

You can create two types of tables.

Table Type 1

Allows you to create a simple table with an optional header row. You cannot set the width of the columns in this table.
Use double bars for a table heading row.

What you need to type:

```
||heading 1||heading 2||heading 3||
|cell A1|cell A2|cell A3|
|cell B1|cell B2|cell B3|
```

What you will get:

heading 1	heading 2	heading 3
cell A1	cell A2	cell A3
cell B1	cell B2	cell B3

You can also use a vertical header.

What you need to type:

```
||heading 1|col A1|col A2|col A3|
||heading 2|col B1|col B2|col B3|
```

What you will get:

heading 1	col A1	col A2	col A3
heading 2	col B1	col B2	col B3

Table Type 2

This method allows you to specify the width of the columns in the table.

What you need to type

```
{section:border=true}
{column:width=30%}
Text for this column goes here. This is  the smaller column with a width of only
30%.
{column}
{column:width=70%}
Text for this column goes here. This is  the larger column with a width of 70%.
{column}
{section}
```

What you will get
Text for this column goes here.
This is the smaller column with a
width of only 30%.
Text for this column goes here. This is the larger column with a width of
70%.

ⓘ For more details please see the Column Macro and the Section Macro.

Advanced Formatting

Colour and Other Formatting

To add colour and other formatting to your tables, you can use the Panel Macro within columns.

More table-formatting options may be available if your Confluence administrator has installed additional macros.

Lists

Here's an example of how to embed lists in a table:

What you need to type

```
||Heading 1||Heading 2||
|* Item 1
* Item 2
* Item 3|# Item 1
# Item 2
# Item 3|
```

What you will get

Heading 1	Heading 2
• Item 1 • Item 2 • Item 3	1. Item 1 2. Item 2 3. Item 3

Text Effects

Use the markup shown in the examples below to format text.

What you need to type	What you will get
strong	**strong**
bold text	**bold text**
emphasis	*emphasis*
italics	*italics* *Hint:* To italicise parts of a word, add braces (curly brackets) around the underscore. For example, `Thing{_}x_` gives you this: Thing*x*
??citation??	*—citation*
-deleted-	~~deleted~~
+inserted+	<u>inserted</u>
Text with^superscript^	Text with^{superscript} Hint: There are two ways to make superscripts work, when used directly after another word or character: • Add a space before the superscript. For example, `kg /m ^3^` gives you this: kg/m ^{3} • Add braces (curly brackets) around the superscript markup. For example, `kg/m{^3^}` gives you this: kg/m^3

Text with~subscript~	Text with_{subscript}
`{{monospaced}}`	`monospaced`
`bq. Here's how you make a paragraph appear as a block quotation.`	> *Here's how you make a paragraph appear as a block quotation.*
{color:red}look ma, red text!{color}	look ma, red text!

Text Breaks

Paragraph Break

In wiki markup, a paragraph is a continuous line of text ending in two carriage returns. This is equivalent to a continuous line of text followed by a blank line.

When rendered into HTML, the result is a line of text wrapped in a set of `<p></p>` tags.

Line Break

Confluence provides two options for forcing a line break within a paragraph of text:

- Implicitly, by entering a single carriage return at its end.
- Explicitly, by entering two consecutive backslashes: \\

When rendered into HTML, the result is a paragraph of text that is split into separate lines by `<br>` tags, wherever a forced line break appears.

ℹ For most purposes, explicit line breaks are not required because a single carriage return is enough.

The examples below show how to use explicit line breaks.

What you need to type	What you will get
here is some text \\ divided \\ using line \\ \\ breaks\\	here is some text divided using line breaks
This is a short list: * Point 1 Text to go with point 1 * Point 2 \\ \\ Text to go with point 2 with a break	This is a short list: - Point 1 Text to go with point 1 - Point 2 Text to go with point 2 with a break

If you wish to use multiple consecutive line breaks, each should be separated by a space character. For example, use this for two consecutive line breaks:
\\ \\

Horizontal Rule

To create a horizontal line across the width of your page or content block, type four dashes (like this: ----) at

the beginning of a line, then press Enter or space.

Make sure that the dashes are on a separate line from the rest of the text.

What you need to type	What you will get
here is some text ---- divided by a horizontal rule	here is some text ――― divided by a horizontal rule

Links

You can use wiki markup to add hyperlinks to your text.

What you need to type	What you will get
[#anchor]	A link to an anchor on the same page.
[Confluence Wiki Markup^attachment.ext]	A link to a file attached to the page.
[pagetitle]	A link to a page.
[pagetitle#anchor]	A link to an anchor on another page.
[pagetitle^attachment.ext]	A link to a file attached to another page.
[spacekey:pagetitle]	A link to a page in another space.
[spacekey:pagetitle#anchor]	A link to an anchor on a page in another space.
[spacekey:pagetitle^attachment.ext]	A link to a file attached to a page in another space.
[/2004/01/12/blogposttitle]	A link to a blog post. Note: blogposttitle is the title of the blog as it appears on the page.
[spacekey:/2004/01/12/blogposttitle]	A link to a blog post in another space. Note: blogposttitle is the title of the blog as it appears on the page.
[/2004/01/12]	A link to a whole day's blog posts.
[spacekey:/2004/01/12]	A link to a whole day's blog posts in another space.
[spacekey:]	A link to the space homepage (or the space summary page of the space.
[~username]	A link to the user profile page of a particular user.
[phrase@shortcut]	A shortcut link to the specified shortcut site. Shortcuts are configured by the site administrator.
[http://confluence.atlassian.com]	A link to an external resource.
[mailto:legendaryservice@atlassian.com]	A link to an email address.
[file://z:/file/on/network/share.txt]	A link to a file on your computer or on a network share that you have mapped to a drive. *This only works on Internet Explorer.*

| [!http://external/image.png!|http://external/link.html] | Displays an external image and links to an external URL. |

Note that Confluence treats headings as anchors, so you can link to headings using this pattern: [spacekey:pagename#headingname], where headingname is case-sensitive and must be entered without spaces.

For each of these link forms:

- You can prepend a link alias, so that alternate text is displayed on the page. Example: [link alias|pagetitle#anchor]
- You can append a link tip, which appears as a tooltip. Example: [pagetitle#anchor|link tip]

Images

You can display images from attached files or remote sources.

What you need to type	What you will get
!http://www.host.com/image.gif!	An image from a remote source is displayed on the page. Uses a fully qualified URL.
!attached-image.gif!	An image file attached to the page is displayed.
!pageTitle^image.gif!	An image file attached to a different page is displayed.
!spaceKey:pageTitle^image.gif!	An image file attached to a page in a different space is displayed.
!/2010/05/23/My Blog Post^image.gif!	An image file attached to a blog post is displayed.
!image.jpg|thumbnail!	The image is displayed as a thumbnail on the page (only works with images that are attached to the page). Users can click on the thumbnail to see the full-sized image. Thumbnails must be enabled by the site administrator for this to work.
!image.gif|align=right, vspace=4!	For any image, you can specify attributes of the HTML image tag as a comma separated list of name=value pairs.

Available HTML image tags include:

Image tag	Details
align	Available values are 'left', 'right', 'bottom', 'center', 'top'.
border	Specifies the width of the border (in pixels).
bordercolor	Use with the 'border' tag. Specify colours by name or hex value.
hspace	Specifies the amount of whitespace to be inserted to the left and right of the image (in pixels).
vspace	Specifies the amount of whitespace to be inserted above and below the image (in pixels).
width	Specifies the width of the image (in pixels). This will override the natural width of the image.
height	Specifies the height of the image (in pixels). This will override the natural height of the image.
title	Specifies alternate text for the image, which is displayed when the pointer hovers over the image.

| alt | Specifies alternate text for the image. This text is retrievable via search, and contributes to accessibility of the page for text-only viewing. |

Page Layouts

There is no wiki markup representation for page layouts.

Macros

Storage format and wiki markup examples have been included in the documentation for each macro.

Files

Share your team's PDFs, Office documents, images, and more in one place by uploading your files to Confluence. Automatic versioning, instant previews, permissions, and full-text search, means shared network drives can be a thing of the past for your team.

When you upload a file it is *attached* to the current page or blog post. This is why files are often referred to as *attachments* in Confluence.

You can attach anything from project plans and design mockups to video and audio files. You and your colleagues can also collaborate by commenting on files displayed on Confluence pages.

On this page:

- Using Files
- Permissions

Related pages:

- Configuring Attachment Size
- Configuring your Attachment Storage

Using Files

- Upload Files
- Display Files and Images
- Manage Files
- Share and Comment on Files
- Edit Office Files

Permissions

The 'Add Attachment' and 'Delete Attachment' permissions are used to control who can upload and delete attachments in a space.

Users with 'Add Page' or 'Add Blog' permissions can insert existing attachments to their pages, but not upload new attachments unless they also have the 'Add Attachment' permission.

There is no permission that controls downloading attachments. See our knowledge base article about disabling the download of attachments if you need to do this.

Upload Files

When you upload a file, such as an image or document, it will be attached to the current page.

You can then choose to display the file on the page as a link, an image or embed it in the page (using a macro).

To upload a file you'll need the 'Add Attachments' space permission.

On this page:

- Upload a file
- Accepted file types and size
- File versions
- Notes

Related pages:

- Files
- Configuring Attachment Size
- Display Files and Images
- Manage Files

Upload a file

There are many ways to attach a file to a page.

In the editor you can:

- Drag the file directly onto the page.

- Go to **Insert** > **Files and images** and upload a file.

When viewing a page you can:

- Drag the file directly onto the page.
- Go to
  •••
  > **Attachments** and upload a file.

You can attach multiple files at a time.

Accepted file types and size

Confluence allows you to attach most file types, but you cannot attach a folder of files (including folders created by applications like Keynote - you'll need to export your presentation to zip or other format).

Although just about any file type can be attached to a page, not all file types can be displayed on or embedded in a page. See Display Files and Images to find out more.
The maximum file size you can upload to Confluence is set by your system administrator. By default it is 10mb, but your administrator may have increased or reduced this limit.

File versions

If you upload a file with the same name as an existing attachment on the same page, Confluence will overwrite the existing attachment. Version history is kept for all attachments. See Manage Files to find out more.

Any changes you make to the source file will not affect the copy that was uploaded to Confluence. To update the Confluence copy, you need to upload the new version of the file.

Notes

⚠ We recommend you don't use special characters in page or attachment names, as the page or attachment may not be found by Confluence search, and may cause some Confluence functions to behave unexpectedly.

Display Files and Images

Files can be displayed on a page as a thumbnail or as a link. There are a few different ways to Upload Files.

You can control how the file appears on your page. The options available depend on the type of file.

Insert a file into your page

There's a few ways to insert a file into your page:

- Go to **Insert** > **Files** on the editor toolbar and select any of the previously uploaded files, or
- Drag the file directly into the editor (this will upload and insert the file in one step), or
- Type ! and choose an attached file from the autocomplete drop down.

Your file will appear on your page as a thumbnail. Click the thumbnail to resize it or to switch to showing the file as a link.

On this page:

- Insert a file into your page
  - Insert a file attached to another page
  - Insert an image file from the web
  - Delete files from your page
- Preview a file
- Office and PDF files
- Image files
- Multimedia files
- Show a list of files on a page

Related pages:

- Links
- Manage Files

Insert a file attached to another page

You can display a file that's attached to a different page of the same Confluence site, if you know the name of the file.

To display an image attached to a different page:

1. Go to **Insert** > **Files** and choose the **Search on other pages**.
2. Enter the name of the file.
3. Choose whether to search the current space or **All Spaces** and choose **Search**.
4. Select the file from the search results and choose **Insert**.

Insert an image file from the web

You can display an image from a remote web page on your Confluence page, without needing to attach it to your page. You need to know the URL for the image, not for the web page it appears on. This is only available for image files, not other types of files (like documents).

To display an image from a web page:

1. While editing the page, position the cursor where you want to place the image.
2. Choose **Insert** > **Files** and choose **Images from the web**.
3. Enter a URL for the image. (example: http://atlassian.wpengine.netdna-cdn.com/wp-content/uploads/AtlassianBushRegeneration-12January2012-083-trunc.jpg
4. Choose **Preview** to check that the URL and image are correct.
5. Choose **Insert**.

Delete files from your page

If you delete a file or image in the editor, the attached file will not be deleted. Go to
•••
> **Attachments** to delete the attachment completely from the page.

Seeing an 'unknown attachment' placeholder on your page? This means that the attached file has been deleted from the page (or another page).

Preview a file

Click an image, file thumbnail or link when viewing a page to launch the preview.

The preview includes images from the web that are displayed on the page and files that are attached to the page (even if they are not currently displayed on the page).

In the preview you can:

- Download the image file.
- Upload a new version of the file (attached files only).
- Comment on the file.
- Zoom in, out or fit the image to the width of your browser.
- Browse like a slideshow using the next and back buttons.
- See other files attached to the page and select a thumbnail to preview that file.
- Switch to a full screen presentation mode.

Many file types can be previewed, including Office files, PDFs and many image types.

ˇ Which files can be previewed?

Images files	Office files	Other files
• JPEG • PNG • TIFF • PSD • WMF • EMF • ICO • ICNS	• DOC • DOCX • PPT • PPTX • XLS • XLSX	• PDF

Download the file, upload a new version, or share with your team

Drag the pin to comment on the file

Also attached to this page

Office and PDF files

Inserting a file in a page is a great way to make useful documents, spreadsheets, presentations and other files available to your team.

As with all file types, you can choose to insert the file as a link, or as a thumbnail. The thumbnail shows a preview of the document's contents, and can be resized.

To view an Office or PDF file, click the link or thumbnail to see the full preview (no need to have Excel, Word or PowerPoint installed). Alternatively, use the Download button in the preview to download the file and view offline.

You can even comment on Office and PDF files.

Image files

When editing the page, select an image to show the image properties panel. The panel allows you to set the display size, add a border and effects and link the image to other pages.

From the image properties panel you can:

• Choose a **preset size** for the image

- Enter a **width** for the image (between 16px and 900px)
- Add a **border** around the image
- **Link** the image to a page or URL
- **Align** the image (you can use the left and right align buttons to make the text wrap around the image too)
- Add a title, which is shown when you hover over the image (go to **Properties** > **Title**)
- Add alt text, which is used by screen readers and when the image can't be shown (go to **Properties** > **Title**)
- Add **effects** to the image such as drop shadow or snapshot (go to **Properties** > **Effects**).

To add a caption to an image using the Instant Camera effect:

- Choose **Effects** in the image properties panel and choose the **Instant Camera** image effect.
- Save the page.
- Go to
  ●●●
  > **Attachments** to go to the 'Attachments' view of the page.
- Choose **Properties** next to the image file.
- Add a **comment** to the attachment. The text in your comment will appear as the image caption.

You will need to re-enter the comment each time you upload a new version of the image.

Note: The Instant Camera effect only works with Latin character languages, due to a lack of handwriting style fonts in multi-byte languages.

Note about image effects

Displaying image effects can be resource intensive. Confluence limits the threads that are dedicated to displaying image effects so that it does not impact your whole instance. If a thread is not available, Confluence will display the image without the effect.

The following knowledge base article provides information if you need to adjust the number of threads - Imag e effects are not displayed in Confluence 5.5 or later.

Multimedia files

You can display multimedia files (such as video, audio and animation) that are attached to your page using the Multimedia Macro.

Display online multimedia (such as YouTube or Vimeo videos) using the Widget Connector Macro.

Show a list of files on a page

There are several ways you can display a list of files on a page. You can:

- Use the Attachments Macro to show files attached to the current page.
- Use the Space Attachments Macro to show all files in a space.
- Use the Gallery Macro to show thumbnails of images attached to a page.

You can also use the File List blueprint for uploading, viewing and managing lists of files.

Manage Files

Files are attached to Confluence pages. See Upload Files to find out about attaching files to pages.

Once attached you can download, delete and edit

these files, for example if you need to upload a new version of the file, or change the page it is attached to.

Download attached files

Any user with permission to view a page can also download any files attached to that page.

To download an individual file:

- Click the **Download** button in the file preview, or
- Go to
  •••
  > **Attachments**
- and then right click on the file name and save the link.

To download all files attached to a page as a zip file:

1. Go to
   •••
   > **Attachments**
2. Click **Download All**.

There's no option to download all attachments in a space.

Delete an attached file

You'll need the 'Delete Attachment' space permission to delete an attached file.

To delete all versions of an attached file:

1. Go to the page that contains the attachment.
2. Go to
   •••
   > **Attachments**
3. Choose **Delete** next to the attachment you want to delete.
4. Choose **Delete** to confirm your action.

Deleted files can be restored from the trash. You'll need to be a space admin to do this.

Space Admins can also delete a specific versions of an attachment:

1. Go to
   •••
   > **Attachments**
2. Click the expand arrow next to the attachment name to see the list of attachment versions
3. Choose **Delete** next to the version you want to delete.

Deleted file versions are not recoverable from the trash.

Screenshot: Attachments and attachment versions

On this page:

- Download attached files
- Delete an attached file
- Upload a new version of an attached file
- Move a file to another page
- Edit properties of an attached file
- View all attached files in a space

Related pages:

- Files
- Display Files and Images

Upload a new version of an attached file

There are two ways up upload a new version of an attached file. You can:

- Upload a file with the same file name to the page.
- Use the **Upload a new version**

  ⬆

  button in the file preview to upload a file with a different name (for images and PDFs only).

To view attachment versions:

- Go to
  •••
  > **Attachments**
- Click the expand arrow next to the attachment name.

All earlier versions of the file will appear.

You can't revert to an earlier version of the file, but you can choose to remove earlier versions if you have Space Administrator permissions.

Move a file to another page

You'll need the 'Add Page', 'Add Attachment' and 'Remove Attachment' space permissions to move an attached file to another page.

To change the page that a file is attached to:

- Go to
  •••
  > **Attachments**
- Choose **Properties** next to the attachment you want to move.
- Enter the name of the page you want to move the attachment to (for example `My Destination Page`).
- Choose **Save**.

If you want to move the file to a page in another space, add the space key before the page name (for example DOC:My Destination Page).

Edit properties of an attached file

You'll need the 'Add Attachment' permission in the space to edit the file properties.

To edit the properties of an attached file:

- Go to
  •••
  > **Attachments**
- Click **Properties** beside the attachment you want to edit.

You can:

- change the file name
- add a comment (used in the version list and also by the Snapshot image effect)
- change the MIME type
- move the attachment to another page
- add a label.

Changing the MIME type may cause your file to display incorrectly.

View all attached files in a space

There are two ways you can view all files in a space. You can:

- Use the Space Attachments macro to display the list of files on a page.
- Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar Then choose **Attachments**.
  If your space uses the Documentation theme go to **Browse** > **Space Operations** > **Attachments**.

You can use the filters to only show files with a particular label or file extension.

Screenshot: Space attachments macro

Share and Comment on Files

Collaboration doesn't just happen on pages; often you'll need to collaborate with your team on documents, presentations, images and spreadsheets. Whether it's mockups for a new marketing campaign or a full project plan, you can simplify your team's feedback loop by working together on files in Confluence.

Share a file

Do you have lots of files on a page and want to get a team member's input on just one of them? You can share the file with them directly.

It works just like sharing a page:

1. Click the thumbnail or link to preview the file
2. Choose the **Share** button
3. Enter an email address, user name or group name, add your message and send

Your team members will get an email with your message and a link to view the file.

> Share notifications are only sent by email, they won't appear in the workbox

On this page:

- Share a file
- Comment on a file
- Edit a file

Related pages:

- Files
- Display Files and Images
- Comment on pages and blog posts

.

Comment on a file

Whether it's an image – like a mockup of the new marketing campaign that needs feedback – a PDF, a presentation, or any other file you can preview in Confluence, you can drop a pin anywhere on the preview and add your comment to start a conversation.

To comment on a file:

1. Click the thumbnail or link to preview the file
2. Drag the pin icon from the bottom of the preview and drop it where you want to comment
3. Add your comment and **Save**

Pinned comments work just like inline comments on pages. You can use @mentions, links, and macros, and drop as many pins as you need on any part of the file. Anyone with permission to add comments to the page can add and reply to comments on a file.

When you preview a file, you'll see pins for any existing comments. Select a pin to view the comment.

Once the conversation is finished, you can resolve the comment to hide it (and any replies) from view. If you need to see resolved comments again, you can reopen them. Go to

> **Resolved comments** in the preview.

Drag a pin onto a file to comment.
Drag more pins to make more commen

You can't comment on files that are hosted on a web server and added to Confluence using their URL, or on files that can't be viewed in the preview (such as videos, zip files, and some other file types).

Edit a file

Need to go beyond commenting on files? You can even edit some files right from the Confluence page. See Edit Office Files to find out about supported file types and applications.

Edit Office Files

The Office Connector allows you to edit attached office files in their native application (such as Word, Excel, PowerPoint or OpenOffice) and save the file right back to the Confluence page. No need to download and re-upload the file.

⚠️ The Office Connector is only available for some browsers, operating systems and applications, check the Office Connector Prerequisites before you start for the combination that will work best for you.

On this page:
- Edit an attached document
- Other ways to edit
- Troubleshooting

Related pages:
- Office Connector Prerequisites
- Office Connector limitations and known issues
- Configuring the Office Connector

Edit an attached document

To edit an Office document attached to a Confluence page:

1. Go to
   •••
   > **Attachments**
2. Choose **Edit in Office** beside the attachment you want edit.
   Your browser will ask you to confirm that you want to open the file.
3. Choose **OK**.
   You may also see a security warning or be asked to log in to your Confluence server - enter your Confluence username and password, then choose **OK**.
4. The file will open in your Office application - make your changes then save the document. It will be saved back to Confluence

Screenshot: Edit in Office option on the attachments page

Other ways to edit

Edit options also appear in the:

- Attachments macro (choose **Edit in Office** beside each attached office file)
- File preview (choose
  •••
  > **Edit in Office**)

Troubleshooting

Check out the Office Connector Prerequisites for the combinations of browsers, operating systems and office applications that are supported. Here's some common issues:

- **Using Chrome?** You can't edit Office documents as Chrome does not support WebDAV clients**.** See

CONF-23322
OPEN

- **Using Firefox?** You'll need to install the WebDAV add-on. See Installing the Firefox Add-On for the Office Connector.
- **Using Internet Explorer?** You can only edit documents in Microsoft Office. OpenOffice is not supported.
- **Using Linux?** You can only edit documents in OpenOffice. Microsoft Office is not supported. See

CONF-17250
OPEN

  .

- **Using Mac OS X?** You can't currently edit documents. See `OPEN` **CONF-25594** .

- **Not seeing the Office Connector options?** Your system administrator may have disabled all or part of the Office Connector.
  See Configuring the Office Connector.

You can find more troubleshooting info in the Office Connector Limitations and Known Issues knowledge base article.

Office Connector Prerequisites

The Office Connector allows you to **import**, and **edit** office documents (such as Word, Excel and PowerPoint) within Confluence.

The browser, operating system and applications required depend on what you are trying to do with an office file.

On this page:
- Viewing Office files
- Importing Word documents as pages
- Editing Office files attached to a page
- Troubleshooting

Viewing Office files

You can upload, insert and preview Office files in the same way as any other file. You don't need to have an Office application installed on your computer in order to view Office files in Confluence.

Importing Word documents as pages

Confluence can import the content from Microsoft Word 97-2013 documents (`.doc` and `.docx`).

Editing Office files attached to a page

The Office Connector allows you to edit Office files that are attached to pages.

You'll need to use a browser, operating system and application (either Microsoft Office or OpenOffice) as described in the compatibility matrix below.

Here's a few common issues:

- **Using Chrome?** You can't edit Office documents as Chrome does not support WebDAV clients**.** See

CONF-23322
OPEN

- **Using Firefox?** You'll need to install the WebDAV add-on. See Installing the Firefox Add-On for the Office Connector.
- **Using Internet Explorer?** You can only edit documents in Microsoft Office. OpenOffice is not supported.
- **Using Linux?** You can only edit documents in OpenOffice. Microsoft Office is not supported. See

CONF-17250
OPEN

  .

- **Using Mac OS X?** You can't currently edit documents. See CONF-25594 OPEN .
- **Not seeing the Office Connector options?** Your system administrator may have disabled all or part of the Office Connector.
  See Configuring the Office Connector.

Configuration matrix

You need one of the following software combinations to edit Office files from your Confluence page.

Software	Operating System	Browser
Microsoft Office 2013	• Windows 7 • Windows 8	• Internet Explorer 11 • Firefox – latest stable version
Microsoft Office XP, 2003, 2007, 2010	• Windows 7 • Windows Vista • Windows XP (Service Pack 2 or 3)	• Internet Explorer supported versions • Firefox – latest stable version
OpenOffice 2.x – 3.x	• Windows 7 • Windows Vista • Windows XP (Service Pack 2 or 3) • Linux	• Firefox – latest stable version • Internet Explorer supported versions (Windows Platforms only)

Note: The only known supported Office editor for Linux is OpenOffice. But in theory it should work with any WebDAV-aware application.

If you experience problems editing documents using the Office Connector (using an application, operating system and browser combination above) please raise an issue and tell us as much as you can about your operating system, application version, document version (if its different to the version of Office / Open Office you're using to open the document) and browser.

Troubleshooting

Having problems with the Office Connector?
- The **WebDAV plugin** must be enabled, because the Office Connector uses WebDAV to transfer information to and from Office documents. The WebDAV plugin is bundled with Confluence, and can be enabled or disabled by the System Administrator. If necessary, refer to the instructions on enabling plugins and configuring the WebDAV options.
- Ensure that your Confluence server's **base URL** is set correctly (see Configuring the Server Base URL to find out how to check this). When a user edits a Confluence page in Word and then uploads the page back to the Confluence server, the base URL determines where the document will be saved. If the base URL is incorrect, the documents may be saved to a different Confluence server.
- **Using Office 2013?** Your administrator will need to enable 'Allow authentication tokens in the URL path' in the Office Connector configuration. See Configuring the Office Connector.

See the Office Connector Limitations and Known Issues knowledge base article for more troubleshooting tips.

Install the Firefox Add-On for the Office Connector

If you are using Firefox as your browser, you will need to install an add-on into Firefox (the Firefox WebDAV Launcher) in order to use some features of the Office Connector. You will be prompted to install the add-on the first time you try to use a function which requires it.

⚠ For more information on the browsers, operation systems and application combinations required to use the Office Connector see Office Connector Prerequisites.

On this page:

- Installing the Firefox Add-On
- Configuring the Add-On
- Security Risks
- Information about this Firefox add-on

Related pages:

- Edit Office Files

Installing the Firefox Add-On

The first time you try to edit an attached Office document in Firefox you'll be prompted to install a WebDAV add-on for your browser. Without this add-on you can't edit Office documents from Firefox.

1. Choose **Edit in Office** or **Edit Document** next to the document you wish to edit.
    You'll see a prompt to download the add-on.

    A plugin is required to use this feature. Would you like to download it?

    OK   Cancel

2. Choose **OK**.
    Your browser will ask you to confirm that you want to allow the add-on.

    Firefox prevented this site (pug.jira-dev.com) from asking you to install software on your computer.

    Allow

3. Choose **Allow**.
    An installation dialog like the one below will appear.

    Software Installation

    Install add-ons only from authors whom you trust.

    Malicious software can damage your computer or violate your privacy.

    You have asked to install the following item:

    WebDAV Launcher *(Atlassian Pty Ltd)*
    https://update.atlassian.com/office-connector/URLLauncher/latest/webdavloader.xp

    Install Now   Cancel

4. Choose **Install Now**.
    Once the installation is complete you'll be prompted to restart Firefox. Make sure you've saved any open pages or other work in your browser before you restart.
5. Next you need to configure the add-on, and tell it which applications to use.

Configuring the Add-On

After installing the add-on you need to configure it, to tell the add-on which desktop applications to launch for each file type.

Configuring the add-on is slightly different in each operating system.

Configuring the add-on in Windows

In most cases the add-on will automatically configure itself, based on information from the system registry. The first time you edit a new file type, the add-on will look up the default editor for that file type and make that the permanent setting.

If you want to override the registry settings, or the automatic configuration is not working, you can configure the Firefox add-on manually.

1. In Firefox locate the **WebDAV Launcher** add-on and choose **Options**.
   The WebDAV Launcher Options dialog, like the one below, will appear.

2. Enter a file extension in the **File Extension** field. Do not include the leading period in the extension name.
   For example, to associate the 'doc' file extension with Microsoft Word 2003 you would type `doc` in the **File Extension** field.

   > **Using Office 2007 or later?**
   > You should configure the WebDAV launcher to open both Office 2003 file extensions (doc, ppt and xls) and Office 2007 and later file extensions (docx, xlsx and pptx).

3. Enter the **Application Path** - you can either:
   - Choose **Auto** to load the associated application from the Windows registry, or
   - Choose **Browse** to find the application on your computer, or
   - Manually enter the path to the application's executable file.

4. Choose **Add** to add the file extension association to the list.
   Repeat this process for all the file extensions you need.

Configuring the Add-On in Mac OS X

The configuration procedure is similar to Windows, as shown above. Note that there is no 'Auto' button in Mac OS X.

You can't currently edit documents in Mac OS X. See **CONF-25594** **OPEN** .

Configuring the Add-On in Linux

The only known supported Office editor for Linux is OpenOffice (see Office Connector Prerequisites).

There is no automatic configuration on Linux. You will need to associate each file type with your Office editor in Firefox. Note that there is no 'Auto' button in Linux.

The configuration procedure is similar to Windows. For both Ubuntu and OpenSUSE, the configuration will look something like this:

File Extension	Application Path
doc and docx	/usr/bin/oowriter
ppt and pptx	/usr/bin/ooimpress
xls and xlsx	/usr/bin/oocalc

Security Risks

> Please be aware that there are security risks in installing this add-on to Firefox. Internet Explorer is exposed to the same risks, because it can directly open Office documents.

Summary of the risks:

- Office documents can contain macro viruses. Before opening an Office document, make sure that you trust the source of the document.

- There are known flaws in the Office file formats and Microsoft Office that an attacker can exploit to gain control of your machine. Microsoft has fixed the known exploits in the latest Service Packs for all Microsoft Office versions, but new exploits may arise at any time. Be sure that you trust the source of a document before opening it.

The add-on reduces the risk by supporting:

- **Same origin policy** - the add-on can only open documents from the same host that initiated the action.
- **Digital signature** - the add-on is digitally signed. When you install the add-on please verify that it is signed by Atlassian.
- **Prompt the user for confirmation** - You will always be warned before a file is opened. **Please read these warnings carefully before opening a file**. The warning tells you the complete URL of the file as well as the complete application path of the application opening the file.

Information about this Firefox add-on

Add-on name:	WebDAV Launcher
Description:	Adds the ability to launch a WebDAV URL directly in a WebDAV-aware client.
Latest version:	See the WebDAV Launcher add-on page at Mozilla
Compatible with:	All versions of Firefox, up to and including the latest stable version.
Author:	Atlassian Pty Ltd

Macros

Using macros helps you to expand the capabilities of your Confluence pages, allowing you to add extra functionality or include dynamic content. For example, use the Attachments macro to list files attached to a page, or use the Widget Connector macro to include things like a YouTube video or Twitter feed.

On this page:

- Add a Macro to your Page
- Confluence macros
- Creating your own macros
- Getting more macros from The Marketplace

Add a Macro to your Page

Including Macros with the Macro Browser

To add a macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Macro Parameters

Many macros have optional parameters you can use to control the macro's output.

With the Attachments Macro, for instance, you have two optional parameters allowing you to:

- Specify the file formats of the attachments displayed

- Choose whether or not you want old versions of the attachments displayed

Macro Placeholders

Macro placeholders are displayed in the editor where you have added a macro to the page.

When editing a page, you can:

- Double-click a macro placeholder (or click the placeholder and choose **Edit**) to open the macro dialog window and edit the macro's parameters
- Select a macro placeholder to cut, copy and paste the macro

Confluence macros

Below is a list of the macros currently bundled with Confluence. Click a macro name for details of the usage, including optional parameters and examples.

- Activity Stream Macro
- Anchor Macro
- Attachments Macro
- Blog Posts Macro
- Change-History Macro
- Chart Macro
- Cheese Macro
- Children Display Macro
- Code Block Macro
- Column Macro
- Content by Label Macro
- Content by User Macro
- Content Report Table Macro
- Contributors Macro
- Contributors Summary Macro
- Create from Template Macro
- Create Space Button Macro
- Excerpt Include Macro
- Excerpt Macro
- Expand Macro
- Favourite Pages Macro
- Gadget Macro
- Gallery Macro
- Global Reports Macro
- HTML Include Macro
- HTML Macro
- IM Presence Macro
- Include Page Macro
- Info, Tip, Note, and Warning Macros
- JIRA Chart Macro
- JIRA Issues Macro
- JUnit Report Macro
- Labels List Macro
- Livesearch Macro
- Loremipsum Macro
- Multimedia Macro
- Navigation Map Macro
- Network Macro
- Noformat Macro
- Office Excel Macro
- Office PowerPoint Macro
- Office Word Macro
- Page Index Macro
- Page Properties Macro
- Page Properties Report Macro
- Page Tree Macro

- Page Tree Search Macro
- Panel Macro
- PDF Macro
- Popular Labels Macro
- Profile Picture Macro
- Recently Updated Dashboard Macro
- Recently Updated Macro
- Recently Used Labels Macro
- Related Labels Macro
- Roadmap Planner Macro
- RSS Feed Macro
- Search Results Macro
- Section Macro
- Space Attachments Macro
- Space Details Macro
- Space Jump Macro
- Spaces List Macro
- Status Macro
- Table of Contents Macro
- Table of Content Zone Macro
- Task Report Macro
- User List Macro
- User Profile Macro
- View File Macro
- Widget Connector Macro

Creating your own macros

Users with System Administrator permissions can create user macros - see Writing User Macros.

If you want to create something more complex, you can develop your own plugin - see Writing Confluence Plugins.

Getting more macros from The Marketplace

You can find a wide range of Atlassian and third party macros at The Marketplace. These are distributed as add-ons and can be installed by a Confluence Administrator.

Activity Stream Macro

The Activity Stream macro is a specific instance of the Gadget macro. It inserts an Activity Stream gadget onto your page. For instructions, see Activity Stream Gadget.

Related pages:

- Activity Stream Gadget
- Gadget Macro

Anchor Macro

Allows you to link to a specific part of a page. For more information on using anchors see Anchors .

Code Examples

The following examples are provided for advanced users who want to edit the underlying markup for a Confluence page.

Macro name: `anchor`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="anchor">
  <ac:parameter ac:name="">here</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{anchor:here}
```

Attachments Macro

The Attachments macro displays a list of files attached to a page.

It also allows users (with appropriate permissions) to:

- upload a file to the page, directly from the list
- edit attachment properties and labels
- delete an attached file (this deletes all versions of the file)
- preview image attachments
- edit attached Office and PDF documents using the Office Connector.
- download all files attached to the page.

Note: you can use the macro parameters to turn off previews if you have very large attachments.

Screenshot: The Attachments macro, showing details of an attachment

Using the Attachments Macro

To add the Attachments macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Filename Patterns (`patterns`)	all	A comma-separated list of regular expressions, used to filter the attachments by file name. Note that the parameter values must be regular expressions. For example: • To match a file suffix of 'jpg', use `.*jpg` (not `*.jpg`). • To match file names ending in 'jpg' or 'png', use `.*jpg,.*png`
Attachment Labels (`labels`)	(none)	A list of labels, used to filter the attachments to display. If you wish to enter more than one label, separate the labels with commas. Confluence will show only attachments that have **all** the labels specified. (The match is an AND, not an OR.) For information on labelling the attachments, see Add, Remove and Search for Labels.
Include Old Attachment Versions (`old`)	false	A value of `true` will include previous attachment versions in the list.
Sort By (`sortBy`)	date	The sort order for attachments. Note that people viewing the page can change the sort order by clicking the column headings. Valid values are: • **date** – sorts by updated date in reverse chronological order (newest first) • **size** – sorts largest to smallest • **name** – sorts alphabetically • **created date** - sorts by creation date in reverse chronological order (newest first)
Sort Order (`sortOrder`)	ascending	Used in combination with the **Sort By** parameter, to sort the attachments in ascending or descending order.
Allow Upload (`upload`)	true	If selected, the list of attachments will include options allowing users to browse for, and attach, new files.
Page Title (`page`)	(none)	Used to display attachments from another page. If you do not enter a page title, the macro will display the files attached to the current page.
Show Previews (`preview`)	true	Used to display a preview of the attached file. If true, preview will be visible when the list item is expanded.

Code examples

The following examples are provided for advanced users who want to edit the underlying markup for a

Confluence page.

Macro name: `attachments`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="attachments">
  <ac:parameter ac:name="old">false</ac:parameter>
  <ac:parameter ac:name="patterns">.*png,.*jpg</ac:parameter>
  <ac:parameter ac:name="sortBy">name</ac:parameter>
  <ac:parameter ac:name="page">
    <ac:link>
      <ri:page ri:content-title="My page about chocolate"/>
    </ac:link>
  </ac:parameter>
  <ac:parameter ac:name="sortOrder">descending</ac:parameter>
  <ac:parameter ac:name="labels">chocolate,cookies</ac:parameter>
  <ac:parameter ac:name="upload">false</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{attachments:old=false|patterns=.*png,.*jpg|sortby=name|page=My page about
chocolate|sortorder=descending|labels=chocolate,cookies|upload=false|preview=false}
```

Blog Posts Macro

The Blog Posts macro allows you to display blog posts on a Confluence page. Clicking on a title takes you to the blog post.

Using the Blog Posts Macro

To add the Blog Posts macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Required	Default	Description

Content Type to Display (`content`)	No	**titles**	Available values: • **titles** — Display the title, creator, space, and created date stamp for each blog post. • **excerpts** — Display a short excerpt from each blog post. If the post contains an Excerpt macro, the Blog Posts macro will display the content defined in the Excerpt macro. If the post does not contain an Excerpt macro, the Blog Posts macro will display the first few sentences of the post. • **entire** - Display the whole content of each blog post.
Time Frame (`time`)	No	no limit	Specify how far back in time Confluence should look for the blog posts to be displayed. Available values: • `m` — Minutes • `h` — Hours, so `'12h'` displays blog posts created in the last twelve hours. • `d` — Days, so `'7d'` displays blog posts created in the last seven days. • `w` — Weeks
Restrict to these Labels (`label`)	No	None	Filter the results by label. The macro will display only the blog posts which are tagged with the label(s) you specify here. You can specify one or more label values, separated by a comma or a space. • To exclude content which matches a given label, put a minus sign (-) immediately in front of that label value. For example: If you specify a label value of `-badpage` you will get only content which is not labelled with 'badpage'. • To indicate that the results **must** match a given label value, put a plus sign (+) immediately in front of that label value. For example: If you specify a label value of `+superpage,+goodpage` you will get only content which has at least two labels, being 'superpage' and 'goodpage'.
Restrict to these Aut hors (`author`)	No	None	Filter the results by author. The macro will display only the blog posts which are written by the author(s) you specify here.

Restrict to these Spaces (`spaces`)	No	@self, i.e. the space which contains the page on which the macro is coded	This parameter allows you to filter content by space. The macro will display only the pages which belong to the space(s) you specify here. You can specify one or more space keys, separated by a comma or a space. • To exclude content in a specific space, put a minus sign (-) immediately in front of that space key. For example: If you specify a space key of `-BADSPACE` you will get only content which is not in the BADSPACE. • To indicate that the results **must** come from a specific space, put a plus sign (+) immediately in front of that space key. For example: If you specify a space key of `+GOODSPACE` you will get only content in GOODSPACE. (Note that this is not particularly useful, because each content item belongs to one space only. If you put a plus sign next to one space key and list other space keys too, the other space keys will be ignored.) Special values: • `@self` — The current space. • `@personal` — All personal spaces. • `@global` — All site spaces. • `@favorite` — The spaces you have marked as favourite. • `@favourite` — The same as `@favorite` above. • `@all` — All spaces in your Confluence site. • `*` — The same as `@all` above. When specifying a personal space, remember to use the tilde (~) sign in front of the username, such as `~jbloggs` or `~jbloggs@example.com`.
Maximum Number of Blog Posts (`max`)	No	15	Specify the maximum number of results to be displayed. Note that the results are sorted first, and then the maximum parameter is applied.
Sort By (`sort`)	No	creation	Specify how the results should be sorted. If this parameter is not specified, the sort order defaults to descending order (newest first) based on the creation date. Values: • **title** — Sort alphabetically by title. • **creation** — Sort by the date on which the content was added. • **modified** — Sort by the date on which the content was last updated.
Reverse Sort (`reverse`)	No	false	Select to change the sort from descending to ascending order (oldest first). Use this parameter in conjunction with the **Sort By** parameter. This parameter is ignored if the **Sort By** parameter is not specified. In storage format and wikimarkup a value of `true` changes the sort order.

Code examples

The following examples are provided for advanced users who want to edit the underlying markup for a Confluence page.

Macro name: `blog-posts`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="blog-posts">
  <ac:parameter ac:name="content">titles</ac:parameter>
  <ac:parameter ac:name="spaces">
    <ri:space ri:space-key="ds"/>
  </ac:parameter>
  <ac:parameter ac:name="author">
    <ri:user ri:userkey="12345678912345678912345678912345"/>
  </ac:parameter>
  <ac:parameter ac:name="time">4w</ac:parameter>
  <ac:parameter ac:name="sort">creation</ac:parameter>
  <ac:parameter ac:name="max">10</ac:parameter>
  <ac:parameter ac:name="labels">chocolate,cookies</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{blog-posts:content=titles|spaces=@self,ds|author=jsmith|time=4w|reverse=true|sort=
creation|max=10|label=chocolate,cookies}
```

Change-History Macro

The Change-History macro shows the history of updates made to a page: version number, author, date and comment. It displays the information inline.

Screenshot: The Change-History macro in Confluence

To add the Change-History macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Note: There are no parameters for this macro.

Code examples

The following examples are provided for advanced users who want to edit the underlying markup for a

Confluence page.

Macro name: `change-history`

Macro body: None.

Parameters: None.

Storage format example

```
<ac:structured-macro ac:name="change-history"/>
```

Wikimarkup example

```
{change-history}
```

Chart Macro

The Chart macro allows you to display a chart based on tabular data. When you add the macro to a page, you:

- supply the data to be charted by the macro as a table in the placeholder of the macro.
- edit the macro parameters in the Macro Browser to configure the format of the chart.

Screenshot 1: A table of data in the Chart macro placeholder

Screenshot 2: The resulting chart

On this page:

- Using the Chart Macro
- Parameters
- Examples
  - Pie Chart
  - Bar Chart
  - 3D Bar Chart
  - Time Series Chart
  - XY Line Chart
  - XY Area Chart
  - Area Charts
  - Gantt Chart
- Code examples
  - Chart type parameters
  - Chart display parameters
  - Chart title and label parameters
  - Chart data parameters
  - Chart colour parameters
  - Chart axis parameters
  - Pie chart Parameters
  - Chart attachment parameters
  - Wiki markup example

Using the Chart Macro

To add the Chart macro to a page:

1. Edit the page and use autocomplete to add the Chart macro. (Type '{chart}'.)
2. Enter your chart data as one or more tables in the body of the macro placeholder. (See the screenshot above.)
3. Click the macro placeholder and choose **Edit**.
4. Select a chart type using the **Type** parameter (see below).
5. Choose other parameter settings in the macro browser, as described below.
6. Click **Refresh** in the 'Preview' area, to check that the chart appears as you expect.
7. Click **Save** to add the chart to your page.
8. Click **Save** again when you are ready to save the page.

Parameters

Chart Type Parameters

These parameters determine the type of chart to display and the way the chart looks.

Parameter	Default	Description
Type	**pie**	The type of chart to display. XY charts have numerical x- and y-axes. The x values may optionally be time-based (see the **Time Series** parameter). *Standard* — pie, bar, line, area *XY Plots* — xyArea, xyBar, xyLine, xyStep, xyStepArea, scatter, timeSeries *Other* — gantt
Display Orientation	**vertical**	Applies to area, bar and line charts. • **vertical** — y-axis is vertical • **horizontal** — x-axis is vertical
Show in 3D	false	Applies to area, bar and line charts.
Stacked Values	false	Applies to area and bar charts.
Show shapes	true	Applies to line charts. Shapes are shown at each data point.
Opacity	• 75 percent for 3D charts • 50 percent for non-stacked area charts • 100 percent for all other charts	A percentage value between 0 (transparent) and 100 (opaque) that determines how opaque the foreground areas and bars are.

Display Control Parameters

Parameter	Default	Description
Width	300	The width of the chart in pixels.
Height	300	The height of the chart in pixels.
Display rendered data	false	Sets whether to display the rendered body of the macro (usually the data tables). By default, the chart data table isn't rendered. • **before** — the data are displayed before the chart. • **after** — the data are displayed after the chart.
Image format	**png**	The image format to be used for the chart. • **png** • **jpg**

Title and Label Parameters

Parameter	Default	Description
Chart Title	none	The title of the chart.
Chart Subtitle	none	A subtitle for the chart, using a smaller font than for **Title**.
Horizontal-axis Label	none	The label for the x-axis (domain).
Vertical-axis Label	none	The label for the y-axis (range).
Show Legend	true	Show a legend or key.

Data Specification Parameters

The data for the chart is taken from tables found when the macro body is rendered. These options control how this data is interpreted. By default, numeric and date values are interpreted according to the Confluence global default language (locale) formats. If conversion fails, other languages defined in Confluence will be tried. Additional conversion options can be specified using the parameters below.

Parameter	Default	Description
Tables	all first level tables	Comma separated list of table ids and/or table numbers (starting at 1) contained within the body of the macro that will be used as the data for the chart. If data tables are embedded in other tables, then table selection will be required. This occurs when more complex formatting is done (for example using section and column macros).
Columns	all columns	Comma separated list of column labels and/or column titles and/or column numbers for tables used for chart data. This applies to all tables processed. Columns are enumerated starting at 1. Column label is the text for the column in the header row. Column title is the (html) title attribute for the column in the header row.
Content Orientation	**horizontal**	**vertical** — data table columns will be interpreted as series. **horizontal** — data tables rows will be interpreted as series.
Time Series	false	true — the x values in an XY plot will be treated as time series data and so will be converted according to date formats.
Date format	Confluence language defined date formats	For time series data, the date format allows for additional customization of the conversion of data to date values. If a **Date format** is specified, it will be the first format used to interpret date values. Specify a format that matches the time series data. See simple date format.
Time Period	**Day**	The time period for time series data. Defines the granularity of how the data is interpreted. Valid values are: **Millisecond**, **Second**, **Minute**, **Hour, Day**, **Week**, **Month**, **Quarter**, **Year**.
Language	none	Use in combination with the **Country** parameter to form a locale. These additional number and date formats will be used for data conversion before the default languages. Valid values are 2 character ISO 639-1 alpha-2 codes.
Country	none	Use in combination with the **Language** parameter to form a locale. Valid values are 2 character ISO 3166 codes.
Forgive	true	- true — the macro tries to convert numeric and date values that do not totally match any of the default or user-specified formats. - false — enforce strict data format. Data format errors will cause the chart to not be produced.

Colour Parameters

Colours are specified using hexadecimal notation or HTML colour names.

Parameter	Default	Description
Background Colour	White	Background of the chart.
Border Colour	no border	Border around the chart.
Colours		Comma-separated list of colours used to customise category, sections, and series colours.

Axis Parameters

Depending on the chart type, the range and domain axis may be customised. These values are automatically generated based on the data but can be overridden by specifying one or more more of these parameters.

Parameter	Default	Description
Range Minimum Value	none	Range axis lower bound.
Range Maximum Value	none	Range axis upper bound.
Range Axis Tick Unit	none	Range axis units between axis tick marks.
Range Axis Label Angle	none	Angle for the range axis label in degrees.
Domain Axis Lower Bound	none	Only applies to XY plots. Domain axis lower bound. For a date axis, this value must be expressed in the date format specified by the **Date format** parameter.
Domain Axis Upper Bound	none	Only applies to XY plots. Domain axis upper bound. For a date axis, this value must be expressed in the date format specified by the **Date format** parameter.
Domain Axis Tick Unit	none	Only applies to XY plots. Domain axis units between axis tick marks. For a date axis, this value represents a count of the units specified in the **Time Period** parameter. The **Time Period** unit can be overridden by specifying a trailing character: y (years), M (months), d (days), h (hours), m (minutes), s (seconds), u (milliseconds).
Domain Axis Label Angle	none	Only applies to XY plots. The angle for the domain axis label, in degrees.

Category Label Position	none	Placement of the axis label text for categories. • `up45` — 45 degrees going upward • `up90` — 90 degrees going upward • `down45` — 45 degrees going downward • `down90` — 90 degrees going downward
Date Tick Mark Position	**start**	Placement of the date tick mark. • `start` — tick mark is at the start of the date period. • `middle` — tick mark is in the middle of the date period. • `end` — tick mark is at the end of the date period.

Pie Chart Parameters

Parameter	Default	Description
Pie Section Label	Show only the pie section key value	Format for how pie section labels are displayed. The format uses a string with special replacement variables: • %0% is replaced by the pie section key. • %1% is replaced by the pie section numeric value. • %2% is replaced by the pie section percent value. Example 1: "%0% = %1%" would display something like "Independent = 20" Example 2: "%0% (%2%)" would display something like "Independent (20%)"
Pie Section Explode	No exploded sections	Comma separated list of pie keys that are to be shown exploded. Note: requires jFreeChart version 1.0.3 or higher.

Attachment Parameters

These are advanced options that can be used for chart versioning, to enable automation and to improve performance. Use these options carefully! Normally, the chart image is regenerated each time the page is displayed. These options allow for the generated image to be saved as an attachment and have subsequent access re-use the attachment. This can be useful especially when combined with the Cache Pluginto improve performance. Depending on the options chosen, chart images can be versioned for historical purposes.

Parameter	Default	Description
Attachment	none	The name and location with which the chart image will be saved as an attachment. The user must be authorised to add attachments to the page specified. • ^attachmentName.png — the chart is saved as an attachment to the current page. • page^attachmentName.png — the chart is saved as an attachment to the page name provided. • space:page^attachmentName.png — the chart is saved as an attachment to the page name provided in the space indicated.
Attachment Version	new	Defines the the versioning mechanism for saved charts. • new — creates new version of the attachment. • replace — replaces all previous versions of the chart. To replace an existing attachment, the user must be authorized to remove attachments for the page specified. • keep — only saves a new attachment if an existing export of the same name does not exist. An existing attachment will not be changed or updated.
Attachment Comment	none	Comment used for a saved chart attachment.

Thumbnail	false	• true — the chart image attachment will be shown as a thumbnail.

Chart Type Parameters | Display Control Parameters | Title and Label Parameters | Data Specification
Parameters | Colour Parameters | Axis Parameters | Pie Chart Parameters | Attachment Parameters

Examples

Pie Chart

Parameters in Macro Browser		Data Table in Macro Placeholder		Rendered Chart
Type	pie	**Fish Type**	**2011**	Fish Sold 2011
Chart Title	Fish Sold 2011	Herring	9,500	
Show Legend	true	Salmon	2,900	
Content Orientation	vertical	Tuna	1,500	

Bar Chart

Parameters in Macro Browser		Data Table in Macro Placeholder			Rendered Chart
Type	bar	**Fish Type**	**2010**	**2011**	Fish Sold
Chart Title	Fish Sold	Herring	9,500	8,300	
Show Legend	true	Salmon	2,900	4,200	
		Tuna	1,500	1,500	

3D Bar Chart

Parameters in Macro Browser		Data Table in Macro Placeholder	Rendered Chart
Type	bar		
Show in 3D	true		
Opacity	50		
Show Legend	true		

Data Table in Macro Placeholder:

	2009	2010	2011
Revenue	12.4	31.8	41.1
Expense	43.6	41.8	31.1

Time Series Chart

Parameters in Macro Browser	
Type	Time Series
Date Format	MM/yyyy
Time Period	Month
Content Orientation	vertical
Range Axis Lower Bound	0
Show Legend	true

Data Tables in Macro Placeholder:

Month	Revenue	Month	Expenses
1/2011	31.8	1/2011	41.1
2/2011	41.8	2/2011	43.8
3/2011	51.3	3/2011	45.3
4/2011	33.8	4/2011	45.0
5/2011	27.6	5/2011	44.6
6/2011	49.8	6/2011	43.8
7/2011	51.8	7/2011	51.8
8/2011	77.3	8/2011	52.3
9/2011	73.8	9/2011	53.8
10/2011	97.6	10/2011	55.6
11/2011	101.2	11/2011	61.2
12/2011	113.7	12/2011	63.7

Rendered Chart:

XY Line Chart

Parameters in Macro Browser		Data Table in Macro Placeholder	Rendered Chart
Type	xyLine		
Show Legend	true		

	12	14	23
Revenue	41.1	31.8	12.4
Expense	31.1	41.8	43.6

XY Area Chart

Parameters in Macro Browser		Data Table in Macro Placeholder	Rendered Chart
Type	xyArea		
Show Legend	true		

	12	14	23
Revenue	41.1	31.8	12.4
Expense	31.1	41.8	43.6

Area Charts

Example 1

Parameters in Macro Browser		Data Table in Macro Placeholder	Rendered Chart
Type	area		

Satisfaction	2009	2010	2011
Very	20	23	34

Show Legend	true
Width	300
Height	300
Opacity	50

satisfied				
Satisfied	40	34	23	
Dissatisfied	25	26	25	
Very dissatisfied	15	17	18	

Example 2

Parameters in Macro Browser		Data Table in Macro Placeholder	Rendered Chart
Type	area		
Show Legend	true		
Width	300		
Height	300		
Stacked Values	true		

Satisfaction	2009	2010	2011
Very satisfied	12	23	31
Satisfied	1	34	36
Dissatisfied	4	6	22
Very dissatisfied	2	7	12

Gantt Chart

Parameters in Macro Browser		Data Table in Macro Placeholder	Rendered Chart
Type	gantt		
Width	300		
Height	200		

Plan	Start	End	Status
Stage 1	6/25/2013	7/10/2013	30%
Stage 2	7/13/2013	11/28/2013	40%

Columns	,,1,2,3,4
Date format	MM/dd/yyyy

	Stage 3	12/1/2013	12/25/2013	

Actual	Start	End	Status
Stage 1	6/25/2013	7/26/2013	100%
Stage 2	7/29/2013	12/01/2013	40%
Stage 3	12/10/2013	12/25/2013	

Code examples

The following examples are provided for advanced users who want to edit the underlying markup for a Confluence page.

Macro name: `chart`

Macro body: Accepts rich text, consisting of tables that hold the chart's data.

This macro recognises a large number of parameters, listed here by type for convenience.
⌄ See all parameters...

Chart type parameters

These parameters determine the type of chart to display and how the chart looks.

Parameter	Required	Default	Description
`type`	No	`pie`	The type of chart to display. XY charts have numerical x- and y-axes. The x values may optionally be time-based. See the `timeSeries` parameter. Available values: • *Standard charts* – `pie`, `bar`, `line`, `area` • *XY plots* – `xyArea`, `xyBar`, `xyLine`, `xyStep`, `xyStepArea`, `scatter`, `timeSeries` • *Other charts* – `gantt`
`orientation`	No	`vertical`	The display orientation. Applies to area, bar and line charts. Available values: • `vertical` – y-axis is vertical • `horizontal` – x-axis is vertical
`3D`	No	`false`	Show in three dimensions. Applies to area, bar and line charts.
`stacked`	No	`false`	Stacked values. Applies to area and bar charts.
`showShapes`	No	true	Applies to line charts. Shapes are shown at each data point.

opacity	No	• 75 percent for 3D charts • 50 percent for non-stacked area charts • 100 percent for all other charts	A percentage value between 0 (transparent) and 100 (opaque) that determines how opaque the foreground areas and bars are.

Chart display parameters

Parameter	Required	Default	Description
width	No	300	The width of the chart in pixels.
height	No	300	The height of the chart in pixels.
dataDisplay	No	false	Determines whether to display the body of the macro, consisting of the data table. By default, the chart data table is not displayed. Available values: • false – the data is not displayed. • true or after – the data is displayed after the chart. • before – the data is displayed before the chart.
imageFormat	No	png	The image format to be used for the chart. Available values: • png • jpg

Chart title and label parameters

Parameter	Required	Default	Description
title	No	(None)	The title of the chart.
subTitle	No	(None)	A subtitle for the chart.
xLabel	No	(None)	The label for the x-axis (domain).
yLabel	No	(None)	The label for the y-axis (range).
legend	No	false	Determines whether to show a legend (key) for the chart.

Chart data parameters

The data for the chart is taken from tables found in the macro body. The parameters below control how this data is interpreted. By default, numeric and date values are interpreted according to the Confluence global default language (locale) formats. If conversion fails, other languages defined in Confluence will be tried. You can specify additional conversion options using the parameters below.

Parameter	Required	Default	Description

tables	No	All first level tables	You can supply a comma-separated list of table IDs and/or table numbers (starting at 1) contained within the body of the macro that will be used as the data for the chart. If data tables are embedded in other tables, then table selection will be required. This occurs when more complex formatting is done (for example using section and column macros).
columns	No	All columns	You can supply a comma-separated list of column labels and/or column titles and/or column numbers for tables used for chart data. This applies to all tables processed. Columns are enumerated starting at 1. Column label is the text for the column in the header row. Column title is the HTML title attribute for the column in the header row.
dataOrientation	No	horizontal	The content orientation. By default, the data tables will be interpreted as columns (horizontally) representing domain and x values. Available values: • `vertical` – data table columns will be interpreted as series. • `horizontal` – data tables rows will be interpreted as series.
timeSeries	No	false	If 'true', the x values in an XY plot will be treated as time series data and so will be converted according date formats.
dateFormat	No	Confluence language defined date formats	For time series data, the date format allows for additional customisation of the conversion of data to date values. If a `dateFormat` is specified, it will be the first format used to interpret date values. Specify a format that matches the time series data. See simple date format.
timePeriod	No	day	The time period for time series data. Defines the granularity of how the data is interpreted. Available values: `millisecond, second, minute, hour, day, week, month, quarter, year`
language	No	(None)	Use in combination with the `country` parameter to form a locale. These additional number and date formats will be used for data conversion before the default languages. Available values are the two-character ISO 639-1 alpha-2 codes.
country	No	(None)	Use in combination with the `language` parameter to form a locale. Valid values are the two-character ISO 3166 codes.

forgive	No	true	Determines whether the macro will forgive (allow) some data formatting errors.

Available values:

- `true` — the macro tries to convert numeric and date values that do not totally match any of the default or user-specified formats.
- `false` — the macro enforces strict data formatting. If there are data format errors, the chart will not be produced.

Chart colour parameters

Colours are specified using hexadecimal notation or HTML colour names.

Parameter	Required	Default	Description
bgColor	No	White	Background colour of the chart.
borderColor	No	No border	Colour of the border around the chart.
colors	No		A comma-separated list of colours used to customise the colours of categories, sections, and series.

Chart axis parameters

Depending on the chart type, the range and domain axis may be customised. These values are automatically generated based on the data but can be overridden by specifying one or more more of these parameters.

Parameter	Required	Default	Description
rangeAxisLowerBound	No	(None)	Minimum value for the range axis.
rangeAxisUpperBound	No	(None)	Maximum value for the range axis.
rangeAxisTickUnit	No	(None)	Range axis units between axis tick marks.
rangeAxisLabelAngle	No	(None)	Angle for the range axis label in degrees.
domainAxisLowerBound	No	(None)	Only applies to XY plots. Domain axis lower bound. For a date axis, this value must be expressed in the date format specified by the `dateFormat` parameter.
domainAxisUpperBound	No	(None)	Only applies to XY plots. Domain axis upper bound. For a date axis, this value must be expressed in the date format specified by the `dateFormat` parameter.
domainAxisTickUnit	No	(None)	Only applies to XY plots. Domain axis units between axis tick marks. For a date axis, this value represents a count of the units specified in the `timePeriod` parameter. The `timePeriod` unit can be overridden by specifying a trailing character: `y` (years), `M` (months), `d` (days), `h` (hours), `m` (minutes), `s` (seconds), `u` (milliseconds).
domainAxisLabelAngle	No	(None)	Only applies to XY plots. The angle for the domain axis label, in degrees.

categoryLabelPosition	No	(None)	Placement of the axis label text for categories. Available values: • `up45` — 45 degrees going upward • `up90` — 90 degrees going upward • `down45` — 45 degrees going downward • `down90` — 90 degrees going downward
dateTickMarkPosition	No	start	Placement of the date tick mark. Available values: • `start` — tick mark is at the start of the date period. • `middle` — tick mark is in the middle of the date period. • `end` — tick mark is at the end of the date period.

Pie chart Parameters

Parameter	Required	Default	Description

pieSectionLabel	No	Show only the pie section key value	Format of pie section labels. The format uses a string with special replacement variables:
			 • `%0%` is replaced by the pie section key.
			 • `%1%` is replaced by the pie section numeric value.
			 • `%2%` is replaced by the pie section percent value.
			
			 Example 1: To display something like 'Independent = 20':
			
			 ``` %0% = %1% ```
			
			 Example 2: To display something like 'Independent (20%)':
			
			 ``` %0% (%2%) ```

pieSectionExplode	No	No exploded sections	A comma-separated list of pie keys that are to be shown exploded. Note: requires jFreeChart version 1.0.3 or higher.

Chart attachment parameters

These are advanced options that can be used for chart versioning, to enable automation and to improve performance. Use these options carefully! Normally, the chart image is regenerated each time the page is displayed. These options allow for the generated image to be saved as an attachment and have subsequent access to re-use the attachment. This can be useful especially when combined with the Cache plugin to improve performance. Depending on the options chosen, chart images can be versioned for historical purposes.

Parameter	Required	Default	Description

attachment	No	(None)	The name and location where the chart image will be saved as an attachment. The user must be authorised to add attachments to the page specified. Available syntax for this parameter: `^attachmentName.png` — the chart is saved as an attachment to the current page. `page name^attachmentName.png` — the chart is saved as an attachment to the page name provided. `spacekey:page name^attachmentName.png` — the chart is saved as an attachment to the page name provided in the space indicated.
attachmentVersion	No	new	Defines the the versioning mechanism for saved charts. Available values: `new` — creates new version of the attachment. `replace` — replaces all previous versions of the chart. To replace an existing attachment, the user must be authorised to remove attachments for the page specified. `keep` — only saves a new attachment if an existing export of the same name does not exist. An existing attachment will not be changed or updated.
attachmentComment	No	(None)	Comment used for a saved chart attachment.
thumbnail	No	false	If `true`, the chart image attachment will be shown as a thumbnail (small, expandable) image.

Storage format example

Below is a simple example of a pie chart. See more examples in Storage Format Examples for Chart Macro.

```
<ac:structured-macro ac:name="chart">
  <ac:parameter ac:name="title">Fish Sold</ac:parameter>
  <ac:parameter ac:name="type">pie</ac:parameter>
  <ac:rich-text-body>
    <table>
      <tbody>
        <tr>
          <th>
            <p>Fish Type</p>
          </th>
          <th>
            <p>2004</p>
          </th>
          <th>
            <p>2005</p>
          </th>
        </tr>
        <tr>
          <th>
            <p>Herring</p>
          </th>
          <td>
            <p>9,500</p>
          </td>
          <td>
            <p>8,300</p>
          </td>
        </tr>
        <tr>
          <th>
            <p>Salmon</p>
          </th>
          <td>
            <p>2,900</p>
          </td>
          <td>
            <p>4,200</p>
          </td>
        </tr>
        <tr>
          <th>
            <p>Tuna</p>
          </th>
          <td>
            <p>1,500</p>
          </td>
          <td>
            <p>1,500</p>
          </td>
        </tr>
      </tbody>
    </table>
  </ac:rich-text-body>
</ac:structured-macro>
```

Wiki markup example

Below is a simple example of a pie chart. See more examples in Wiki Markup Examples for Chart Macro.

```
{chart:type=pie|title=Fish Sold}
|| Fish Type || 2004 || 2005 ||
|| Herring | 9,500 | 8,300 |
|| Salmon | 2,900 | 4,200 |
|| Tuna | 1,500 | 1,500 |
{chart}
```

Confluence Storage Format Examples for Chart Macro

This page is an extension of the documentation for the Chart Macro. This page contains additional examples for the Chart macro.

Pie chart

Here is a simple example of a pie chart.

Storage format

```
{chart:type=pie|title=Fish Sold}
|| Fish Type || 2004 || 2005 ||
|| Herring | 9,500 | 8,300 |
|| Salmon | 2,900 | 4,200 |
|| Tuna | 1,500 | 1,500 |
{chart}
```

```
<ac:structured-macro ac:name="chart">
  <ac:parameter ac:name="title">Fish Sold</ac:parameter>
  <ac:parameter ac:name="type">pie</ac:parameter>
  <ac:rich-text-body>
    <table>
      <tbody>
        <tr>
          <th>
            <p>Fish Type</p>
          </th>
          <th>
            <p>2004</p>
          </th>
          <th>
            <p>2005</p>
          </th>
        </tr>
        <tr>
          <th>
            <p>Herring</p>
          </th>
          <td>
            <p>9,500</p>
          </td>
          <td>
            <p>8,300</p>
          </td>
        </tr>
        <tr>
          <th>
            <p>Salmon</p>
          </th>
          <td>
            <p>2,900</p>
          </td>
          <td>
            <p>4,200</p>
          </td>
        </tr>
        <tr>
          <th>
            <p>Tuna</p>
          </th>
          <td>
            <p>1,500</p>
          </td>
          <td>
            <p>1,500</p>
          </td>
        </tr>
      </tbody>
    </table>
  </ac:rich-text-body>
</ac:structured-macro>
```

Resulting chart

Fish Sold

■ 2004 (1,500 - 50%) ■ 2005 (1,500 - 50%)

Bar chart

Here is a simple example of a bar chart.

Storage format

```
<ac:macro ac:name="chart">
  <ac:parameter ac:name="title">Fish Sold</ac:parameter>
  <ac:parameter ac:name="type">bar</ac:parameter>
  <ac:rich-text-body>
    <table>
      <tbody>
        <tr>
          <th>
            <p>Fish Type</p>
          </th>
          <th>
            <p>2004</p>
          </th>
          <th>
            <p>2005</p>
          </th>
        </tr>
        <tr>
          <th>
            <p>Herring</p>
          </th>
          <td>
            <p>9,500</p>
          </td>
          <td>
            <p>8,300</p>
          </td>
        </tr>
        <tr>
          <th>
            <p>Salmon</p>
          </th>
          <td>
            <p>2,900</p>
          </td>
          <td>
            <p>4,200</p>
          </td>
        </tr>
        <tr>
          <th>
            <p>Tuna</p>
          </th>
          <td>
            <p>1,500</p>
          </td>
          <td>
            <p>1,500</p>
          </td>
        </tr>
      </tbody>
    </table>
  </ac:rich-text-body>
</ac:macro>
```

Resulting chart

Fish Sold

Time series chart

Here is an example of a time series chart.

Storage format

```
<ac:structured-macro ac:name="chart">
  <ac:parameter ac:name="timePeriod">Month</ac:parameter>
  <ac:parameter ac:name="dataOrientation">vertical</ac:parameter>
  <ac:parameter ac:name="rangeAxisLowerBound">0</ac:parameter>
  <ac:parameter ac:name="dateFormat">MM/yyyy</ac:parameter>
  <ac:parameter ac:name="type">timeSeries</ac:parameter>
  <ac:parameter ac:name="domainaxisrotateticklabel">true</ac:parameter>
  <ac:rich-text-body>
    <table>
      <tbody>
        <tr>
          <th>
            <p>Month</p>
          </th>
          <th>
            <p>Revenue</p>
          </th>
        </tr>
        <tr>
          <td>
            <p>1/2005</p>
          </td>
          <td>
            <p>31.8</p>
          </td>
        </tr>
        <tr>
          <td>
            <p>2/2005</p>
          </td>
          <td>
            <p>41.8</p>
          </td>
        </tr>
        <tr>
          <td>
```

```
      <p>3/2005</p>
    </td>
    <td>
      <p>51.3</p>
    </td>
  </tr>
  <tr>
    <td>
      <p>4/2005</p>
    </td>
    <td>
      <p>33.8</p>
    </td>
  </tr>
  <tr>
    <td>
      <p>5/2005</p>
    </td>
    <td>
      <p>27.6</p>
    </td>
  </tr>
  <tr>
    <td>
      <p>6/2005</p>
    </td>
    <td>
      <p>49.8</p>
    </td>
  </tr>
  <tr>
    <td>
      <p>7/2005</p>
    </td>
    <td>
      <p>51.8</p>
    </td>
  </tr>
  <tr>
    <td>
      <p>8/2005</p>
    </td>
    <td>
      <p>77.3</p>
    </td>
  </tr>
  <tr>
    <td>
      <p>9/2005</p>
    </td>
    <td>
      <p>73.8</p>
    </td>
  </tr>
  <tr>
    <td>
      <p>10/2005</p>
    </td>
    <td>
      <p>97.6</p>
    </td>
  </tr>
  <tr>
```

```
        <td>
          <p>11/2005</p>
        </td>
        <td>
          <p>101.2</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>12/2005</p>
        </td>
        <td>
          <p>113.7</p>
        </td>
      </tr>
    </tbody>
  </table>
  <table>
    <tbody>
      <tr>
        <th>
          <p>Month</p>
        </th>
        <th>
          <p>Expenses</p>
        </th>
      </tr>
      <tr>
        <td>
          <p>1/2005</p>
        </td>
        <td>
          <p>41.1</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>2/2005</p>
        </td>
        <td>
          <p>43.8</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>3/2005</p>
        </td>
        <td>
          <p>45.3</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>4/2005</p>
        </td>
        <td>
          <p>45.0</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>5/2005</p>
        </td>
      </td>
```

```
        <td>
          <p>44.6</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>6/2005</p>
        </td>
        <td>
          <p>43.8</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>7/2005</p>
        </td>
        <td>
          <p>51.8</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>8/2005</p>
        </td>
        <td>
          <p>52.3</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>9/2005</p>
        </td>
        <td>
          <p>53.8</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>10/2005</p>
        </td>
        <td>
          <p>55.6</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>11/2005</p>
        </td>
        <td>
          <p>61.2</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>12/2005</p>
        </td>
        <td>
          <p>63.7</p>
        </td>
      </tr>
    </tbody>
```

```
      </table>
    </ac:rich-text-body>
  </ac:structured-macro>
```

Resulting chart

XY line chart

Here is an example of an XY line chart.

Storage format

```
<ac:macro ac:name="chart">
  <ac:parameter ac:name="type">xyline</ac:parameter>
  <ac:rich-text-body>
    <table>
      <tbody>
        <tr>
          <th>
            <p> </p>
          </th>
          <th>
            <p>12</p>
          </th>
          <th>
            <p>14</p>
          </th>
          <th>
            <p>23</p>
          </th>
        </tr>
        <tr>
          <td>
            <p>Revenue</p>
          </td>
          <td>
            <p>41.1</p>
          </td>
          <td>
            <p>31.8</p>
          </td>
          <td>
            <p>12.4</p>
          </td>
        </tr>
        <tr>
          <td>
            <p>Expense</p>
          </td>
          <td>
            <p>31.1</p>
          </td>
          <td>
            <p>41.8</p>
          </td>
          <td>
            <p>43.6</p>
          </td>
        </tr>
      </tbody>
    </table>
  </ac:rich-text-body>
</ac:macro>
```

Resulting chart

XY bar chart

Here is an example of an XY bar chart.

Storage format

```
<ac:structured-macro ac:name="chart">
  <ac:parameter ac:name="opacity">60</ac:parameter>
  <ac:parameter ac:name="type">xybar</ac:parameter>
  <ac:rich-text-body>
    <table>
      <tbody>
        <tr>
          <th>
            <p> </p>
          </th>
          <th>
            <p>2005</p>
          </th>
          <th>
            <p>2006</p>
          </th>
          <th>
            <p>2007</p>
          </th>
        </tr>
        <tr>
          <td>
            <p>Revenue</p>
          </td>
          <td>
            <p>41.1</p>
          </td>
          <td>
            <p>31.8</p>
          </td>
          <td>
            <p>12.4</p>
          </td>
        </tr>
        <tr>
          <td>
            <p>Expense</p>
          </td>
          <td>
            <p>31.1</p>
          </td>
          <td>
            <p>41.8</p>
          </td>
          <td>
            <p>43.6</p>
          </td>
        </tr>
      </tbody>
    </table>
  </ac:rich-text-body>
</ac:structured-macro>
```

Resulting chart

XY area chart

Here is an example of an XY area chart.

Storage format

```
<ac:structured-macro ac:name="chart">
  <ac:parameter ac:name="type">xyarea</ac:parameter>
  <ac:rich-text-body>
    <table>
      <tbody>
        <tr>
          <th>
            <p> </p>
          </th>
          <th>
            <p>12</p>
          </th>
          <th>
            <p>14</p>
          </th>
          <th>
            <p>23</p>
          </th>
        </tr>
        <tr>
          <td>
            <p>Revenue</p>
          </td>
          <td>
            <p>41.1</p>
          </td>
          <td>
            <p>31.8</p>
          </td>
          <td>
            <p>12.4</p>
          </td>
        </tr>
        <tr>
          <td>
            <p>Expense</p>
          </td>
          <td>
            <p>31.1</p>
          </td>
          <td>
            <p>41.8</p>
          </td>
          <td>
            <p>43.6</p>
          </td>
        </tr>
      </tbody>
    </table>
  </ac:rich-text-body>
</ac:structured-macro>
```

Resulting chart

Area chart

Here are two examples of area charts.

Storage format for area chart 1

```
<ac:structured-macro ac:name="chart">
  <ac:parameter ac:name="height">300</ac:parameter>
  <ac:parameter ac:name="legend">true</ac:parameter>
  <ac:parameter ac:name="width">300</ac:parameter>
  <ac:parameter ac:name="opacity">50</ac:parameter>
  <ac:parameter ac:name="type">area</ac:parameter>
  <ac:parameter ac:name="dataDisplay">true</ac:parameter>
  <ac:rich-text-body>
    <table>
      <tbody>
        <tr>
          <th>
            <p>Satisfaction</p>
          </th>
          <th>
            <p>2002</p>
          </th>
          <th>
            <p>2003</p>
          </th>
          <th>
            <p>2004</p>
          </th>
        </tr>
        <tr>
          <td>
            <p>Very satisfied</p>
          </td>
          <td>
            <p>20</p>
          </td>
          <td>
            <p>23</p>
          </td>
          <td>
            <p>34</p>
```

```
        </td>
      </tr>
      <tr>
        <td>
          <p>Satisfied</p>
        </td>
        <td>
          <p>40</p>
        </td>
        <td>
          <p>34</p>
        </td>
        <td>
          <p>23</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>Disatisfied</p>
        </td>
        <td>
          <p>25</p>
        </td>
        <td>
          <p>26</p>
        </td>
        <td>
          <p>25</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>Very disatisfied</p>
        </td>
        <td>
          <p>15</p>
        </td>
        <td>
          <p>17</p>
        </td>
        <td>
          <p>18</p>
        </td>
      </tr>
    </tbody>
```

```
      </table>
    </ac:rich-text-body>
  </ac:structured-macro>
```

Resulting area chart 1

Satisfaction	2002	2003	2004
Very satisfied	20	23	34
Satisfied	40	34	23
Disatisfied	25	26	25
Very disatisfied	15	17	18

Storage format for area chart 2

```
<ac:structured-macro ac:name="chart">
  <ac:parameter ac:name="stacked">true</ac:parameter>
  <ac:parameter ac:name="height">300</ac:parameter>
  <ac:parameter ac:name="legend">true</ac:parameter>
  <ac:parameter ac:name="width">300</ac:parameter>
  <ac:parameter ac:name="type">area</ac:parameter>
  <ac:parameter ac:name="dataDisplay">true</ac:parameter>
  <ac:rich-text-body>
    <table>
      <tbody>
        <tr>
          <th>
            <p>Satisfaction</p>
          </th>
          <th>
            <p>2002</p>
          </th>
          <th>
            <p>2003</p>
          </th>
          <th>
            <p>2004</p>
          </th>
```

```
      </tr>
      <tr>
        <td>
          <p>Very satisfied</p>
        </td>
        <td>
          <p>12</p>
        </td>
        <td>
          <p>23</p>
        </td>
        <td>
          <p>31</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>Satisfied</p>
        </td>
        <td>
          <p>1</p>
        </td>
        <td>
          <p>34</p>
        </td>
        <td>
          <p>36</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>Disatisfied</p>
        </td>
        <td>
          <p>4</p>
        </td>
        <td>
          <p>6</p>
        </td>
        <td>
          <p>22</p>
        </td>
      </tr>
      <tr>
        <td>
          <p>Very disatisfied</p>
        </td>
        <td>
          <p>2</p>
        </td>
        <td>
          <p>7</p>
        </td>
        <td>
          <p>12</p>
        </td>
      </tr>
    </tbody>
```

```
            </table>
        </ac:rich-text-body>
    </ac:structured-macro>
```

Resulting area chart 2

Satisfaction	2002	2003	2004
Very satisfied	12	23	31
Satisfied	1	34	36
Disatisfied	4	6	22
Very disatisfied	2	7	12

Wiki Markup Examples for Chart Macro

This page is an extension of the documentation for the Chart Macro. This page contains additional examples for the Chart macro.

Pie chart

Here is a simple example of a pie chart.

Wiki markup

```
{chart:type=pie|title=Fish Sold}
|| Fish Type || 2004 || 2005 ||
|| Herring | 9,500 | 8,300 |
|| Salmon | 2,900 | 4,200 |
|| Tuna | 1,500 | 1,500 |
{chart}
```

Resulting chart

Fish Sold

■ 2004 (1,500 - 50%) ■ 2005 (1,500 - 50%)

Bar chart

Here is a simple example of a bar chart.

Wiki markup

```
{chart:type=bar|title=Fish Sold}
|| Fish Type || 2004 || 2005 ||
|| Herring | 9,500 | 8,300 |
|| Salmon | 2,900 | 4,200 |
|| Tuna | 1,500 | 1,500 |
{chart}
```

Resulting chart

Time series chart

Here is an example of a time series chart.

Wiki markup

```
{chart:type=timeSeries|dateFormat=MM/yyyy|timePeriod=Month|
dataOrientation=vertical|rangeAxisLowerBound=0|domainaxisrotateticklabel=true}
|| Month || Revenue ||
| 1/2005 | 31.8 |
| 2/2005 | 41.8 |
| 3/2005 | 51.3 |
| 4/2005 | 33.8 |
| 5/2005 | 27.6 |
| 6/2005 | 49.8 |
| 7/2005 | 51.8 |
| 8/2005 | 77.3 |
| 9/2005 | 73.8 |
| 10/2005 | 97.6 |
| 11/2005 | 101.2 |
| 12/2005 | 113.7 |

|| Month || Expenses ||
| 1/2005 | 41.1 |
| 2/2005 | 43.8 |
| 3/2005 | 45.3 |
| 4/2005 | 45.0 |
| 5/2005 | 44.6 |
| 6/2005 | 43.8 |
| 7/2005 | 51.8 |
| 8/2005 | 52.3 |
| 9/2005 | 53.8 |
| 10/2005 | 55.6 |
| 11/2005 | 61.2 |
| 12/2005 | 63.7 |
{chart}
```

Resulting chart

XY line chart

Here is an example of an XY line chart.

Wiki markup

```
{chart:type=xyline}
|| || 12 || 14 || 23 ||
| Revenue | 41.1 | 31.8 | 12.4 |
| Expense | 31.1 | 41.8 | 43.6 |
{chart}
```

Resulting chart

XY bar chart

Here is an example of an XY bar chart.

Wiki markup

```
{chart:type=xybar|opacity=60}
|| || 2005 || 2006 || 2007 ||
| Revenue | 41.1 | 31.8 | 12.4 |
| Expense | 31.1 | 41.8 | 43.6 |
{chart}
```

Resulting chart

XY area chart

Here is an example of an XY area chart.

Wiki markup

```
{chart:type=xyarea}
|| || 12 || 14 || 23 ||
| Revenue | 41.1 | 31.8 | 12.4 |
| Expense | 31.1 | 41.8 | 43.6 |
{chart}
```

Resulting chart

Area chart

Here are two examples of area charts.

Wiki markup for area chart 1

```
{chart:type=area|dataDisplay=true|legend=true|width=300|height=300|opacity=50}
|| Satisfaction || 2002 || 2003 || 2004 ||
| Very satisfied | 20 | 23 | 34 |
| Satisfied | 40 | 34 | 23 |
| Disatisfied | 25 | 26 | 25 |
| Very disatisfied | 15 | 17 | 18 |
{chart}
```

Resulting area chart 1

Satisfaction	2002	2003	2004
Very satisfied	20	23	34
Satisfied	40	34	23
Disatisfied	25	26	25
Very disatisfied	15	17	18

Wiki markup for area chart 2

```
{chart:type=area|dataDisplay=true|legend=true|width=300|height=300|stacked=true}
|| Satisfaction || 2002 || 2003 || 2004 ||
| Very satisfied | 12 | 23 | 31 |
| Satisfied | 1 | 34 | 36 |
| Disatisfied | 4 | 6 | 22 |
| Very disatisfied | 2 | 7 | 12 |
{chart}
```

Resulting area chart 2

Satisfaction	2002	2003	2004
Very satisfied	12	23	31
Satisfied	1	34	36
Disatisfied	4	6	22
Very disatisfied	2	7	12

Cheese Macro

The Cheese macro simply displays the words "I like cheese!" You can use this macro to test the Confluence macro functionality. 🙂

To add the Cheese macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Note: There are no parameters for this macro.

Code examples

The following examples are provided for advanced users who want to edit the underlying markup for a Confluence page.

Macro name: `cheese`

Macro body: None.

Parameters: None

Storage format example

```
<ac:structured-macro ac:name="cheese"/>
```

Wiki markup example

```
{cheese}
```

Children Display Macro

Use the Children Display macro to list the child pages of a page and the further descendants (children's children). By default, the macro displays links to the child pages as shown in the screenshot below. People viewing the page will see only the links for pages that they have permission to view.

Screenshot: The Children Display macro in Confluence

This is the Children Display macro

- Step 1 - download the game
- Step 2 - Play
- Step 3 - share your favourite nerd
- Sample page
- Games as a communication medium

Using the Children Display macro

To add the Children Display macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Example

This list of child pages is generated by a Children Display macro on this page:

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (example).

Parameter	Default	Description
Show Descendants (all)	false	Choose whether to display all the parent page's descendants. If true shows the complete tree of pages underneath the parent page, regardless of **Depth of Descendants**
Parent Page (page)	current	Specify the page to display children for, from either the current space or a different space. Enter: • '/' — to list the top-level pages of the current space, i.e. those without parents. • 'pagename' — to list the children of the specified page. • 'spacekey:' — to list the top-level pages of the specified space. • 'spacekey:pagename' — to list the children of the specified page in the specified space.

Number of Children (`first`)	none	Restrict the number of child pages that are displayed at the top level.
Depth of Descendants (`depth`)	none	Enter a number to specify the depth of descendants to display. For example, if the value is 2, the macro will display 2 levels of child pages. This setting has no effect of **Show Descendants** is enabled.
Heading Style (`style`)	none	Choose the style used to display descendants.
Include Excerpts (`excerpt`)	none	Allows you to include a short excerpt under each page in the list. Choose between: • **None** - no excerpt will be displayed • **Simple** - displays the first line of text contained in an Excerpt macro any of the returned pages. If there is not an Excerpt macro on the page, nothing will be shown. • **Rich content** - displays the contents of an Excerpt macro, or if there is not an Excerpt macro on the page, the first part of the page content, including formatted text, images and some macros.
Sort Children By (`sort`)	Manual if manually ordered, otherwise alphabetical	Optional. Choose: • **creation** — to sort by content creation date • **title** — to sort alphabetically on title • **modified** — to sort of last modification date.
Reverse Sort (`reverse`)	false	Use with the **Sort Children By** parameter. When set, the sort order changes from ascending to descending.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `children`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="children">
  <ac:parameter ac:name="reverse">true</ac:parameter>
  <ac:parameter ac:name="sort">creation</ac:parameter>
  <ac:parameter ac:name="style">h4</ac:parameter>
  <ac:parameter ac:name="page">
    <ac:link>
      <ri:page ri:content-title="Home"/>
    </ac:link>
  </ac:parameter>
  <ac:parameter ac:name="excerpt">none</ac:parameter>
  <ac:parameter ac:name="first">99</ac:parameter>
  <ac:parameter ac:name="depth">2</ac:parameter>
  <ac:parameter ac:name="all">true</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{children:reverse=true|sort=creation|style=h4|page=Home|excerpt=none|first=99|depth
=2|all=true}
```

Code Block Macro

The Code Block macro allows you to display source code in your page with syntax highlighting, like the one below:

```
public static void main(String[] args)
{
   System.out.println("Hello World!");
}
```

Add the Code Block Macro

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

You type the code block directly into the macro placeholder in the editor. Note that any white space contained in the placeholder *is not* manipulated in any way by the Code Block macro. This is to provide the writer with flexibility over code indentation.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description

Syntax highlighting (`language`)	java	Specifies the language (or environment) for syntax highlighting. The default language is Java but you can choose from one of the following languages/environments: • **actionscript3** • **bash** • **csharp** (C#) • **coldfusion** • **cpp** (C++) • **css** • **delphi** • **diff** • **erlang** • **groovy** • **html/xml** • **java** • **javafx** • **javascript** • **none** (no syntax highlighting) • **perl** • **php** • **powershell** • **python** • **ruby** • **scala** • **sql** • **vb**
Title	none	Adds a title to the code block. If specified, the title will be displayed in a header row at the top of the code block.
Collapsible (`collapse`)	false	If selected, the code macro's content will be collapsed upon visiting or refreshing the Confluence page. Clicking the `expand source` link allows you to view this content. If false, the code macro's content is always displayed in full.
Show line numbers (`linenumbers`)	false	If selected, line numbers will be shown to the left of the lines of code.
First line number (`firstline`)	1	When **Show line numbers** is selected, this value defines the number of the first line of code.
Theme	Default	Specifies the colour scheme used for displaying your code block. Many of these themes are based on the default colour schemes of popular integrated development environments (IDEs). The default theme is **Confluence** (also known as **Default**), which is typically black and coloured text on a blank background. However, you can also choose from one of the following other popular themes: • **DJango** • **Emacs** • **FadeToGrey** • **Midnight** • **RDark** • **Eclipse** • **Confluence**

Configure the Code Block macro

You can configure the Code Block macro to use a specific language and theme by default and also upload new lanuages. You need **Confluence Administrator** permissions to change the default theme and language and **System Administrator** permissions to upload new languages.

To set the default appearance of code blocks in your site:

1. Choose the **cog icon**
   ⚙ ▾
   , then choose **General Configuration** under Confluence Administration
2. Choose **Configure Code Macro**
3. Select a **Default Theme** and **Default Language**
4. Choose **Save**

All new code blocks will use the default theme and language unless you specify otherwise. Existing code blocks will be unchanged.

To add an additional language:

1. Choose the **cog icon**
   ⚙ ▾
   , then choose **General Configuration** under Confluence Administration
2. Choose **Configure Code Macro**
3. Choose **Add a new language**
4. Locate your language file and enter a **Name** for the new language (this will appear when selecting the language)
5. Choose **Add**

Language files must be correctly formatted JavaScript files and adhere to the Custom Brush syntax. You can find some examples of language files here.

To disable or remove a user-installed language:

1. Choose the **cog icon**
   ⚙ ▾
   , then choose **General Configuration** under Confluence Administration
2. Choose **Manage Add-ons**
3. Go to **User-installed Add-ons** and locate the add-on for your uploaded language - it will appear like this 'Custom Code Macro Highlighting for...'
4. Choose **Uninstall** or **Disable**

The language will no longer appear in the Code Macro.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `code`

Macro body: Accepts plain text.

Storage format example

The following example shows all parameters and a body:

```
<ac:structured-macro ac:name="code">
  <ac:parameter ac:name="title">This is my title</ac:parameter>
  <ac:parameter ac:name="theme">FadeToGrey</ac:parameter>
  <ac:parameter ac:name="linenumbers">true</ac:parameter>
  <ac:parameter ac:name="language">xml</ac:parameter>
  <ac:parameter ac:name="firstline">0001</ac:parameter>
  <ac:parameter ac:name="collapse">true</ac:parameter>
  <ac:plain-text-body><![CDATA[<b>This is my code]]></b></ac:plain-text-body>
</ac:structured-macro>
```

Below are three examples of the Code Block macro with various optional parameters used.

Description	Markup	What you will get
Code block macro with a body and no optional parameters	``` <ac:structured-macro ac:name="code"> <ac:plain-text-body><![CDATA[this is my code]]></ac:plain-text-body> </ac:structured-macro> ```	this is my code
Code block macro with a body and the optional *language* parameter defined	``` <ac:structured-macro ac:name="code"> <ac:parameter ac:name="language">html/xml</ac:parameter> <ac:plain-text-body><![CDATA[this is my code]]></ac:plain-text-body> </ac:structured-macro> ```	this is my code
Code block macro with a body and optional *title*, *line numbers* and *language* parameters defined	``` <ac:structured-macro ac:name="code"> <ac:parameter ac:name="title">This is my title</ac:parameter> <ac:parameter ac:name="linenumbers">true</ac:parameter> <ac:parameter ac:name="language">html/xml</ac:parameter> <ac:plain-text-body><![CDATA[this is my code]]></ac:plain-text-body> </ac:structured-macro> ```	**This is my title** this is my code

Wiki markup example

```
{code:title=This is my
title|theme=FadeToGrey|linenumbers=true|language=html/xml|firstline=0001|collapse=t
rue}
This is my code
{code}
```

Column Macro

Used with the Section macro to define columns on a page. See   Working with page layouts and columns and sections .

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `column`

Macro body: Accepts rich text.

Parameter name	Required	Default	Parameter description and accepted values
`width`	No	100% of the page width, divided equally by the number of columns in the section.	The width of the column. Can be specified either in pixels (for example, `400px`) or as a percentage of the available page width (for example, `50%`).

Storage format example

```
<ac:structured-macro ac:name="column">
  <ac:parameter ac:name="width">100px</ac:parameter>
  <ac:rich-text-body>
    <p>This is the content of <strong>column 1</strong>.</p>
  </ac:rich-text-body>
</ac:structured-macro>
```

Wiki markup example

```
{column:width=100px}
This is the content of *column 1*.
{column}
```

Content by Label Macro

The Content by Label macro is used to display lists of pages, blog posts or attachments that have particular labels. It's great for collecting related pages together and filtering out content that you don't want to see.

For example, you could use this macro to display a list of all pages that have the label 'feature-shipped' and include the word 'Blueprint', or to list any pages with the label 'meeting-notes' that you've been mentioned in.

Here's how the macro looks on your page:

And here's how you would set it up in the macro browser:

Using the Content by Label Macro

To add the Content by Label macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

CQL filters

CQL (Confluence Query Language) is a query language developed for Confluence, which you can use in some macros and the Confluence search. Confluence search and CQL-powered macros allow you to add filters to build up a search query, adding as many filters as you need to narrow down the search results.

Use the **Add a filter** link to add more filters to your query.

∨ Use AND, OR, and NOT operators

- For an OR search, specify multiple values in the same field.
  So to show pages with 'label-a', 'label-b' or both you'd put 'label-a' and 'label-b' in the same Label field, like this:

  Label *

  | label-a × | label-b × |

- For an AND search, add more than one filter and specify a single value in each.
  To show only pages with label-a and label-b you'd put 'label-a' in one label field, then add a second Label field to the macro, and put 'label-b' in the second one, like this:

  Label *

  | label-a × |

  Label   ⊗

  | label-b × |

  Put simply, OR values are entered in the same filter, AND values are entered in different filter.
  Only some filters support AND. If the filter doesn't support the AND operator, you won't be able to add that filter more than once.
- For a NOT search, enter a minus sign (-) before the label. This'll exclude everything with that label.

You can use the following CQL filters to build your query:

Filter	Description	Operators
Label*	Include pages, blog posts or attachments with these labels.	OR (multiple values in the same filter) AND (multiple Label filters)
With ancestor	Include pages that are children of this page. This allows you to restrict the macro to a single page tree.	OR (multiple values in the same filter)
Contributor**	Include pages or blog posts that were created or edited by these people.	OR (multiple values in the same filter)
Creator	Include items created by these people.	OR (multiple values in the same filter)
Mentioning user	Include pages and blog posts that @mention these people.	OR (multiple values in the same filter)
With parent	Include only direct children of this page (further sub-pages won't be included)	EQUALS (one page only)
In space**	Include items from these spaces.	OR (multiple values in the same filter)
Including text**	Include items that contain this text.	CONTAINS (single word or phrase)
With title	Include items that contain this text in the title.	CONTAINS (single word or phrase)

| Of type** | Include only pages, blogs or attachments. | OR (multiple values in the same filter) |

* This field is required in CQL-powered macros.

** You can add these filters in CQL-powered macros but in search they're part of the standard search filters, so they don't appear in the **Add a filter** menu.

Macro display options

These options control how the macro appears on your page.

Parameter	Default	Description
Sort by	Modified	Sort the list by title, the date it was created, or the date it was last modified. If you don't select an option, CQL default ordering by relevancy is used.
Reverse sort	False	Sort the list descending instead of ascending (Z - A, earliest - latest)
Maximum number of pages	15	Limit the number of items to include in the list.
List title	Blank	Include an optional heading for the macro
Show labels for each page	True	Show or hide the labels applied to each item
Show space name for each page	True	Show or hide the space name for each item
Display excerpts	False	Allows you to include a short excerpt under each page in the list. Choose between: • **None** - no excerpt will be displayed. • **Simple** - displays the first line of text contained in an Excerpt macro any of the returned pages. If there is not an Excerpt macro on the page, nothing will be shown. • **Rich content** - displays the contents of an Excerpt macro, or if there is not an Excerpt macro on the page, the first part of the page content, as formatted text, including images and some macros.

Content by User Macro

The Content by User macro generates a tabulated list of the content items, throughout the Confluence installation, that have been created by a specified Confluence user. The list includes all current pages, comments and spaces created by the user. Each item in the table is linked to its corresponding page, page comment or space dashboard.

Note that items for page comments contain a link to the page, followed by a second link to the comment itself, with these separated by a greater-than sign (>).

Using the Content by User macro

To add the Content by User macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**

2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Required	Default	Description
Username `(name)`	yes	none	The Confluence username for a person who has created content. Parameter is unnamed in wikimarkup.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `content-by-user`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="content-by-user">
  <ac:parameter ac:name="">
    <ri:user ri:userkey="12345678912345678912345678912345"/>
  </ac:parameter>
</ac:structured-macro>
```

Wikimarkup example

```
{content-by-user:jsmith}
```

Content Report Table Macro

The Content Report Table macro displays a set of pages and blog posts in tabular format, based on the labels specified in the macro parameters.

A working example

Below is a working example of the Content Report Table macro, displaying content with the label 'LDAP'.

Title	Creator	Modified
Connecting to LDAP or JIRA applications or Other Services via SSL	Nick Faiz	Oct 07, 2015
Requesting Support for External User Management	Matt Ryall	Oct 07, 2015

User Management Limitations and Recommendations	Sarah Maddox [Atlassian]	Oct 07, 2015
Diagrams of Possible Configurations for User Management	Sarah Maddox [Atlassian]	Oct 07, 2015
Configuring User Directories	Sarah Maddox [Atlassian]	Oct 07, 2015
Synchronising Data from External Directories	Sarah Maddox [Atlassian]	Aug 18, 2015
Managing Nested Groups	Sarah Maddox [Atlassian]	Aug 18, 2015
Connecting to an Internal Directory with LDAP Authentication	Sarah Maddox [Atlassian]	Aug 18, 2015
Configuring the LDAP Connection Pool	Sarah Maddox [Atlassian]	Aug 18, 2015
Connecting to an LDAP Directory	Sarah Maddox [Atlassian]	Aug 18, 2015
Managing Multiple Directories	Sarah Maddox [Atlassian]	Jan 28, 2015
Configuring an SSL Connection to Active Directory	Sarah Maddox [Atlassian]	Sep 25, 2013

Using the Content Report Table Macro

To add the Content Report Table macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Required	Default	Description
Label(s) `(labels)`	Yes	None	This parameter is required. Specify one or more labels, separated by a comma. The macro will display the content tagged with any of the label(s) specified here. For example, if you specify labels 'A' and 'B', the macro will display all pages that have the label 'A', and all pages that have the label 'B', and all pages that have both those labels.

Space(s) **(`spaces`)**	No	(All spaces)	Specify one or more space keys, separated by a comma or a space. The macro will display only the content which belongs to the space(s) specified here. When specifying a personal space, remember to use the tilde (~) sign in front of the username, such as `~jbloggs` or `~jbloggs@example.com`.
Maximum Number of Pages **(`maxResults`)**	No	20	Define the maximum number of pages that the macro will show in a single set of results. If there are more pages to be shown, the macro will display a link labelled 'Find more results'. People viewing the page can choose the link to go to a search view, which shows all pages tagged with the specified label(s). Which pages will appear? Before displaying the results, Confluence will sort them by the date the page was last modified. The most-recently created/updated pages will appear first.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `content-report-table`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="content-report-table">
  <ac:parameter ac:name="spaces">
    <ri:space ri:space-key="DOC"/>
  </ac:parameter>
  <ac:parameter ac:name="labels">LDAP</ac:parameter>
  <ac:parameter ac:name="analytics-key">meeting-notes</ac:parameter>
  <ac:parameter ac:name="maxResults">5</ac:parameter>
</ac:structured-macro>
```

Wiki markup

Wiki markup is not available for this macro. You cannot add this macro via wiki markup.

Contributors Macro

The Contributors macro displays a list of Confluence users who have made a contribution of some type to a page. It can also be used to list watchers of the page.

The scope of this macro can be extended to include the immediate children or descendants of the specified page. The list of contributors can be based on people who have:

- authored or edited the page(s)
- contributed comments or added labels to the page(s), or
- are simply watching the page(s).

Screenshot: Example list of Contributors

This is the Contributors macro

- Rach Admin
- Josh User
- Sophie Staunton

In this example, the **Display Format** parameter has been set to **list**.

Using the Contributors Macro

To add the Contributors macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Contribution Type (`include`)	authors	Filters by either the type of contribution made to a page (and optionally its descendant pages), or the watches on the page. Contribution types are: • `authors` - includes people who created or have edited the page(s) • `comments` - includes people who have added comments to the page(s) • `labels` - includes people who have added labels to the page(s) • `watches` - includes people who are watching the page(s). You can specify one or more contribution types, separated by commas.
Sort By (`order`)	count	Specifies the criteria used to sort contributors. Sort criteria are: • **count** - sorts people based on the total number of contributions to the page(s) • **name** - sorts people into alphabetical order • **update** - sorts people based on the date of their last contribution to the page(s).
Reverse Sort (`reverse`)	false	Reverses the sort order of contributors in the list. Must be used in conjunction with the **Sort By** parameter.
Maximum Number of Contributors (`limit`)	no limit	Limits the number of contributors in the list. If a number is not specified, all contributors are included.

Display Format **(mode)**	inline	Sets how the list of contributor's names is formatted: • **inline** — a comma-separated list • **list** — a bullet list.
Show Anonymous Contributions? **(showAnonymous)**	false	Sets whether to include those who contributed anonymously to a page.
Show Count? **(showCount)**	false	Sets whether to show the number of times each person made a contribution of the specified **Contribution Type**.
Show Last Contribution Time? **(showLastTime)**	false	Sets whether to show the last time each person made a contribution of the specified **Contribution Type**.
Page Name **(page)**	current	Specifies the page to use when generating the list of contributors. If **Page Name** and **Space(s)** are left blank, the current page is assumed.
Label(s) **(labels)**	none	Filters the list of contributors to those who created the specified labels from a page. You can specify one or more labels, separated by commas.

Space(s) **(spaces)**	current	Specifies the space key of the Confluence space that contains the page set in **Page Name** or alternatively, specifies the spaces to search. Space keys are case-sensitive.
		
		This parameter also takes special values, including:
		
		• `@global` — All site spaces.
		• `@personal` — All personal spaces.
		• `@all` — All spaces in your Confluence site.
		
		   You can specify one or more space keys or special values, separated by commas.
		
		   If no **Page Name** and **Label(s)** are specified, all pages from the specified set of spaces are included.

Content Type **(contentType)**	both pages and blog posts	Restricts the content type to use when generating the list of contributors: • **pages** — pages • **blogposts** — blog posts.
Blog Post Date **(publishDate)**	none	Specifies the publish date for a blog post. The date format required is: YYYY/MM/DD.
Include Page Hierarchy **(scope)**	specified page only	Specifies additional pages to include when generating the list of contributors: • **children** — just the child pages of the specified page • **descendants** — all descendants of the specified page.
Show Selected Pages **(showPages)**	false	Sets whether to show a list of the pages used to generate the list of contributors.
Custom "None Found" Message **(noneFoundMessage)**	default message	Specifies the message to be used to override the default message that is displayed when no contributors are found.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `contributors`

Macro body: None.

Storage format example

This example specifies a content type of blog posts:

```
<ac:macro ac:name="contributors">
    <ac:parameter ac:name="limit">10</ac:parameter>
    <ac:parameter ac:name="spaces">ds,@personal</ac:parameter>
    <ac:parameter ac:name="reverse">true</ac:parameter>
    <ac:parameter ac:name="labels">chocolate,cake</ac:parameter>
    <ac:parameter ac:name="showPages">true</ac:parameter>
    <ac:parameter ac:name="noneFoundMessage">Oh dear, no contributors
found</ac:parameter>
    <ac:parameter ac:name="showCount">true</ac:parameter>
    <ac:parameter ac:name="contentType">blogposts</ac:parameter>
    <ac:parameter ac:name="include">authors,comments,labels,watches</ac:parameter>
    <ac:parameter ac:name="mode">list</ac:parameter>
    <ac:parameter ac:name="showAnonymous">true</ac:parameter>
    <ac:parameter ac:name="order">update</ac:parameter>
    <ac:parameter ac:name="showLastTime">true</ac:parameter>
    <ac:parameter ac:name="publishDate">2012/06/30</ac:parameter>
  </ac:macro>
```

Wikimarkup example

This example specifies a content type of blog posts:

```
{contributors:limit=10|spaces=ds,@personal|reverse=true|labels=chocolate,cake|showP
ages=true|noneFoundMessage=Oh dear, no contributors
found|showCount=true|contentType=blogposts|include=authors,comments,labels,watches|
mode=list|showAnonymous=true|order=update|showLastTime=true|publishDate=2012/06/30}
```

This example specifies a content type of pages:

```
{contributors:limit=10|spaces=ds,@personal|reverse=true|scope=descendants|labels=ch
ocolate,cake|showPages=true|noneFoundMessage=Oh dear, no contributors
found|showCount=true|contentType=pages|include=authors,comments,labels,watches|mode
=list|showAnonymous=true|order=update|page=ds:Advanced Topics|showLastTime=true}
```

Contributors Summary Macro

The Contributors Summary macro displays a table of contribution-based statistics for a set of pages. These statistics can be grouped according to individual pages or individual contributors.

The default scope for this macro is an individual page, but this can be extended to include the immediate children or descendants of a specified page. The statistics cover the following types of contributions:

- edits to the page(s)
- comments added to the page(s)
- labels added to the page(s)
- people watching the page(s)

A simple example of the Contributors Summary macro is shown in the screenshot below. It lists statistics for the number of times each contributor has edited, added comments and added labels to this page.

Screenshot: Example Contributors Summary table of statistics

User	Edits	Comments	Labels
Sarah Maddox	2	2	0
Edwin Dawson	2	0	0
Giles Gaskell	2	0	0
Rosie Jameson	2	0	0
Andrew Lui	2	0	1
Charles Miller	1	0	0
Andrew Prentice	0	0	1

In this example, all default parameter settings are used. For more information about this macro's parameters, refer to the Parameters section below.

Using the Contributors Summary Macro

To add the Contributors Summary macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Group By (`groupby`)	contributors	Specifies the basis for grouping contribution-based statistics: • **contributors** — group by the people who have contributed • **pages** — group by the pages used to find contributors.

Columns to Display (`columns`)	edits,comments,labels	Sets the columns that should appear in the table. The statistics or type of information presented depends on the basis for grouping set with the **Group By** parameter. Statistics may be calculated for: • `edits` — the number of times each contributor has edited the page(s) or the number of edits made to each page. • `edited` — a list of the pages edited by each contributor or a list of contributors who have edited each page. • `comments` — the number of times each contributor has added comments to the page(s) or the number of comments on each page. • `commented` — a list of pages to which each contributor has added comments or a list of contributors who have commented on each page. • `labels` — the number of times each contributor has added labels to the page(s) or the number of labels on each page. • `labeled` — a list of pages to which each contributor has added labels or a list of contributors who have added a label to each page. • `labellist` — a list of labels either added by each contributor or on each page. • `watches` — the number of pages being watched by each contributor/person or the number of contributors/people watching each page. • `watching` — a list of pages being watched by each contributor/person or a list of contributors/people watching each page. • `lastupdate` — the last time each contributor made an update or when each page was last updated. Valid updates can include edit, comment or label modifications to a page. One or more columns can be used.
Sort By (`order`)	edits	Sets the criterion used for sorting items in the table. The items sorted depend on the basis for grouping set with the **Group By** parameter. Sort criteria are: • `edits` — sorts items in the table based on the total number of edits made, either by a contributor or to a page. • `name` — sorts items in the table in alphabetical order, either by contributor or page name. • `editTime` — sorts items in the table based on when the contributor last edited a page (or a specified set of pages) or when the page was lasted edited. • `update` — sorts items in the table based on when the contributor last made any contribution to a page (or a specified set of pages) or when the page last had any contribution made to it.
Reverse Sort (`reverse`)	false	Reverses the sort order of items in the table, as specified using the **Sort By** parameter. (Used only in conjunction with the **Sort By** parameter.)
Maximum Number of Items (`limit`)	no limit	Limits the number of contributors or pages in the table to the value specified. If no number is specified, all items are included.

Show Anonymous Contributions? (`showAnonymous`)	false	Includes individuals who have made anonymous contributions to a page.
Show Zero Counts? (`showZeroCounts`)	false	Sets whether contributors or pages are included for which a calculated statistic is zero.
Page Name (`page`)	current	Sets the page for which to calculate the contribution-based statistics. If no values for **Page Name** and **Space(s)** are specified, the current page is assumed.
Label(s) (`labels`)	none	Restricts the contribution-based statistics to the specified labels only. You can specify one or more labels, separated by commas.
Space(s) (`spaces`)	current	Specifies the space key of the Confluence space which contains the specified page name or alternatively, specifies a scope of spaces to search. Space keys are case-sensitive. This parameter also takes special values, including: • `@global` — All site spaces. • `@personal` — All personal spaces. • `@all` — All spaces in your Confluence site. You can specify one or more space keys or special values, separated by commas. If no **Page Name** and **Label(s)** are specified, all pages from the specified set of spaces are included.
Content Type (`contentType`)	both pages and blog posts	Restricts page types to either pages (**pages**) or blog posts (**blogposts**). If no value is specified in the Macro Browser, both pages and blog posts are included. Available values `pages` and `blogposts`.
Blog Post Date (`publishDate`)	none	Specifies the publish date for a blog post. The date format required is: YYYY/MM/DD.
Include Page Hierarchy (`scope`)	specified page only	Includes either the immediate **children** or all **descendants** of the specified page. If no value is indicated in the Macro Browser, only the specified page is included.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `contributors-summary`

Macro body: None.

Storage format example

This example specifies a content type of blog posts:

```
<ac:structured-macro ac:name="contributors-summary">
  <ac:parameter ac:name="limit">10</ac:parameter>
  <ac:parameter ac:name="spaces">
    <ri:space ri:space-key="ds"/>
    <ri:space ri:space-key="@personal"/>
  </ac:parameter>
  <ac:parameter ac:name="reverse">true</ac:parameter>
  <ac:parameter ac:name="showAnonymous">true</ac:parameter>
  <ac:parameter ac:name="order">update</ac:parameter>
  <ac:parameter ac:name="labels">chocolate,cake</ac:parameter>
  <ac:parameter ac:name="columns">edits,comments,labels,lastupdate</ac:parameter>
  <ac:parameter ac:name="groupby">pages</ac:parameter>
  <ac:parameter ac:name="contentType">blogposts</ac:parameter>
  <ac:parameter ac:name="showZeroCounts">true</ac:parameter>
  <ac:parameter ac:name="publishDate">2012/06/07</ac:parameter>
</ac:structured-macro>
```

This example specifies a content type of pages:

```
<ac:structured-macro ac:name="contributors-summary">
  <ac:parameter ac:name="limit">10</ac:parameter>
  <ac:parameter ac:name="spaces">
    <ri:space ri:space-key="ds"/>
    <ri:space ri:space-key="@personal"/>
  </ac:parameter>
  <ac:parameter ac:name="reverse">true</ac:parameter>
  <ac:parameter ac:name="showAnonymous">true</ac:parameter>
  <ac:parameter ac:name="scope">descendants</ac:parameter>
  <ac:parameter ac:name="order">update</ac:parameter>
  <ac:parameter ac:name="page">
    <ac:link>
      <ri:page ri:content-title="Advanced Topics" ri:space-key="ds"/>
    </ac:link>
  </ac:parameter>
  <ac:parameter ac:name="labels">chocolate,cake</ac:parameter>
  <ac:parameter ac:name="columns">edits,comments,labels,lastupdate</ac:parameter>
  <ac:parameter ac:name="groupby">pages</ac:parameter>
  <ac:parameter ac:name="contentType">pages</ac:parameter>
  <ac:parameter ac:name="showZeroCounts">true</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

This example specifies a content type of blog posts:

```
{contributors-summary:limit=10|spaces=ds,@personal|reverse=true|showAnonymous=true|
order=update|labels=chocolate,cake|columns=edits,comments,labels,lastupdate|groupby
=pages|contentType=blogposts|showZeroCounts=true|publishDate=2012/06/07}
```

This example specifies a content type of pages:

```
{contributors-summary:limit=10|spaces=ds,@personal|reverse=true|showAnonymous=true|
scope=descendants|order=update|page=ds:Advanced
Topics|labels=chocolate,cake|columns=edits,comments,labels,lastupdate|groupby=pages
|contentType=pages|showZeroCounts=true}
```

Create from Template Macro

The Create from Template macro displays a button on a page, linked to a specific template. When someone clicks the button, the macro opens the editor, ready to add a new page, and adds content to the page based on the given template.

When adding the macro to the page, you can specify a blueprint or a user-created template in the macro. You will also specify the name of the button displayed, and the space in which the new page will appear.

Example

Screenshot: A page with three buttons, all displayed by the 'Create from Template' macro

Using the Create from Template Macro

To add the Create from Template macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Required	Default	Description
Button Text		'Create from Template'	The description that people will see when viewing this macro on the page.
Template Name		None	Select the template or blueprint to base the new page on. Only global and user-created templates for the current space appear (unless you have specified a different space in the 'Space Key' field).
Template Title		Blank	Specify a default title for pages created using this macro (optional). You can include @currentDate, @spaceName and @spaceKey variables in the title.
Space Key		The space where the current page is located	Supply the unique space identifier (space key), to determine where the new page will be created when someone uses this macro to create a page.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `create-from-template`

Macro body: None.

Parameters for storage format differ from those available in the macro browser as follows.

Parameter name	Required	Default	Parameter description and accepted values
`templateId`	Required for user-created templates	(None)	The ID of a user-created template. This is the u... Confluence assigns when you create a template `630593`. To find the ID of a template, edit the te... the URL in your browser. The template ID is giv... parameter named `entityId`.
`blueprintModuleCompleteKey`	Required for blueprints	(None)	
`templateName`	Yes	(None)	The ID of a user-created template (for example, qualified name of the add-on that defines the bl... example, `com.atlassian.confluence.plu...` `e-shared-files-plugin:file-list-blu...`
`buttonLabel`	Yes	'Create from Template'	The description that people will seeing when vie... the page.
`spaceKey`	Yes	The space where the current page is located	The unique space identifier, to determine where... be created when someone uses this macro to c...
`Title`	No	(None)	The title for pages created using this macro. Yo... `rentDate`, `@spaceName` and `@spaceKey` vari... This title will override any title specified in a blu...

Storage format example

This example specifies a user-created template:

```
<ac:structured-macro ac:name="create-from-template">
  <ac:parameter ac:name="templateId">299630593</ac:parameter>
  <ac:parameter ac:name="buttonLabel">Blitz test</ac:parameter>
  <ac:parameter ac:name="spaceKey">
    <ri:space ri:space-key="DOCTHEME"/>
  </ac:parameter>
  <ac:parameter ac:name="templateName">299630593</ac:parameter>
</ac:structured-macro>
```

This example uses a blueprint:

```
<ac:structured-macro ac:name="create-from-template">
  <ac:parameter
ac:name="blueprintModuleCompleteKey">com.atlassian.confluence.plugins.confluence-sh
ared-files-plugin:file-list-blueprint</ac:parameter>
  <ac:parameter ac:name="buttonLabel">Shared files</ac:parameter>
  <ac:parameter ac:name="spaceKey">
    <ri:space ri:space-key="DOCTHEME"/>
  </ac:parameter>
  <ac:parameter
ac:name="templateName">com.atlassian.confluence.plugins.confluence-shared-files-plu
gin:file-list-blueprint</ac:parameter>
</ac:structured-macro>
```

Wikimarkup

Wiki markup is not available for this macro. You cannot add this macro via wiki markup.

Create Space Button Macro

The Create Space Button macro displays a create space icon that links to the 'Create Space' page. To see this icon, viewers need the 'Create Space' permission which is assigned by a site administrator.

Using the Create Space Button macro

To add the Create Space Button macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Icon Size `size`	large	Specify whether to use large or small icon. Available values: • **large** • **small**

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `create-space-button`

Macro body: None.

The following additional paramteters are available in storage format and wikimarkup.

Parameter name	Required	Default	Parameter description and accepted values
`width`	No	Natural size of icon (1:1 pixel ratio)	The width of the icon to be displayed, specified in pixels. Confluence will stretch or shrink the width of the icon to the number of pixels specified. Note: This parameter is not available via the macro browser.
`height`	No	Natural size of icon (1:1 pixel ratio)	The height of the icon to be displayed, specified in pixels. Confluence will stretch or shrink the height of the icon to the number of pixels specified. Note: This parameter is not available via the macro browser.

Storage format example

```
<ac:structured-macro ac:name="create-space-button">
  <ac:parameter ac:name="size">small</ac:parameter>
  <ac:parameter ac:name="height">50px</ac:parameter>
  <ac:parameter ac:name="width">50px</ac:parameter>
</ac:structured-macro>
```

Wikimarkup example

```
{create-space-button:size=small}
{create-space-button:height=50px|width=50px}
```

Excerpt Include Macro

The Excerpt Include macro is used to display 'excerpted' (that is, a segment of) content from one page in another.

Before you can use this macro, the excerpt must have been defined using the Excerpt macro. Note that you can have more than one Excerpt Include macro on a page (although you can have only one Excerpt macro on a page).

Example

The paragraph below shows an example of an Excerpt Include macro, containing content from an excerpt which we have defined on the Excerpt Macro page. On the Excerpt Include macro below, we have set the options to show both the title of the page and the panel surrounding the content.

Define a part of a page as the page's 'excerpt' which can then be displayed in another page.

Using the Excerpt Include Macro

To add the Excerpt Include macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Page Containing the Excerpt `(default-parameter)`	none	Type the name of the page that contains the excerpt to be displayed. You can use an excerpt from a page in the same space or another space in the same wiki.
		When you type the name of the page into the Excerpt Include macro dialog, Confluence will offer a list of matching pages, including those from other spaces.
		Alternatively, you can type the space key followed by a colon (:) and the page name, like this:
		`SPACEKEY:Page name`
		This parameter is unnamed in wikimarkup.
Remove Surrounding Panel `(nopanel)`	false	Determines whether Confluence will display a panel around the excerpted content. The panel includes the title of the page containing the excerpt, and the border of the panel. By default, the panel and title are shown.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `excerpt-include`

Macro body: None.

Storage format example

```
<ac:macro ac:name="excerpt-include">
  <ac:parameter ac:name="nopanel">true</ac:parameter>
  <ac:default-parameter>My page name</ac:default-parameter>
</ac:macro>
```

Wikimarkup example

```
{excerpt-include:My page name|nopanel=true}
```

Excerpt Macro

The Excerpt macro is used to mark a part of a page's content for re-use. Defining an excerpt enables other

macros, such as the Excerpt Include and Blog Posts macros, to display the marked content elsewhere.

You can only define one excerpt per page. In other words, you can only add the Excerpt macro once to a page.

Using the Excerpt Macro

To add the Excerpt macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

To add reusable content to the macro:

1. Add your content inside the Excerpt macro placeholder.
2. Choose the macro placeholder to see the options panel, and select the option to **Display on new line** or to **Display inline**. The default is to display the content of the macro on a new line. If you choose the inline option, the content of the macro will form part of the same paragraph as the text preceding and following it. Note that this option affects only the page that contains the Excerpt macro. It does not affect any pages where the content is reused.

Screenshot: The Excerpt macro placeholder and options panel

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Hide Excerpted Content (`hidden`)	false	Controls whether the page content contained in the Excerpt macro placeholder is displayed on the page. Note that this option affects only the page that contains the Excerpt macro. It does not affect any pages where the content is reused.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `excerpt`

Macro body: Accepts rich text.

The following additional parameter is available in storage format and wikimarkup. It performs the same function as the options panel in the editor.

Parameter name	Required	Default	Parameter description and accepted values
`atlassian-macro-output-type`	No	`BLOCK`	Determines whether the content of the Excerpt macro body is displayed on a new line or inline. Available values: • `BLOCK` – Displays the content of the macro on a new line . • `INLINE` – Displays the the content of the macro as part of the same paragraph as the text preceding and following it. Note that this option affects only the page that contains the Excerpt macro. It does not affect any pages where the content is reused.

Storage format example

```
<ac:structured-macro ac:name="excerpt">
  <ac:parameter ac:name="hidden">true</ac:parameter>
  <ac:parameter ac:name="atlassian-macro-output-type">BLOCK</ac:parameter>
  <ac:rich-text-body>

<p>This is the <strong>text</strong> I want to reuse in other pages. This text is
inside an Excerpt macro.</p>
  </ac:rich-text-body>
</ac:structured-macro>
```

Wikimarkup example

```
{excerpt:hidden=true|atlassian-macro-output-type=BLOCK}
This is the *text* I want to reuse in other pages. This text is inside an Excerpt
macro.
{excerpt}
```

Expand Macro

The Expand macro displays an expandable/collapsible section of text on your page.

Here is an example:
˅ Expand me...
    This text is hidden until you expand the section.

Using the Expand Macro

To insert the Expand macro into a page using the macro browser:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open,

where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Title	Click here to expand...	Defines the text that appears next to the expand/collapse icon.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `expand`

Macro body: Accepts rich text.

Storage format example

```
<ac:structured-macro ac:name="expand">
  <ac:parameter ac:name="title">Expand me...</ac:parameter>
  <ac:rich-text-body>

<p>This text is hidden until you expand the section.</p>
  </ac:rich-text-body>
</ac:structured-macro>
```

Wikimarkup example

```
{expand:This is my message}
This text is _hidden_ until you expand it.
{expand}
```

Notes

- **Text is expanded in PDF and HTML exports.** When you export the page to PDF or HTML, the text between the macro tags is expanded so that readers can see it in the PDF and HTML versions of the page.
- **Nesting your Expand macros.** You can put one Expand macro inside another, and Confluence will correctly show and hide the contents of all Expand macros, including the nested ones.

Favourite Pages Macro

Use the Favourite Pages macro to display a list of your favourite pages on a Confluence page. You can also get to your favourite pages from the Dashboard.

Screenshot: The Favourite Pages Macro in Confluence

This is the Favourite Pages macro

Favourite Pages

Displaying pages recently added to your favourites list.

Why I loved Angry Nerds (Angry Nerds) ⭐
Which nerd are you? (Angry Nerds) ⭐
Our new logo (Angry Nerds) ⭐

Using the Favourite Pages Macro

To insert the favourite pages macro into a page using the Macro Browser:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

There are no parameters for this macro.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `favpages`

Macro body: None.

Parameters: None.

Storage format example

```
<ac:structured-macro ac:name="favpages"/>
```

Wiki markup example

```
{favpages}
```

Gadget Macro

Gadgets are small applications that can offer dynamic content. They are typically served from a web application server and can be re-used in many other web applications. In Confluence, use the Gadget macro to add gadgets to pages or blog posts.

Confluence comes bundled with a few of its own gadgets that you can add to your pages or blog posts. However, you can access additional gadgets in this list if your Confluence Administrator has:

- Installed additional gadgets in Confluence (typically as a Confluence plugin) or
- Registered gadgets served from an external web application or website (such as those from a JIRA application).

Unlike other macros, the name of each gadget macro is unique and follows the convention **"<gadget-name> macro"**, where **<gadget-name>** is the name supplied by the gadget itself.

Inserting gadgets into a Confluence page or blog post

To add a gadget to a page:

1. Edit your page or blog post.
2. Choose **Insert** > **Other Macros**.
3. Click **External Content** to see a list of gadgets configured for use in your Confluence installation.
   (Some Confluence macros like the JIRA Issues, RSS Feed, and Widget Connector macros also appear in this category because they can also access external content.)
4. Click the desired gadget to access its parameters and properties.
   Almost all gadgets allow you to set basic parameters (listed below), which appear in the macro dialog. Some gadgets also have their own set of parameters, that can only be set in the gadget itself. You can use the preview in the macro browser to access them.
5. Set the parameters to your requirements.
6. Click **Refresh** to preview your changes.
7. Click **Insert** to add the gadget to the page.

Editing gadgets on a Confluence page or blog post

To edit an existing gadget on a page or blog post:

1. Edit your page or blog post.
2. Click the Gadget macro placeholder and choose **Edit** (or double-click the placeholder).
3. Set the gadget's parameters to your requirements.
4. Click **Refresh** to preview your changes.
5. **Save** the gadget.

Standard gadget parameters

Almost all gadgets allow you to set basic parameters (listed below), which appear on the right of the macro dialog. Each gadget may also have its own set of parameters, which appear on the left of the macro dialog.

Parameter	Default	Description
Width	450 pixels	Set the width of the gadget, using one of the following conventions: - Width in pixels, using `px` or plain numbers. For example, `500px` or `500` - A percentage of the page width, using `%`. For example, `50%` - Automatic resizing of the gadget to fit 100% of the page width: `auto`
Border	true	Places a thin grey border around the gadget.

Contents of a Gadget macro

In the addition to the standard parameters, the gadget macro contains 'parameter-like' content, which represents specific property settings that are particular to each gadget. Hence they are not documented here. Typically, this content would only be changed by customising the gadget's default properties using the macro browser.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `gadget`

Macro body: None.

The following additional parameters are available in storage format and wikimarkup.

Parameter name	Required	Default	Parameter description and accepted values
`url`	Yes	(none)	This is the location of the gadget specification (XML file).
`preferences`	No	(Gadget-dependent)	Specific property settings that are particular to each gadget.

A note about editing a gadget's properties (preferences) in markup: It is possible to edit the values of these properties directly in the wiki markup or storage format. However, this will allow the entry of invalid values. If a gadget property supports a certain set of values, the macro browser will restrict the user to selecting only valid values for that property. For that reason, we recommend that you use the macro browser to edit a gadget's properties.

Storage format example

This example shows the Confluence Page gadget:

```
<ac:structured-macro ac:name="gadget">
  <ac:parameter ac:name="width">500</ac:parameter>
  <ac:parameter ac:name="border">false</ac:parameter>
  <ac:parameter
ac:name="url">rest/gadgets/1.0/g/com.atlassian.confluence.plugins.gadgets:confluenc
e-page-gadget/gadgets/confluence-page-gadget.xml</ac:parameter>
  <ac:parameter
ac:name="preferences">spaceName=Documentation&amp;spaceKey=DOC&amp;quickfind-space=
Documentation&amp;pageId=753666&amp;pageName=Documentation%20Home&amp;quickfind-pag
e=Documentation%20Home&amp;isEditable=true&amp;isConfigured=true&amp;refresh=15&amp
;showLink=false</ac:parameter>
</ac:structured-macro>
```

This example shows the Confluence News gadget:

```
<ac:structured-macro ac:name="gadget">
  <ac:parameter
ac:name="url">rest/gadgets/1.0/g/com.atlassian.confluence.plugins.gadgets:confluenc
e-news-gadget/gadgets/confluence-news-gadget.xml</ac:parameter>
  <ac:parameter ac:name="preferences"/>
</ac:structured-macro>
```

Wiki markup example

This example shows the Confluence Page gadget:

```
{gadget:width=500|border=false|url=rest/gadgets/1.0/g/com.atlassian.confluence.plug
ins.gadgets:confluence-page-gadget/gadgets/confluence-page-gadget.xml}
spaceName=Documentation&spaceKey=DOC&quickfind-space=Documentation&pageId=753666&pa
geName=Documentation%20Home&quickfind-page=Documentation%20Home&isEditable=true&isC
onfigured=true&refresh=15&showLink=false {gadget}
```

This example shows the Confluence News gadget:

```
{gadget:url=rest/gadgets/1.0/g/com.atlassian.confluence.plugins.gadgets:confluence-
news-gadget/gadgets/confluence-news-gadget.xml} {gadget}
```

Gallery Macro

The Gallery macro displays a collection of thumbnail images in a table, based on the images attached to a Confluence page. When viewing the page, a user can click a thumbnail image to zoom into a screen-sized image and then view the images as a slide show.

Overview:

- The images shown in the gallery are taken from the files attached to the Confluence page. You can also specify a different page where the attachments are located. For information about attaching images to a page, see Upload Files.
- The captions below the images are drawn from the comments on the attachments. For information about adding comments to attachments, see Upload Files.
- By default, the gallery will include all the images attached to the page. You can also exclude or include images using parameters.
- You can sort your images into a particular order.
- You can specify a title for the gallery and also configure how many columns you want for the table in which your images are displayed.

Illustration: Live example of the Gallery macro

Some office photos, and a waterfall

Using the Gallery macro

To insert the Gallery macro onto a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
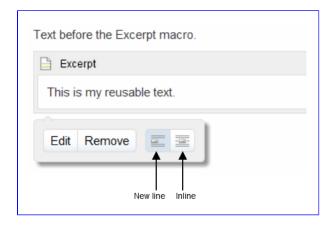2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

The full list of parameters is shown in the following table. If the name of an attached file or page contains a comma, you can refer to it in the relevant parameters below by enclosing it in single or double quotes, for example `"this,that.jpg", theother.png`

Parameter	Default	Description
Gallery Title (`title`)	Nothing	Specify a title for your gallery.
Number of Columns (`columns`)	4	Specify the number of columns for your table.

Images to Exclude (`exclude`)	No exclusions. Include all the pictures on the page.	The gallery will ignore any pictures specified. You can specify more than one picture, separated by commas. Note: The filename and filetype for this parameter are case-sensitive. For example, 'my picture.PNG' will not be recognised as 'my picture.png'.
Include these Images Only (`include`)	Include all the pictures on the page.	If you specifically include one or more pictures, the gallery will show only those pictures. You can specify more than one picture, separated by commas. Note: The filename and filetype for this parameter are case-sensitive. For example, 'my picture.PNG' will not be recognised as 'my picture.png'.
Exclude Images with these Labels (`excludeLabel`)	No exclusions. Include all the pictures on the page.	The gallery will ignore any pictures that have the specified label. You can specify more than one label, separated by commas. For information on labelling the attachments, see Add, Remove and Search for Labels.
Include Images with these Labels Only (`includeLabel`)	None. The images are not filtered by label.	Filters the images to display, based on a list of labels. If you wish to enter more than one label, separate the labels with commas. Confluence will show only images that have **all** the labels specified. (The match is an AND, not an OR.) For information on labelling the attachments, see Add, Remove and Search for Labels.
Use Images in these Pages (`page`)	If no page is specified, the gallery macro displays the images attached to the page on which the macro is used.	Specify the title of the page which contains the images you want displayed. You can specify more than one page name, separated by commas. To specify a page in a different space, use the following syntax: `SPACEKEY:Page Title`
Sort Images By (`sort`)	None. The sort order is unspecified and therefore unpredictable.	Specify an attribute to sort the images by. Sort order is ascending, unless you select the **Reverse Sort** parameter (see below). Options are: • **name** – file name. • **comment** – comment linked to the attached file. • **date** – date/time last modified. • **size** – size of the attached file.
Reverse Sort (`reverse`)	Off. Sort order is ascending	Used in combination with the **Sort Images By** parameter above. Use **Reverse Sort** to reverse the sort order, from ascending to descending. Available values in storage format and wikimarkup: • `true` – Sort order is descending. • `false` – Sort order is ascending.

Image file formats

You can attach image files of any format to a page. Confluence supports the following image formats in the Gallery macro and when displaying an image on a page:
- gif
- jpeg
- png
- bmp (depending on browser support)

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `gallery`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="gallery">
   <ac:parameter ac:name="title">My holiday pictures</ac:parameter>
   <ac:parameter ac:name="reverse">true</ac:parameter>
   <ac:parameter ac:name="sort">size</ac:parameter>
   <ac:parameter ac:name="page">My page1, ds:Welcome to Confluence</ac:parameter>
   <ac:parameter ac:name="excludeLabel">badlabel1, badlabel2</ac:parameter>
   <ac:parameter ac:name="columns">3</ac:parameter>
   <ac:parameter ac:name="exclude">badpicture.png</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{gallery:title=My holiday pictures|reverse=true|sort=size|page=My page1, ds:Welcome
to Confluence|excludeLabel=badlabel1, badlabel2|columns=3|exclude=badpicture.png}
```

Global Reports Macro

The Global Reports macro displays a list of links to some reports about content on your site.

These reports include a list of new or updated pages, orphan pages, undefined pages, and RSS feeds for new pages and blog posts.

Screenshot: The Global Reports macro

Using the Global Reports Macro

To add the Global Reports macro to a page:

1. In the Confluence editor, choose **Insert > Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open,

where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Width of Table **(width)**	99%	Specify the width of the table in which the links are displayed, as a percentage of the window width.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `global-reports`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="global-reports">
  <ac:parameter ac:name="width">50%</ac:parameter>
</ac:structured-macro>
```

Wikimarkup example

```
{global-reports:width=50%}
```

HTML Include Macro

The HTML Include macro allows you to include the contents of an external HTML file (a webpage) in a Confluence page.

> **HTML macros are disabled by default**
> The HTML macro will only be available if it has been enabled by an administrator. Enabling these macros can make your Confluence site vulnerable to cross-site scripting attacks.

Using the HTML Include Macro

To insert the HTML Include macro into a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Required	Default	Description
HTML Page's URL (`url`)	Yes	None	The URL of the page to include.

Enabling the HTML Include Macro

The HTML Include macro is disabled by default. You'll need Confluence Administrator or System Administrator permissions to enable this macro.

⚠️ Enabling these macros can make your Confluence site vulnerable to cross-site scripting attacks. You should only turn on these macros if you trust all your users not to attempt to exploit them. We strongly recommend leaving this macro disabled if you allow self-signed up or anonymous users to create content.

To enable the HTML Include macro:

1. Go to
   ⚙ ▾
   > **Add-ons**
2. Select **System** from the drop down and search for the **Confluence HTML Macros** add-on.
3. Expand the add-on and enable the **html-include (html-include-xhtml)** module.

Administrators can also choose to use the whitelist to restrict URLs that can be displayed in the HTML Include macro.

Troubleshooting

- Administrators can define a whitelist of trusted URLs. If a URL is not in the whitelist, you will see an error message in the HTML Include macro.
- You can only use the HTML Include macro for pages with absolute links. If you use the macro to include an HTML page that has relative links, you will see a 'Page Not Found' error. See

  CONF-6567 - HTML Include macro should rewrite relative links to point to remote site
  RESOLVED
  .

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `html-include`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="html-include">
  <ac:parameter ac:name="url">
    <ri:url ri:value="http://www.example.com"/>
  </ac:parameter>
</ac:structured-macro>
```

Wikimarkup example

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

HTML Macro

The HTML macro allows you to add HTML code to a Confluence page.

> **HTML macros are disabled by default**
> The HTML macro will only be available if it has been enabled by an administrator. Enabling these macros can make your Confluence site vulnerable to cross-site scripting attacks.

Using the HTML Macro

To add the HTML macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

There are no parameters for this macro.

Enter the HTML code into the body of the macro placeholder.

Enabling the HTML Macro

The HTML macro is disabled by default. You'll need Confluence Administrator or System Administrator permissions to enable this macro.

⚠ Enabling these macros can make your Confluence site vulnerable to cross-site scripting attacks. You should only turn on these macros if you trust all your users not to attempt to exploit them. We strongly recommend leaving this macro disabled if you allow self-signed up or anonymous users to create content.

To enable the HTML macro:

1. Go to
   ⚙ ▾
   > **Add-ons**.
2. Select **System** from the drop down and search for the **Confluence HTML Macros** add-on.
3. Expand the add-on and enable the **html (html-xhtml)** module.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `html`

Macro body: Text, consisting of HTML code.

Parameters: None.

Storage format example

```
< ac:structured-macro ac:name = "html" >
  < ac:plain-text-body > <![CDATA[<a href="http://www.atlassian.com">Click here</a>
to see the <b>Atlassian</b> website.]]> </ ac:plain-text-body >
```

```
</ ac:structured-macro >
```

Wikimarkup example

```
{html}Click here to see the Atlassian website.{html}
```

IM Presence Macro

The IM Presence macro indicates graphically when a contact is signed into an Instant Messaging (IM) service. The IM Presence macro appears as a small icon on the page.

Using the IM Presence Macro

To add the IM Presence macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Description
User ID/Screen Name	Identify the user by their ID, account name or screen name.
Service (`service`)	aim – AOL Instant Messenger gtalk – Google Talk icq – ICQ jabber – Jabber msn – MSN Instant Messenger sametime – IBM Lotus Sametime skype – Skype. Note: Skype requires 'Show my status on the web' to be checked under 'Privacy' preferences skypeme – Skype wildfire – Openfire Server yahoo – Yahoo! Messenger
Show User ID (`showid`)	Shows or hides the User ID of the contact.

Code examples

- The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `im`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="im">
  <ac:parameter ac:name="showid">false</ac:parameter>
  <ac:parameter ac:name="service">skype</ac:parameter>
  <ac:parameter ac:name="">MySkypeName</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{im:MySkypeName|service=skype|showid=false}
```

Include Page Macro

You can use the Include Page macro to display the contents of one Confluence page or blog post in another page or blog post.

Using the Include Page Macro

To add the Include Page macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Page to Include	None	This is the name of the Confluence page or blog post that you want to include in the current page. Start typing a page title, and Confluence will suggest matching pages from the current space and other spaces. Alternatively you can specify the page as follows: • If the page or blog post is located in another space, add the space key and a colon in front of the page name. For example, `DOC:My page name`. The space key is case sensitive. • To include a blog post, specify the date as well as the title of the blog post. For example: `/2010/12/01/My blog post`. • You can include pages from personal spaces using `~username` as the space key, where 'username' is the person's username. For example, `~jsmith:My page name`.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `include`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="include">
  <ac:parameter ac:name="">
    <ac:link>
      <ri:page ri:content-title="My chocolate page" ri:space-key="DOC"/>
    </ac:link>
  </ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{include:DOC:My chocolate page}
```

Notes

- If you want to include part of a page rather than the whole page, use the Excerpt and Excerpt Include macros.
- To display a page's contents, you need 'View' permission for that page. Similarly, people who view the page will need 'View' permissions for the embedded page as well as the page into which it is embedded. See space permissions or contact your Confluence space administrator for more information.
- If you want to embed an external page into a Confluence page, you need the HTML Include Macro.

Sample Include Page

Start of sample page content

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Aliquam fermentum vestibulum est. Cras rhoncus. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Sed quis tortor. Donec non ipsum. Mauris condimentum, odio nec porta tristique, ante neque malesuada massa, in dignissim eros velit at tellus. Donec et risus in ligula eleifend consectetuer. Donec volutpat eleifend augue. Integer gravida sodales leo. Nunc vehicula neque ac erat. Vivamus non nisl. Fusce ac magna. Suspendisse euismod libero eget mauris.

End of sample page content

Info, Tip, Note, and Warning Macros

The Info, Tip, Note, and Warning macros allow you to highlight information on a Confluence page. They create a coloured box surrounding your text, like this:

> ⓘ **Info macro example**
>
> This text is rendered inside the info macro.

> ⚠ **Note macro example**
>
> This text is rendered inside the note macro.

> ⊘ **Tip macro example**
>
> This text is rendered inside the tip macro.

> ⊙ **Warning macro example**
>
> This text is rendered inside the warning macro.

To add an Info, Tip, Note, or Warning Macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Optional Title **(title)**	none	The title of the box. If specified, the title text will be displayed in bold next to the icon.
Show information/tip/Exclamation Mark/Warning Icon **(icon)**	true	If "false", the icon will not be displayed.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `info/tip/note/warning`

Macro body: Accepts rich text.

Storage format example

The following example shows the info macro with all parameters and a body:

```
<ac:structured-macro ac:name="info">
  <ac:parameter ac:name="icon">false</ac:parameter>
  <ac:parameter ac:name="title">This is my title</ac:parameter>
  <ac:rich-text-body>
<p>
    <span>This is </span> <em>important</em> <span> information.</span>
  </p>
  </ac:rich-text-body>
</ac:structured-macro>
```

Wiki markup example

```
{info:title=This is my title|icon=false}
This is _important_ information.
{info}
```

JIRA Issues Macro

If your Confluence site is connected to a JIRA application, you can create and display JIRA issues on

Confluence pages. You can connect Confluence to any JIRA application, including JIRA Software and JIRA Service Desk.

Before you can use this macro, your Confluence and JIRA application must be connected via Application Links. People viewing the page will see the publicly accessible issues from the JIRA site. If your JIRA site has restricted viewing (that is, people need permission to view issues) then they will need to authenticate before seeing the restricted issues. See more about restricted JIRA issues below.

What you can do with the JIRA Issues macro

Using the JIRA Issues macro, you can:

- Display a list of issues on your page, based on a JIRA Query Language (JQL) search, filter or URL.
- Display a single issue.
- Display a count of issues.
- Create new issues and display the issues on your page without leaving Confluence.

Adding and updating the JIRA Issues macro – an overview

Quick guide to using the macro on a Confluence page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Screenshot: Example of JIRA issues macro on a Confluence page

Key	Summary	T	Updated ▾	P	Status	Resolution
CONF-30894	User is not prompted before overwriting changes of another user if attempting to save after an editor error is thrown	▣	Oct 22, 2013	↑	RESOLVED	Fixed
CONF-30866	TableOfContents Macros navigation is not working in Firefox 23.x / 24.x and Opera 12.16	▣	Sep 24, 2013	↓	RESOLVED	Fixed
CONF-30663	Left panel overlaps custom banner in Chrome	▣	Sep 23, 2013	↓	RESOLVED	Fixed
CONF-30690	Dashboard popular stream won't render when it contains content from an anonymous user	▣	Sep 19, 2013	↑	RESOLVED	Fixed

4 issues ↻ Refresh

Displaying issues via a JIRA Query Language (JQL) search

You can use the macro to display a table of issues on your page, based on the results of a search using JIRA Query Language (JQL).

JQL is a simple query language that is similar to SQL. A basic JQL query consists of a *field*, followed by an *oper ator* (such as = or >), followed by one or more *values* or *functions*.

Examples:

- The following query will find all issues in the 'TEST' project:

```
project = "TEST"
```

- The following query will find all issues in the 'documentation' component of the 'CONF' project:

```
project = CONF and component = documentation
```

For more information about JQL syntax, see Advanced searchingin the JIRA Software documentation.

To display a table of issues based on a JQL search:

1. Insert the JIRA Issues macro onto your Confluence page, as described above.
2. Choose a JIRA server next to the **Search** button.
3. If prompted, log in to the JIRA server.
4. Enter the JQL query into the **Search** box.
5. Choose **Search**.
6. If you want to customise the display, choose **Display options** and adjust the columns and number of issues that will appear in your table of issues.
7. Choose **Insert**.

Screenshot: Display options in the JIRA Issues macro browser.

Displaying issues via a JIRA URL

You can paste any of the following JIRA application URLs into the JIRA Issues macro. Confluence will immediately convert the URL to a JQL search.

- Any URL for an issue search or filter.
- A URL for a single issue.
- The URL of the XML view of a search.

Auto-convert: You can paste URLs directly into the Confluence editor (without calling up the macro browser). Confluence will automatically convert the URL into a JIRA Issues macro.

Displaying a single issue, or selected issues

To display a single JIRA issue, choose one of the following methods:

- Paste the URL of the issue directly onto the Confluence page. (There is no need to use the macro browser.) Confluence will auto-convert the link to a JIRA Issues macro.

- Or: Add the JIRA issues macro to the page as described above, and choose **Recently Viewed** to see the issues you have visited recently. Select an issue and choose **Insert**.
- Or: Add the JIRA issues macro to the page as described above, and paste the issue URL into the search box in the macro browser.
- Or: Add the JIRA issues macro to the page, define your search criteria in the macro browser via JQL as described above, then select the check box next to the issue in the search results, within the macro browser.

You can choose to show just the issue key, or the issue key and a summary. Select the macro placeholder and choose Show Summary or Hide Summary.

To display a subset of JIRA issues from your search results:

1. Add the JIRA issues macro to the page.
2. Define your search criteria in the macro browser via JQL, as described above.
3. Select the check boxes next to the required issues in the search results, within the macro browser.

Screenshot: Selecting a subset of issues to display

Displaying a count of issues

You can choose to display the number of issues returned by your search, rather than a table of issues. The JIRA Issues macro will display a count of issues, linked to the search in your JIRA application.

Screenshot: The JIRA Issues macro displaying an issue count on a Confluence page

To display an issue count:

1. Add the JIRA Issues macro to the page.

2. Define your search criteria in the macro browser via JQL, as described above.
3. Choose **Display options**, then choose **Total issue count** next to 'Display options' in the macro browser.
4. Choose **Insert**.

Creating a new issue

While editing a Confluence page, you can create an issue in JIRA and display it on your Confluence page, without leaving the Confluence editor.

To create an issue and add it to your page:

1. Add the JIRA Issues macro to the page, as described above.
2. Choose **Create New Issue**.
3. Supply the information about your JIRA server, project, and issue, as prompted.
4. Choose **Insert**.

Confluence will send a request to your JIRA application, to create the issue, then display the newly created issue on your page.

Limitations

The JIRA Issues macro will notify you if it is unable to create an issue in the selected project. This may be because the project has a required field, field configuration or other customisation that is not supported by the JIRA Issues macro. In this situation you will need to create the issue directly in your JIRA application.

Configuring application links to display restricted issues

Before you can use this macro, your Confluence and JIRA sites must be connected via Application Links.

If the JIRA site allows anonymous users to view issues, you must configure an application link, but there's no need to configure any incoming or outgoing authentication between the JIRA application and Confluence. People viewing the Confluence page will see the publicly accessible issues.

If your JIRA site has restricted viewing, or if some projects or issues are restricted to viewing by certain people, then people will need to log in before seeing the restricted issues.

In this situation, the *outgoing authentication* in the Confluence Application Links determines how the JIRA Issues macro handles restricted issues:

- If the outgoing authentication is set to **Trusted Applications**, people can see restricted issues if their username is the same in the JIRA application and Confluence, and if they have permission to see the issue.
- If the outgoing authentication is set to **OAuth**, people may need to choose **Login & Approve**, to gain access to the JIRA application and restricted issues.
- If the outgoing authentication is set to **Basic Access**, people can see the issues that are visible to the user account configured in JIRA's *outgoing authentication* setting.

Rendering HTML from JIRA applications

Formatted fields from JIRA can be displayed in Confluence if you set up a Confluence-to-JIRA application link. Otherwise, such formatted fields will be escaped within the output of the JIRA issues macro. This is to prevent the possibility of malicious HTML being served by an untrusted JIRA server. The most likely field where you will notice this is in the description field.

This example shows how a description column may be displayed in JIRA:

Description

This is
 	
  • the description
  • of my issue

If there is no application link between JIRA and Confluence, the description will appear in the JIRA issues macro

like this:

Description
\ This is\ \ - the description\ \ - of my issue\ \ \

Disabling the JIRA Issues macro

The functionality is provided by an add-on (plugin) called 'JIRA Macros'. To make the macro unavailable on your site, you can disable the add-on. See Disabling and enabling add-ons.

Notes

HTTPS: The JIRA Issues macro can access a JIRA application running under SSL provided the Confluence server is configured to accept the JIRA SSL certificate. See Connecting to LDAP or JIRA applications or Other Services via SSL.

Custom fields can be added as columns to the table simply by using the name of the field with no quotes. Earlier versions of the macro required you to use the custom field id, e.g. `customfield_10100`.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `jiraissues`

Macro body: None.

Note: A number of additional parameters that are not available via the macro browser are available in storage format and wikimarkup.

Parameter name	Required	Default	Parameter description and accepted values
`anonymous`	No	`false`	If this parameter is set to 'true', your JIRA application will return only the issues which allow unrestricted viewing. That is, the issues which are visible to anonymous viewers. If this parameter is omitted or set to 'false', then the results depend on how your administrator has configured the communication between the JIRA application and Confluence. By default, Confluence will show only the issues which the user is authorised to view. **Note:** This parameter is available only if you insert the macro via wiki markup or by editing the storage format of the page. The macro browser does not offer this parameter.
`baseurl`	No	The value of the 'url' parameter	If you specify a 'baseurl', then the link in the header, pointing to your JIRA application, will use this base URL instead of the value of the 'url' parameter. This is useful when Confluence connects to JIRA with a different URL from the one used by other users.

columns	No	By default, the following columns are shown: - type - key - summary - assignee - reporter - priority - status - resolution - created - updated - due	A list of JIRA column names, separated by semi-colons (;). You can include any columns recognised by your JIRA application, including custom columns.
count	No	`false`	If this parameter is set to 'true', the issue list will show the number of issues in JIRA. The count will be linked to your JIRA site.
cache	No	`on`	The macro maintains a cache of the issues which result from the JIRA query. If the 'cache' parameter is set to 'off', the relevant part of the cache is cleared each time the macro is reloaded. (The value 'false' also works and has the same effect as 'off'.) **Note:** This parameter is available only if you insert the macro via wiki markup or by editing the storage format of the page. The macro browser does not offer this parameter.
height	No	`480` (if render mode is dynamic)	The height in pixels of the table displaying the issues. Note that this height specification is ignored in the following situations: - If the 'renderMode' parameter (see below) is set to 'static'. - When the issues are displayed in a PDF or Word document, in an email message or in an RSS feed. **Note:** This parameter is available only if you insert the macro via wiki markup or by editing the storage format of the page. The macro browser does not offer this parameter.
renderMode	No	`static`	If the value is 'dynamic', the JIRA Issues macro offers an interactive display which people can manipulate as follows: - Click the column headers to sort the output. - Drag and drop the columns into a different order. - Temporarily remove a column from the display. - View a page of issues at a time, for faster response times. A value of 'static' will disable the dynamic display features. **Note:** This parameter is available only if you insert the macro via wiki markup or by editing the storage format of the page. The macro browser does not offer this parameter.
title	No	JIRA Issues	You can customise the title text at the top of the issues table with this parameter. For instance, setting the title to 'Bugs-to-fix' will replace the default 'JIRA Issues' text. This can help provide more context to the list of issues displayed. **Note:** This parameter is available only if you insert the macro via wiki markup or by editing the storage format of the page. The macro browser does not offer this parameter.

url	Yes	none	The URL of the XML view of your selected issues.
			Note: If the URL in the 'url' parameter does not contain a `tempMax` argument, then the value of `tempMax` will default to 500. If your JIRA server is version 3.12 or earlier, this means that the JIRA Issues macro will return a maximum of 500 issues. If your JIRA server is version 3.13 or later, a value of 500 means that the JIRA Issues macro will return a maximum of 500 issues per page.
width	No	100%	The width of the table displaying the issues. Can be entered as a percentage (%) or in pixels (px).
			Note: This parameter is available only if you insert the macro via wiki markup or by editing the storage format of the page. The macro browser does not offer this parameter.

Storage format example

Example using a URL that points to the XML view of a JIRA search:

```
<ac:structured-macro ac:name="jiraissues">
  <ac:parameter ac:name="anonymous">true</ac:parameter>
  <ac:parameter ac:name="columns">type;key;summary</ac:parameter>
  <ac:parameter ac:name="url">
    <ri:url
ri:value="http://jira.atlassian.com/sr/jira.issueviews:searchrequest-xml/temp/Searc
hRequest.xml?jqlQuery=project+%3D+CONF+AND+%28summary+%7E+jiraissues+OR+description
+%7E+jiraissues+OR+comment+%7E+jiraissues%29&amp;tempMax=10"/>
  </ac:parameter>
</ac:structured-macro>
```

Example using JQL – note that Confluence will insert the `server` and `serverId` parameters, based on settings in Application Links:

```
<ac:structured-macro ac:name="jira">
  <ac:parameter
ac:name="columns">key,summary,type,created,assignee,status</ac:parameter>
  <ac:parameter ac:name="server">Atlassian JIRA</ac:parameter>
  <ac:parameter
ac:name="serverId">144880e9-a1111-333f-9412-ed999a9999fa</ac:parameter>
    <ac:parameter ac:name="jqlQuery">project = CONF AND component = documentation
AND resolution = unresolved</ac:parameter>
  </ac:structured-macro>
```

Wiki markup example

```
{jiraissues:anonymous=true|url=http://jira.atlassian.com/sr/jira.issueviews:searchr
equest-xml/temp/SearchRequest.xml?jqlQuery=project+%3D+CONF+AND+%28summary+%7E+jira
issues+OR+description+%7E+jiraissues+OR+comment+%7E+jiraissues%29&tempMax=10|column
s=type;key;summary|title=My List of Issues}
```

JUnit Report Macro

The JUnit Report macro displays a summary of JUnit test results from a directory accessible by the Confluence server. JUnit is a unit testing framework which allows programmers to ensure that individual units of Java source

code are functioning correctly.For security reasons JUnit test result files can only be imported from a specified location in the Confluence server's file system. We recommend administrators create a folder in their Confluence home directory, add the system property `confluence.junit.report.directory` and specify the location for JUnit test result files to be imported from. JUnit Test result files cannot be imported from the server until this system property is set.The JUnit Report macro appears as shown in the screenshot below.

Screenshot: The JUnit Report macro in Confluence

Test			Time	Messages
MoneyTest	100%	[green bar]	00:03.391	Tests:Failures:Exceptions: 1 0 0
▷ testAdd			00:03.0	

Note: When generating reports from the JUnit Report macro, set the Apache Ant formatter to 'XML'.

Using the JUnit Report macro

To add the JUnit Report macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Required	Default	Description
URL of the test result XML file (`url`)	Must include either the `directory` or the `url` parameter	None	URL of a particular test result XML file. Is overridden by the **Directo your test result files** parameter if you use both. Example, from a Confluence instance: `http://yourConfluence om/download/attachments/<page id>/file.xml` For Confluence installations without anonymous user access, you c logon credentials as part of this parameter in the form of URL param • `os_username` — The username of a Confluence user with per access to the JUnit test results. • `os_password` — The password of the Confluence user specifie `username` parameter.
Directory (URL) of your test result files (`directory`)	Must include either the `directory` or the `url` parameter	None	URL of a directory containing your test result files. This must be a d and **not** the XML file itself. Overrides the **URL of the test result XN** eter if you use both. Example, `file:///C:/TEMP/` Example, for a network drive: `http://*host*/*path*`
Report Detail (`reportdetail`)	No	**all**	Detail for the report. Can be **all**, **fixture**, **summary** or **failuresonly**.

Debug (`debug`)	No	None	Shows the content of failures, as well as the error messages.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `junitreport`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="junitreport">
  <ac:parameter ac:name="reportdetail">summary</ac:parameter>
  <ac:parameter ac:name="directory">
    <ri:url ri:value="http://confluence.com/download/attachments/123/"/>
  </ac:parameter>
  <ac:parameter ac:name="debug">true</ac:parameter>
  <ac:parameter ac:name="url">
    <ri:url ri:value="http://confluence.com/download/attachments/123/file.xml"/>
  </ac:parameter>
</ac:structured-macro>
```

Wiki markup examples

Loading JUnit reports from a local drive:

```
{junitreport:directory=file:///C:/TEMP/}
```

Loading JUnit reports from a network drive:

```
{junitreport:url=http://*host*/*path*}
```

Loading JUnit reports from a Confluence site:

```
{junitreport:url=http://yourConfluenceInstance.com/download/attachments/<page
id>/file.xml}
```

Loading JUnit reports from a Confluence site that requires authentication:

If your Confluence site is not accessible by anonymous users, specify login credentials with the `os_username` a nd `os_password` URL parameters (as part of the macro's `url` parameter). In this case, we are specifying a username of 'admin' and a password of 'secret'.

```
{junitreport:url=http://yourConfluenceInstance.com/download/attachments/<page
id>/file.xml?os_username=admin&os_password=secret}
```

Labels List Macro

The Labels List macro displays a list of all labels within the current space. Each label in the list links to a page that displays all pages in the current space that contain that label.

Using the Labels List macro

To add the Labels List macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter name	Required	Default	Parameter description and accepted values
Restrict to this Space Key (`spaceKey`)	No	Current space	The key of the space whose labels you want to display.
Excluded label(s) (`excludedLabels`)	No	Blank	The labels that you do not want to appear in the list.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `listlabels`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="listlabels">
    <ac:parameter ac:name="spaceKey">
      <ri:space ri:space-key="DOC"/>
    </ac:parameter>
    <ac:parameter ac:name="excludedLabels">not-this-label</ac:parameter>
  </ac:structured-macro>
```

Wiki markup example

```
{listlabels:spaceKey=DOC}
```

Livesearch Macro

The Livesearch macro allows you to add a search box to a Confluence page. When users enter a search term into the search box, Confluence will dynamically display matching results as they type.

Using the Livesearch macro

To add the Livesearch macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Restrict to this Space Key (`spaceKey`)	all spaces	Specify a space key to limit the search to the given space. Case-sensitive. Alternatively, use `@self` to restrict the search to the current space.
Restrict to label(s) (`labels`)		Specify labels to limit the search to content with that label. If unspecified will search all content regardless of label.
Size (`size`)	medium	Choose a medium or large search field size.
Placeholder text (`placeholder`)		Specify the placeholder text to appear in the search field, for example 'Search this space'
Type (`type`)	all	Specify the content types to be included in the search - choose from pages, blogs, comments, space descriptions, or all content types.
Additional (`additional`)	space name	Display the space name, a page excerpt or nothing under the search result.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `livesearch`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="livesearch">
  <ac:parameter ac:name="additional">page excerpt</ac:parameter>
  <ac:parameter ac:name="placeholder">Search this space</ac:parameter>
  <ac:parameter ac:name="labels">myLabel</ac:parameter>
  <ac:parameter ac:name="spaceKey">
    <ri:space ri:space-key="SS"/>
  </ac:parameter>
  <ac:parameter ac:name="type">page</ac:parameter>
  <ac:parameter ac:name="size">large</ac:parameter>
</ac:structured-macro>
```

Wikimarkup example

```
{livesearch:spaceKey=DOC|size=large|placeholder=Search this space}
```

Loremipsum Macro

The Loremipsum macro displays paragraphs of pseudo-Latin text (more information). You can use this macro to generate more-or-less meaningless text for demonstration purposes in pages showing a draft layout or arrangement of page elements. The text is deliberately non-meaningful so that it does not influence the viewer's perception of the page arrangement or design.

A basic example of the Loremipsum text:

> Lorem ipsum dolor sit amet, consectetur adipiscing elit. Aliquam fermentum vestibulum est. Cras rhoncus. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Sed quis tortor. Donec non ipsum. Mauris condimentum, odio nec porta tristique, ante neque malesuada massa, in dignissim eros velit at tellus. Donec et risus in ligula eleifend consectetur. Donec volutpat eleifend augue. Integer gravida sodales leo. Nunc vehicula neque ac erat. Vivamus non nisl. Fusce ac magna. Suspendisse euismod libero eget mauris.

Using the Loremipsum macro

To add the Loremipsum macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Number of Paragraphs	3	Determines the amount of pseudo-Latin (space-filler) text to display. The macro will display a maximum number of 30 paragraphs.
		Parameter is unnamed in storage format and wikimarkup.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `loremipsum`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="loremipsum">
  <ac:parameter ac:name="">2</ac:parameter>
</ac:structured-macro>
```

Wikimarkup example

```
{loremipsum:2}
```

Multimedia Macro

The multimedia macro is used to embed attached video, animation and other multimedia files on a Confluence page.

Confluence supports the following multimedia formats:

- Adobe Flash (.swf)
- Apple QuickTime (.mov)
- Windows Media (.wma, .wmv)
- Real Media (.rm, .ram)
- MP3 and MP4 files (.mp3, .mp4)
- MPEG files (.mpeg, .mpg)
- AVI files (.avi) *You may need to enable an avi decoder within your browser.*

See the Widget Connector Macro if you want to display online multimedia content, such as YouTube and Vimeo videos on a page.

To add the Multimedia macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

You can also drag a supported multimedia file directly on to the page. Confluence will attach the file and add the macro for you.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter name	Required	Default	Parameter description and accepted values

Page Name `(page / space)`	Yes	Current page	Name of the page to which the multimedia file is attached. Start typing the name of the page and then select it from list of suggested pages. Include the spacekey if you want to specify a page in another space (for example `MYSPACE:My Page Title`)
Attachment `(name)`	Yes	None	File name of the attached multimedia file.
Width	No	If not specified, the browser will determine the width based on the file type.	Width of the movie window to be displayed on the page. By default, this value is specified in pixels. You can also choose to specify a percentage of the window's width, or any other value accepted by HTML.
Height	No	If not specified, the browser will determine the height based on the file type.	Height of the movie window to be displayed on the page. By default, this value is specified in pixels. You can also choose to specify a percentage of the window's height, or any other value accepted by HTML.
Autoplay `(autostart)`	No	`false`	If the parameter is set to `true` then the video or audio file will start playing as soon as the page is loaded. If this option is set to `false` then the file will not play until the user clicks the icon or image on the page.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `multimedia`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="multimedia">
  <ac:parameter ac:name="width">500</ac:parameter>
  <ac:parameter ac:name="name">
    <ri:attachment ri:filename="Ninjas.MOV"/>
  </ac:parameter>
</ac:structured-macro>
```

Wikimarkup example

```
{multimedia:space=DOC|page=My macros|name=ninjas.swf|autostart=true}
```

Notes

- **You will need the relevant multimedia plugin for your browser.** Your browser may need a plugin to play the video or audio file on a Confluence page.

- **Autoplay may not always work as expected.** Some browsers may not autoplay the attached file.

- If you get the error, 'Unable to embed content of type application/octet-stream', this means the MIME type is not recognised.

- **Advanced users can try styling via CSS.** By default, each embedded object is wrapped in a `div` tag. If you wish to style the `div` and its contents, override the `embeddedObject` CSS class. Specifying an ID as a property also allows you to style different embedded objects differently. CSS class names in the format `embeddedObject-ID` are used.

Navigation Map Macro

The Navigation Map macro displays a navigable map of the pages tagged with a given label.

Using the Navigation Map Macro

To add the Navigation Map macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Label	none	Specify the label associated with the pages you want to show in the navigation map. This parameter is unnamed in storage format and wikimarkup.
Map Title `(title)`	none	Specify a title for the navigation map.
Number of Cells Per Row `(wrapAfter)`	5	Specify the number of cells in a row
Cell Width (Pixels) `(cellWidth)`	90	Specify the cell width (enter a number only, don't include `px`)
Cell Height (Pixels) `(cellHeight)`	60	Specify the cell height (enter a number only, don't include `px`)

Navigation Map Theme (`theme`)	Confluence	Define a theme for the navmap (Server only). If you want to create your own navmap 'look and feel' (for example, one with rounded corners), you need to add a customised navmap macro theme file to the `WEB-INF/classes/templates/macros` directory. The file name convention to use is `navmap-mytheme.vm`. Use the name of your choice for the `mytheme` part of the file name, which is also the value you use for this parameter. Hence, if your theme was called `navmap-roundededges.vm`, use the value of `roundededges` for this parameter.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `navmap`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="navmap">
  <ac:parameter ac:name="wrapAfter">4</ac:parameter>
  <ac:parameter ac:name="title">My map name</ac:parameter>
  <ac:parameter ac:name="cellHeight">50px</ac:parameter>
  <ac:parameter ac:name="theme">navmap-mytheme.vm</ac:parameter>
  <ac:parameter ac:name="cellWidth">80px</ac:parameter>
  <ac:parameter ac:name="">mylabel</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{navmap:mylabel|wrapAfter=4|title=My map
name|cellHeight=50px|theme=navmap-mytheme.vm|cellWidth=80px}
```

Network Macro

The Network macro displays a list of Network activity on a Confluence page or blog post. You can specify the user whose network activity you wish to show. These interactions include the users that the specified user is following or users who are following the specified user. The Network macro shows each listed user by their profile picture. It also provides a choice of two themes and the ability to limit the number of users in the list.

Screenshot: Network macro

Using the Network macro

To add the Network macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Username	Current user's username	The username of the Confluence user whose network interactions you wish to show. If no username is specified, then current user's (that is, your) network interactions are shown.
Mode	**following**	Determines which users are listed, with respect to the specified user: • **following** – those who the user is following. • **followers** – those who are following the user. This parameter is unnamed in storage format and wikimarkup.
Theme	**full**	Determines how the user's network is displayed: • **full** – shows a large version of user's profile pictures and, if the **following** mode is set, provides an entry field function to follow more users. • **tiny** – shows only the small version of user's profile pictures.
Maximum Results **(max)**	No limit imposed up to a maximum of 30	Restricts the number of users displayed. If the number of users exceeds the specified maximum, then a **Show All** link is provided. This link leads to the specified user's Network view, showing the complete list of network interactions.

Disabling the Network macro

The Network macro is provided by the 'network' module in the 'Profile Macros' plugin. To remove the macro from your site, you can disable the module in the plugin. See Disabling and enabling add-ons.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `network`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="network">
  <ac:parameter ac:name="username">
    <ri:user ri:userkey="12345678912345678912345"/>
  </ac:parameter>
  <ac:parameter ac:name="max">10</ac:parameter>
  <ac:parameter ac:name="theme">full</ac:parameter>
  <ac:parameter ac:name="">followers</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{network:followers|username=admin|max=10|theme=full}
```

Noformat Macro

The Noformat macro displays a block of text in monospace font with no other formatting.

Using the Noformat Macro

To add the Noformat macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Enter the content that is to be unformatted into the body of the macro placeholder.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
No Panel (`nopanel`)	False	Removes the panel around the content.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `noformat`

Macro body: Accepts plain text.

Storage format example

```
<ac:structured-macro ac:name="noformat">
  <ac:parameter ac:name="nopanel">true</ac:parameter>
  <ac:plain-text-body><![CDATA[http://www.example.com]]></ac:plain-text-body>
</ac:structured-macro>
```

Wiki markup example

```
{noformat:nopanel=true}http://www.example.com{noformat}
```

Office Excel Macro

The Office Excel macro is going away in a future release

Read more about this change or find out about a better way to insert and preview Excel files.

The Office Excel macro displays the content of an Excel spreadsheet on a wiki page. First attach the spreadsheet to a Confluence page, then use the macro to display the content. When viewing the page, people will see the content of the spreadsheet, without needing to have Office installed.

For details, see the View File Macro.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `viewxls`

Macro body: None.

Parameter name	Required	Default	Parameter description and accepted values
`name`	Yes	(None.)	The file name of the Excel spreadsheet to be displayed. The document must be attached to a page on your Confluence site.
`page`	No	The page containing the macro	The name of a Confluence page to which the Excel spreadsheet is attached.
`col`	No	Last column with content	The number of the last column you want displayed, starting from '0' as the first column. **Hint for reducing the size of the spreadsheet:** Use the `col` and `row` parameters to reduce the size of the spreadsheet displayed on the wiki page. This is especially useful to prevent the display from showing empty cells. This will also help to prevent 'out of memory' errors.

row	No	Last row with content	The number of the last row you want displayed, starting from '0' as the first row.
grid	No	true	Determines whether the macro will show borders around the cells in the spreadsheet. Accepted values: `true` – Shows borders. `false` – Does not show borders.
sheet	No	The most-recently viewed worksheet	The name of the worksheet that you want displayed.

Storage format example

```
<ac:structured-macro ac:name="viewxls">
  <ac:parameter ac:name="col">E</ac:parameter>
  <ac:parameter ac:name="name">
    <ri:attachment ri:filename="Spreadsheet.xls"/>
  </ac:parameter>
  <ac:parameter ac:name="sheet">Sheet1</ac:parameter>
  <ac:parameter ac:name="row">5</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{viewxls:col=5|page=Docs|name=My document.xls|grid=false|sheet=mysheet|row=5}
```

Office PowerPoint Macro

> **The Office PowerPoint macro is going away in a future release**
>
> Read more about this change or find out about a better way to insert and preview PowerPoint files.

The Office PowerPoint macro displays the content of a PowerPoint presentation on a wiki page. First attach the presentation to a Confluence page, then use the macro to display the content. When viewing the page, people will see the content of the presentation, without needing to have Office installed.

For details, see the View File Macro.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: viewppt

Macro body: None.

Parameter name	Required	Default	Parameter description and accepted values
name	Yes	(None.)	The file name of the PowerPoint presentation to be displayed. The document must be attached to a page on your Confluence site.
page	No	The page containing the macro	The name of a Confluence page to which the PowerPoint presentation is attached.
height	No		The height of the macro display, specified in pixels (for example: 10 px, or just 10) or as a percentage (for example: 20%) of the window's height.
width	No		The width of the macro display, specified in pixels (for example: 10 px, or just 10) or as a percentage (for example: 20%) of the window's height.
slide	No	All slides, starting with the first, as a slide show	The number of the slide that you want displayed on the Confluence page, where the first slide is numbered zero. Instead of a slide show, the page will display just the single slide, represented as a JPEG image.

Storage format example

```
<ac:structured-macro ac:name="viewppt">
  <ac:parameter ac:name="height">250</ac:parameter>
  <ac:parameter ac:name="width">250</ac:parameter>
  <ac:parameter ac:name="name">
    <ri:attachment ri:filename="101_samplepresentation.pptx"/>
  </ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{viewppt:height=20%|page=Docs|width=20%|name=My document.ppt|slide=4}
```

Office Word Macro

The Office Word macro is going away in a future release

Read more about this change or find out about a better way to insert and preview Word files.

The Office Word macro displays the content of a Word document on a wiki page. First attach the document to a Confluence page, then use the macro to display the content. When viewing the page, people will see the content of the document, without needing to have Office installed.

For details, see the View File Macro.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `viewdoc`

Macro body: None.

Parameter name	Required	Default	Parameter description and accepted values
`name`	Yes	(None.)	The file name of the Word document to be displayed. The document must be attached to a page on your Confluence site.
`page`	No	The page containing the macro	The name of a Confluence page to which the Word document is attached.

Storage format example

```
<ac:structured-macro ac:name="viewdoc">
  <ac:parameter ac:name="name">
    <ri:attachment ri:filename="Chocolate.doc"/>
  </ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{viewdoc:page=Docs|name=My document.doc}
```

Page Index Macro

The Page Index macro creates a hyperlinked alphabetical index of all page titles within the current space.

The top section of the index contains a cell for each letter of the alphabet, with cells for numbers and symbols. Each cell indicates how many pages are in the corresponding list.

The lower section contains lists of page titles followed by the first few sentences of content on that page.

Each letter, number or symbol in the top section is hyperlinked to the corresponding cell in the lower section. Each page title in the lower section is hyperlinked to the page in the space.

Screenshot: Page Index macro (partial view)

Space Index

0-9 ... 1	A ... 1	B ... 1	C ... 1	D ... 0	E ... 0
F ... 0	G ... 1	H ... 0	I ... 1	J ... 0	K ... 0
L ... 0	M ... 0	N ... 0	O ... 1	P ... 0	Q ... 0
R ... 0	S ... 8	T ... 1	U ... 0	V ... 0	W ... 0
X ... 0	Y ... 0	Z ... 0	!@#$ ... 0		

0-9

📄 25 reasons people like Confluence
   Integration with JIRA Great product Its a wiki User experience design

A

🏠 Angry Nerds Home
   Welcome to Angry Nerds, a sample Confluence space. We think you're gonna like it here!

B

📄 Blog ideas
   Do users need to be told left side / right side of screen Macros and what you can do with them

C

📄 Confluence hints and tips
   What hints and tips do you share with other Confluence users?

D

E

F

G

📄 Games as a communication medium
   Can we use games as a communication medium?

Using the Page Index macro

To add the Page Index macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

This macro accepts no parameters.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `index`

Macro body: None.

Parameters: None.

Storage format example

```
<ac:structured-macro ac:name="index"/>
```

Wiki markup example

```
{index}
```

Page Properties Macro

The Page Properties and Page Properties Report macro work together to enable you to show summary information from one page on a another page. You can see examples of these two macros in action in the Decisions Blueprint and Product Requirements blueprints.

This macro was previously known as the Metadata Details macro.

Adding the Page Properties macro to a page

The Page Properties macro uses key value pairs. The body of the macro must be set up correctly for your information to appear in the Page Properties Report.

To add the Page Properties macro to a page:

1. In the editor, choose **Insert** > **Other Macros** > **Page Properties**.
2. In the macro body create a two column table and remove the header row.
3. In the left column list your 'keys' - these will be the column headings in your report table.
4. In the right column list the value for each key - these will populate the rows in your report table.
5. Add a **label** to your page - you will need to specify this label in the page properties report macro.
6. Save your page.

On this page:
- Adding the Page Properties macro to a page
- Using multiple Page Properties macros on one page
- Parameters

Related pages:
- Page Properties Report Macro
- Decisions Blueprint
- Product Requirements Blueprint

⚠ Don't forget to add a label to your page, or your page will not appear in the Page Properties Report macro.

Next you need to add the Page Properties Report macro to another page.

Screenshot: Page Properties macro in the editor. Deadline, Current Status and Team will be column headings in the report.

Screenshot: The example above as it appears in the Page Properties Report

Title	Current status	Deadline	Team
Project DEF	Not started	1 June	Purple Monkeys
Project ABC	In Progress	15 November	Green Parrots

Using multiple Page Properties macros on one page

You can add multiple Page Properties macros on a single page, and choose whether to include all or only specific macros in the report. You might use multiple macros because you want the information in the macro to display in context with the rest of the page, or because you want to be able to report on individual Page Properties macros separately.

The Page Properties macro includes an optional ID parameter that can be used to identify specific Page Properties macros.

To show the contents of **all** Page Properties macros in the report:

- Add a label to the page containing the Page Properties macros.
- Specify this label in the Page Properties Report macro.

To show the contents of **selected** Page Properties macros in the report:

- Add a label to the page containing the Page Properties macros.
- Specify an ID in the Page Properties macro that you want to report on.
- Specify both the label and ID in the Page Properties Report macro.

Note: The Page Properties Report macro can only accept one page label, and one ID.

> See the Page Properties macro in action in How to document product requirements in Confluence. This powerful macro lets you create a summary page that pulls in information from multiple pages.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Page Properties ID	(None)	Optional ID used to identify a particular Page Properties macro on a page. Specify this ID in the Page Properties Report to include summary information from macros with this ID only.
Hidden	False	Determines whether the data in the Page Properties macro will be displayed on the current page. This setting does not affect the display of the detail in the Page Properties Report macro.

Notes

- You can't use text formatting or macros in the left column as the data in this column is used to populate the column headings in your Page Properties Report macro.
- If your table has a header row, this row will be ignored by the Page Properties Report macro. You should remove the header row.
- It is not possible to reference the metadata using the metadata key from within the page, or anywhere else on a Confluence page.

Page Properties Report Macro

The Page Properties and Page Properties Report macro work together to enable you to show summary information from one page on a another page. You can see examples of these two macros in action on the Decision and Product Requirements blueprints.

This macro was previously known as the Details Summary macro.

Adding the Page Properties Report macro to a page

To add the Page Properties Report macro to a page:

1. In the editor, choose **Insert** > **Other Macros** > **Page Properties Report**.
2. Enter the **Labels** you want to report on - this is the label added to pages containing the Page Properties macro.
3. Further narrow down your search by adding more fields, or specifying a Page Properties ID (more info on this below)
4. Choose **Insert**.

Here's how the macro looks on your page:

Title	Current status	Deadline	Team
Project DEF	Not started	1 June	Purple Monkeys
Project ABC	In Progress	15 November	Green Parrots

And here's how you would set it up in the macro browser:

Reporting on specific Page Properties macros

It's possible to add multiple Page Properties macros on a page, and choose whether to include all or only specific macros in the report. The Page Properties macro includes an optional ID parameter that can be used to identify specific Page Properties macros.

To show the contents of:

- **Selected Page Properties macros** in the report - specify the label for the page and the ID of the particular Page Properties macro (under Options)
- **All Page Properties macros** in the report - specify just the label for the page - leave the Page Properties ID field blank.

Note: The Page Properties Report macro can only accept **one** ID.

CQL fields

CQL (Confluence Query Language) is a query language developed for Confluence, which you can use in some macros and the Confluence search. Confluence search and CQL-powered macros allow you to add filters to build up a search query, adding as many filters as you need to narrow down the search results.

Use the **Add a filter** link to add more filters to your query.

⌄ Use AND, OR, and NOT operators
- For an OR search, specify multiple values in the same field.
    So to show pages with 'label-a', 'label-b' or both you'd put 'label-a' and 'label-b' in the same Label field, like this:

- For an AND search, add more than one filter and specify a single value in each.
    To show only pages with label-a and label-b you'd put 'label-a' in one label field, then add a second Label field to the macro, and put 'label-b' in the second one, like this:

Put simply, OR values are entered in the same filter, AND values are entered in different filter. Only some filters support AND. If the filter doesn't support the AND operator, you won't be able to add that filter more than once.

- For a NOT search, enter a minus sign (-) before the label. This'll exclude everything with that label.

You can use the following CQL filters to build your query:

Filter	Description	Operators
Label*	Include pages, blog posts or attachments with these labels.	OR (multiple values in the same filter) AND (multiple Label filters)
With ancestor	Include pages that are children of this page. This allows you to restrict the macro to a single page tree.	OR (multiple values in the same filter)
Contributor**	Include pages or blog posts that were created or edited by these people.	OR (multiple values in the same filter)
Creator	Include items created by these people.	OR (multiple values in the same filter)
Mentioning user	Include pages and blog posts that @mention these people.	OR (multiple values in the same filter)
With parent	Include only direct children of this page (further sub-pages won't be included)	EQUALS (one page only)
In space**	Include items from these spaces.	OR (multiple values in the same filter)
Including text**	Include items that contain this text.	CONTAINS (single word or phrase)
With title	Include items that contain this text in the title.	CONTAINS (single word or phrase)
Of type**	Include only pages, blogs or attachments.	OR (multiple values in the same filter)

* This field is required in CQL-powered macros.

** You can add these filters in CQL-powered macros but in search they're part of the standard search filters, so they don't appear in the **Add a filter** menu.

Macro display options

These options control how the macro appears on your page.

Parameter	Default	Description

Page Properties ID	Blank	If not specified, the report will show data from all Page Properties macros on a page, where there are multiple macros. Specify an ID to include only data from Page Properties macros with the same ID.
Title column heading	Title	The heading to display on the first column in the report table. This column contains links to pages displayed by the report. The default column heading is 'Title'.
Columns to show		If not specified, the report will show all columns. You can specify a comma separated list of columns to include.
		If your column heading includes commas, use double quotes around the column name. If your column heading includes quotes, use double quotes. For example, A column, "My ""new"" column, yes", Third column
Number of items to display	30	Number of items to display in the table before displaying pagination options for additional items.
Sort by	Modified	Sort the table by a specific column heading. Enter the column name, exactly as it appears in the corresponding Page Properties macro.
		Select the **Reverse Sort** check box to sort the table in reverse order.
Show Comments Count	No	Displays the number of comments for each page in the table.
Show Likes Count	No	Displays the number of likes for each page in the table.

Troubleshooting

If your report is empty, check:

- You have entered the label correctly and that the label does appear on pages containing a Page Properties macro.
- The Page Properties macros on each page are configured correctly.
- Any other fields you have specified have not narrowed your search too far (for example there are no pages with that label under the Parent page you've specified).

Page Tree Macro

The Page Tree macro displays a dynamic, hierarchical list of pages starting from a specified parent (root) page. It can act as a table of contents or a list of related topics. When viewing the page tree, your reader can click a link to view the relevant page. The page's current position is highlighted in the page tree.

Note: The Page Tree macro is used in the Confluence Documentation theme, to create a navigation panel showing a table of contents for your space. Read the instructions on Configure the Documentation Theme.

Using the Page Tree Macro

To add the Page Tree macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Click **Refresh** in the Macro Browser to see the effect of changes to the macro parameters.

Macro Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Root Page **(root)**	The home page of the space	Specify the parent page for the page tree. The tree will include all children of the given page, plus their children and grand-children etc. The tree will *not* include the root page itself. Specify the page title or a special value as follows: • Your page title — to specify a page name for the parent or root of the tree. The tree will include all children and grand-children of the specified root. The tree will *not* include the specified root page itself. • '@home' — will include all pages under the home page of the space (default). • '@self' — will include all pages under the current page. • '@parent' — will include all pages under the parent of the current page, including the current page. • '@none' — will include all pages in the space, including orphaned pages and the home page.
Sort Pages By **(sort)**	**position**	Specify the order to display the pages in the tree. This sort order is for display purposes only. It does not permanently re-arrange the page order. The value may be one of the following: • **bitwise** — sort alphabetically, for example: title1, title10, title2. • **creation** — sort by date of creation. • **modified** — sort by order of date last modified. • **natural** — sort in 'natural' alphabetical order, for example: title1, title2, title10. • **position** — sort by the default Confluence sorting rules. If your pages have been ordered manually, this sort will respect the defined order. Otherwise the pages will be displayed in the 'natural' alphabetical order, such as: title1, title2, title10.
Include Excerpts in Page Tree **(excerpt)**	false	Select if you want the page tree to show excerpts from each page. The excerpts must be defined on each page by the Excerpt macro.
Reverse Order **(reverse)**	false	Select to show the pages in reverse (descending) natural order. Must be used in combination with the **Sort Pages By** parameter.
Include Search Box above Page Tree **(searchBox)**	false	Select if you want to include a search box above the page tree. The search box allows your readers to search within the page tree for the specified value.
Show Expand/Collapse Links **(expandCollapseAll)**	false	Select if you want to display the 'expand all' and 'collapse all' links at the top of your page tree. Your readers can click these links to open or close all branches of the tree at once. Available values in wikimarkup and storage format: • `true` – Show the 'expand all' and 'collapse all' options. • `false` – Do not show the options.

Start Depth **(startDepth)**	1	Enter any number greater than 0 to set how many levels of children the tree should show when it opens for the first time.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: pagetree

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="pagetree">
  <ac:parameter ac:name="reverse">false</ac:parameter>
  <ac:parameter ac:name="sort">natural</ac:parameter>
  <ac:parameter ac:name="root">
    <ac:link>
      <ri:page ri:content-title="Page Name"/>
    </ac:link>
  </ac:parameter>
  <ac:parameter ac:name="startDepth">3</ac:parameter>
  <ac:parameter ac:name="excerpt">true</ac:parameter>
  <ac:parameter ac:name="searchBox">true</ac:parameter>
  <ac:parameter ac:name="expandCollapseAll">true</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{pagetree:root=Page
Name|sort=natural|excerpt=true|reverse=false|startDepth=3|expandCollapseAll=true|se
archBox=true}
```

Page Tree Search Macro

Using the Page Tree Search macro, you can add a search box to your Confluence page. When a viewer enters a search term, Confluence will search a hierarchy of pages starting from a specified parent (root) page and return the search results.

You can also add a search box as part of a dynamic page tree, which looks like a table of contents. See the Page Tree macro .

Using the Page Tree Search macro

To add the Page Tree Search macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Name of Root Page (`root`)	none	The name of the root page whose hierarchy of pages will be searched by this macro. If this not specified, the root page is the current page. **Note:** Unlike the Page Tree macro, the Page Tree Search macro does not accept the special values that start with an @ sign, such as @home or @self.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `pagetreesearch`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="pagetreesearch">
  <ac:parameter ac:name="root">My page name</ac:parameter>
</ac:structured-macro>
```

Wikimarkup example

```
{pagetreesearch:root=My page name}
```

Panel Macro

The Panel macro displays a block of text within a customisable panel, like this:

> **My Title**
>
> A formatted panel.

Handy hint: You can use panels within columns, in table cells and in the sections defined by page layouts.

Using the Panel macro

To add the Panel macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Panel Title (`title`)	none	The title of the panel. If specified, this title will be displayed in its own title row.
Border Style (`borderStyle`)	solid	The style of the panel's border. Accepted values are `solid`, `dashed` and other valid CSS border styles.
Border Colour (`borderColor`)		The colour of the panel's border. Colours can be specified as HTML colour names or hexadecimal codes.
Border Pixel Width (Value Only) (`borderWidth`)		The width of the panel's border (in pixels).
Background Colour (`bgColor`)		The background colour of the panel. Colours can be specified as HTML colour names or hexadecimal codes.
Title Background Colour (`titleBGColor`)		The background colour of the title row of the panel. Colours can be specified as HTML colour names or hexadecimal codes.
Title Text Colour (`titleColor`)		The colour of the text in the title row of the panel. Colours can be specified as HTML colour names or hexadecimal codes.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `panel`

Macro body: Accepts rich text.

Storage format example

```
<ac:structured-macro ac:name="panel">
  <ac:parameter ac:name="bgColor">#72bc72</ac:parameter>
  <ac:parameter ac:name="titleBGColor">#00a400</ac:parameter>
  <ac:parameter ac:name="title">My title</ac:parameter>
  <ac:parameter ac:name="borderStyle">dashed</ac:parameter>
  <ac:parameter ac:name="borderColor">blue</ac:parameter>
  <ac:parameter ac:name="titleColor">white</ac:parameter>
  <ac:rich-text-body>

<p>A formatted panel</p>
  </ac:rich-text-body>
</ac:structured-macro>
```

Wiki markup example

```
{panel:title=My
title|borderStyle=dashed|borderColor=blue|titleBGColor=#00a400|titleColor=white|bgC
olor=#72bc72}
A formatted panel
{panel}
```

PDF Macro

The PDF macro is going away in a future release

Read more about this change or find out about a better way to insert and preview PDF files.

The PDF macro displays the content of a PDF document on a page. First attach the document to a Confluence page, then use the macro to display the document.

For details, see the View File Macro.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `viewpdf`

Macro body: None.

Parameter name	Required	Default	Parameter description and accepted values
`name`	Yes	(None.)	The file name of the PDF document to be displayed. The document must be attached to a page on your Confluence site.
`page`	No	The page containing the macro	The name of a Confluence page to which the PDF document is attached.

Storage format example

```
<ac:structured-macro ac:name="viewpdf">
    <ac:parameter ac:name="name">
      <ri:attachment ri:filename="My_document.pdf"/>
    </ac:parameter>
  </ac:structured-macro>
```

Wiki markup example

```
{viewpdf:page=Docs|name=My document.pdf}
```

Popular Labels Macro

The Popular Labels macro displays the most popular labels used throughout your Confluence site or within a

space. A popular label is a label that has been added to many pages.

Using the Popular Labels Macro

To add the Popular Labels macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Number of Labels to Display (`count`)	100	Specifies the total number of labels to display in the heatmap.
Restrict Labels to this Space Key (`spaceKey`)	none	Restricts the list of popular labels to the specified space.
Style of Labels (`style`)	list	• **list** – displays the popular labels as a bulleted list, ordered by popularity (highest first). • **heatmap** – displays the popular labels using different font sizes for each label depending on the label's popularity, ordered by label names.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `popular-labels`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="popular-labels">
    <ac:parameter ac:name="count">20</ac:parameter>
    <ac:parameter ac:name="spaceKey">
      <ri:space ri:space-key="ds"/>
    </ac:parameter>
  </ac:structured-macro>
```

Wikimarkup example

```
{popular-labels:style=heatmap|count=20|spaceKey=ds}
```

Profile Picture Macro

You can use the Profile Picture macro to display a user's profile picture on a page, and it's useful for things like creating team pages that show all members of a project team.

Hover your mouse-over the picture to see the Hover Profile for the user, and choose the user's picture or name to view their user profile. When editing the page, you can also select the macro and choose **Vie w User Profile** to see the profile for the user.

Here's an example of how to use the Profile Picture macro: How to build a release planning page in Confluence.

Add the Profile Picture Macro

1. In the Confluence editor, choose **Insert** > **Oth er Macros**
2. Find and select the 'Profile picture' macro
3. Search for and select the user and choose **Sa ve**

Speeding up macro entry with autocomplete: Typ e **{** and the beginning of the macro name, then select the macro from a list of suggested macros. For more information, see Using Autocomplete.

Code examples

The following example is provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `profile-picture`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="profile-picture">
 <ac:parameter ac:name="User">
   <ri:user ri:userkey="123456789123456789012345"/>
 </ac:parameter>
</ac:structured-macro>
```

Wiki markup

Wiki markup is not available for this macro. You cannot add this macro via wiki markup.

Recently Updated Dashboard Macro

The Recently Updated Dashboard macro displays a list of the most recently changed pages, blogs and files in your site. It is similar to the Recently Updated macro but with tabs to allow you to see all updates, updates from your favourite spaces or updates from your network (the people you follow).

Using the Recently Updated Dashboard macro

To add the Recently Updated Dashboard macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of

suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Space(s) `(spaces)`	The space which contains the page on which the macro is added	Filter content by space. The macro will display only the pages etc which belong to the space(s) you specify here. You can specify one or more space keys, separated by commas. Use '*' for all spaces.
Include these Content Types Only `(types)`	all types	Filter content by type. You can specify one or more types, separated by commas. Available types are: `page`, `blogpost` or `news`, `spacedesc`, `attachment`, `comment`, `mail`, `userinfo`.
Label(s) `(labels)`	none	Filter content by label. The macro will display only the pages etc which are tagged with the label(s) you specify here. You can specify one or more labels, separated by commas. *Note:* If there are **no** pages matching any of the specified labels, then Confluence will ignore the labels and will list all recently updated pages.
User(s) `(users)`	all users	Filter by username of the user who updated the content. The macro will only display content created and updated by the user(s) you specify here. You can specify one or more usernames separated by commas.
Width of Table `(width)`	100%	Specify the width of the macro display, as a percentage of the window width.
Show User Profile Pictures `(showProfilePic)`	false	Select whether profile pictures of the users who updated the content are displayed.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `recently-updated-dashboard`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="recently-updated-dashboard">
  <ac:parameter ac:name="spaces">
    <ri:space ri:space-key="ss"/>
  </ac:parameter>
  <ac:parameter ac:name="users">
    <ri:user ri:userkey="12345678912345678912345"/>
  </ac:parameter>
  <ac:parameter ac:name="width">50%</ac:parameter>
  <ac:parameter ac:name="labels">choc</ac:parameter>
  <ac:parameter ac:name="types">page</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{recently-updated-dashboard:spaces=ds|users=admin|width=50%|showProfilePic=true|labels=choc|types=page}
```

Recently Updated Macro

The Recently Updated macro displays a list of the most recently changed content within Confluence.

Using the Recently Updated Macro

To add the Recently Updated macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Author(s) by username (`author`)	None specified. That is, display all content	Filter the results by author. The macro will display only the pages etc which were last modified by the author(s) you specify here. You can specify multiple users.

Space(s) **(spaces)**	@self That is, the space which contains the page on which the macro is used	This parameter allows you to filter content by space. The macro will display only the pages etc which belong to the space(s) you specify here. You can specify one or more space keys, separated by a comma or a space. • To exclude content in a specific space, put a minus sign (-) immediately in front of that space key. For example: If you specify a space key of -BADSPACE you will get only content which is not in the BADSPACE. • To indicate that the results **must** come from a specific space, put a plus sign (+) immediately in front of that space key. For example: If you specify a space key of +GOODSPACE you will get only content in GOODSPACE. (Note that this is not particularly useful, because each content item belongs to one space only. If you put a plus sign next to one space key and list other space keys too, the other space keys will be ignored.) Special values: • @self — The current space. • @personal — All personal spaces. • @global — All site spaces. • @favorite — The spaces you have marked as favourite. • @favourite — The same as @favorite above. • @all — All spaces in your Confluence site. • * — The same as @all above. When specifying a personal space, remember to use the tilde (~) sign in front of the username, such as ~jbloggs or ~jbloggs@example.com.
Label(s) **(labels)**	None specified i.e. display all content	Filter the results by label. The macro will display only the pages etc which are tagged with the label(s) you specify here. You can specify one or more label values, separated by a comma or a space. • To exclude content which matches a given label, put a minus sign (-) immediately in front of that label value. For example: If you specify a label value of -badpage you will get only content which is not labelled with 'badpage'. • To indicate that the results **must** match a given label value, put a plus sign (+) immediately in front of that label value. For example: If you specify a label value of +superpage,+goodpage you will get only content which has at least two labels, being 'superpage' and 'goodpage'. The labels parameter only applies to the page and blog content types.
Width of Table **(width)**	100%	Specify the width of the macro display, as a percentage of the window width.

Include these Content Types Only (`types`)	All types	This parameter allows you to filter content by content type. The macro will display only the content of the type you specify here. You can specify one or more types, separated by a comma or a space. To exclude content of a given content type, put a minus sign (-) immediately in front of that content type. For example: If you specify a content type of `-blogpost` you will get pages and all other content except for blog posts. Available values: • `page` — Pages. • `blogpost` or `news` — Blog posts, also known as news items. • `comment` — Comments on pages and blog posts. • `attachment` — Attachments.
Maximum Number of Results (`max`)	15	Specify the maximum number of results to be displayed. If this parameter is omitted, then a maximum of 15 results are displayed. The theoretical maximum value that this parameter can accept is 2 to the power of 31, minus 1 (or 2147483647), though this has been limited to 200 in the code, for performance reasons. More details are here.
theme (`theme`)	`concise`	Choose the appearance of this macro: • `concise` — the default list, showing the names of pages which were updated or commented on, the users who made the page modifications and time when the modifications occurred. • `social` — lists recent modifications in reverse chronological order, but groups them by user into short time segments. A 'sub' list appears within each user's time segment, showing the names of pages which they updated or commented on and time when these modifications occurred. • `sidebar` — lists recent updates in reverse chronological order, showing the names of pages which were updated or commented on and time when the page modifications occurred. This theme does not show authorship.
Show User Profile Pictures (`showProfilePic`)	`false`	Specify `showProfilePic=true` to display the profile pictures of the users who updated the content.
Hide Title (`hideHeading`)	False	Determines whether the macro hides or displays the text 'Recently Updated' as a title above the list of content. Only available in wikimarkup and storage format. Accepted values: • `true` – Title is hidden. • `false` – Title is shown.

Notes

- The Recently Updated Dashboard macro is similar to this macro, but is intended for display on the Confluence dashboard.
- If you would like to change the wording displayed by the Recently Updated macro, please refer to the document on modifying the Confluence interface text.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a

Confluence page.

Macro name: `recently-updated`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="recently-updated">
  <ac:parameter ac:name="spaces">
    <ri:space ri:space-key="ss"/>
  </ac:parameter>
  <ac:parameter ac:name="author">
    <ri:user ri:userkey="12345678912345678912345"/>
  </ac:parameter>
  <ac:parameter ac:name="users"/>
  <ac:parameter ac:name="max">10</ac:parameter>
  <ac:parameter ac:name="width">50%</ac:parameter>
  <ac:parameter ac:name="theme">sidebar</ac:parameter>
  <ac:parameter ac:name="labels">choc</ac:parameter>
  <ac:parameter ac:name="types">page</ac:parameter>
</ac:structured-macro>
```

Wikimarkup example

```
{recently-updated:spaces=ds|author=admin|max=10|hideHeading=true|width=50%|theme=si
debar|showProfilePic=true|labels=choc|types=page}
```

Recently Used Labels Macro

The Recently Used Labels macro displays a list of the labels that have been most recently added to a page, blog post or attachment. You can define the number of labels to be displayed and the scope (the current space, your personal space or site spaces, also known as 'global' spaces).

Using the Recently Used Labels macro

To add the Recently Used Labels macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters of this macro

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description

Number of Labels to Display (`count`)	10	Specifies the total number of labels to display in the list.
Scope for Retrieving Labels (`scope`)	global	Specifies the scope of labels to be displayed in the list. Valid values include: • **global** — covers all site spaces (non-personal) in the Confluence installation. • **space** — the current space. • **personal** — your own personal space.
List Style (`style`)	list	• **list** – displays the list of labels horizontally. • **table** – includes additional information such as the page to which the label was added and the user who added it.
Table Title (`title`)	none	Adds a title to the top of the list in table style. Titles are only visible when the **List Style** parameter has been set to **table**.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `recently-used-labels`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="recently-used-labels">
    <ac:parameter ac:name="title">My title</ac:parameter>
    <ac:parameter ac:name="scope">space</ac:parameter>
    <ac:parameter ac:name="style">table</ac:parameter>
    <ac:parameter ac:name="count">20</ac:parameter>
  </ac:structured-macro>
```

Wiki markup example

```
{recently-used-labels:title=My title|scope=space|style=table|count=20}
```

Related Labels Macro

The Related Labels macro lists all labels from every page which has one or more labels in common with the current page.

Using the Related Labels Macro

To add the Related Labels macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Restrict to these Labels `(labels)`	none	Specify the labels for which you want to view related labels. For example, `doc umentation,my:stuff`.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: related-labels

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="related-labels">
    <ac:parameter ac:name="labels">choc,cake</ac:parameter>
  </ac:structured-macro>
```

Wiki markup example

```
{related-labels:labels=choc,cake}
```

Roadmap Planner Macro

Create simple, visual timelines that are useful for planning projects, software releases and much more with the Roadmap macro.

Roadmaps are made up of:

- **bars** to indicate phases of work
- **lanes** to differentiate between teams, products or streams
- **markers** to highlight important dates and milestones
- a **timeline** showing months or weeks.

You can provide more information about items on your roadmap by linking a bar to a page.

Screenshot: Roadmap macro as it appears on a page

Adding the Roadmap macro

To add the Roadmap macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Editing your Roadmap

To edit your roadmap:

- Select the roadmap and choose **Edit**.
- **Add** lanes, bars and markers.
- Drag lanes, bars and markers to the desired location on the roadmap.
- Select lanes, bars and markers to add text, change colours and remove from the roadmap.
- Select bars to add links to existing pages, create new pages or add a description.
- Set the start and end dates for the roadmap and choose to display it by weeks or months.

Screenshot: Roadmap macro in the editor

Parameters

This macro does not use the macro browser to set parameters. You also cannot add this macro via wiki markup or by editing the storage format directly.

Notes

The Roadmap macro was previously available as an add-on from The Marketplace. The macro has changed significantly. If you had an older version of the macro installed you will be able to view your existing roadmaps but not edit them.

RSS Feed Macro

The RSS Feed macro embeds an RSS feed on a page. It can display the contents of external feeds, or of internal feeds generated by Confluence. To display blog posts or to list recently updated pages in a space, use the Feed Builder to create an internal feed, then render it using this macro.

> **CAUTION: Including unknown HTML inside a webpage is dangerous.**
> HTML inside an RSS feed can contain active scripting components. This means that it would be possible for a malicious attacker to present a user of your site with script that their web browser would believe came from you. Such code could be used, for example, to steal a user's authentication cookie and give the attacker their Confluence login password.

The RSS Feed macro may be disabled by your Confluence administrator. Also, your Confluence administrator

can define a whitelist of trusted URLs. You will see an error message on the Confluence page, if the included URL is not in the whitelist.

Using the RSS Feed macro

To add the RSS Feed macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.
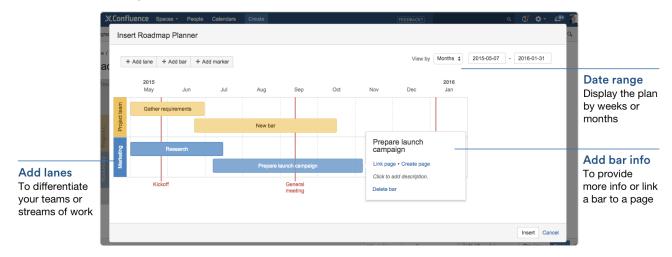
To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
RSS Feed URL (`url`)	none	The URL of the RSS feed link you want to show.
Maximum Number of Entries (`max`)	15	Limit the number of entries displayed.
Show Item Titles Only (`showTitlesOnly`)	false	Show only the titles of the news items, not the content.
Show Name/Title of RSS Feed (`titleBar`)	true	Hide the feeds title bar.

How up to date is the feed?

By default, the RSS Feed macro caches the feed results for 60 minutes before fetching the data again.

If you wish to change the default caching, use the Cache macro to define how often the RSS Feed macro fetches the feed updates. You will need to install the Cache plugin onto your Confluence site.

What happens to a page containing a disallowed URL?

Your Confluence Administrator can set up a whitelist of allowed URLs. If this is the case, you may see an error on the pages which contain the RSS Feed macro.

A user can add the RSS Feed macro or the HTML-include macro to a Confluence page. The macro code includes a URL from which the content is drawn. When the page is displayed, Confluence will check the URL against the whitelist. If the URL is not allowed, Confluence will display an error message on the page.

The error message says that Confluence "could not access the content at the URL because it is not from an allowed source" and displays the offending URL. If the person viewing the page is a Confluence Administrator, they will also see a link to the Administration page where they can configure the URL whitelist.

Here is an example of the error message, including the link shown only to Confluence Administrators:

Here is an example of the error message, but without the link.

Authentication

Private feeds from external sites

RSS feeds which require authentication cannot be accessed using the RSS Feed macro.

Accessing internal HTTPS feeds

This applies only to Confluence instances which have enabled HTTPS for all content. If your site is fully HTTPS, the RSS Feed macro cannot access internal feeds. To enable the RSS Feed macro to access internal feeds without affecting your HTTPS setup, enable local-only HTTP access:

1. Shut down Confluence.
2. Consult the SSL guide to enable HTTP access to Confluence. You'll want to ensure that you have an HTTP connector and an SSL connector, both commented in. This means that Confluence will be accessible via both HTTP and HTTPS. However, you should *not* have a redirect port, nor rules in web.xml to redirect all traffic.
3. Instead of using web.xml to redirect traffic, insert a firewall rule to redirect all HTTP requests not from the Confluence server to the equivalent HTTPS URL. This ensures that users will only be able to access Confluence via HTTPS, as intended. If you have still left HTTP access for attachments enabled (to avoid the IE download bug) you must selectively enable those URLS as well.
4. Modify your Confluence RSS Feed macro feed link to use the HTTP URL, and restart Confluence.

Enabling the RSS Feed macro

The RSS Feed macro is disabled by default.

To enable the RSS Feed macro:

1. Go to

   > **Add-ons**.
2. Select **System** from the drop down and search for the **Confluence HTML Macros** add-on.
3. Expand the add-on and enable the **rss (rss-xhtml)** module.

Code examples

Macro name: rss

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="rss">
    <ac:parameter ac:name="max">10</ac:parameter>
    <ac:parameter ac:name="showTitlesOnly">true</ac:parameter>
    <ac:parameter ac:name="url">
      <ri:url ri:value="http://myblog.com/feed"/>
    </ac:parameter>
    <ac:parameter ac:name="titleBar">false</ac:parameter>
  </ac:structured-macro>
```

Wikimarkup example

```
{rss:max=10|showTitlesOnly=true|url=http://myblog.com/feed|titleBar=false}
```

Search Results Macro

The Search Results macro searches your Confluence site based on search terms specified in the macro parameters, and displays the results on the wiki page.

Using the Search Results macro

To add the Search Results macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Search Terms (`query`)	none	The search terms which this macro will use to generate its results. You can refine your search query by using operators such as 'AND' and 'OR'. For example: `my_query1 AND my_query2` For more information, take a look at the documentation on the Confluence search syntax.
Maximum Number of Results (`maxLimit`)	10	Set a limit to the number of search results displayed.
Restrict to this Space Key	all	Start typing the space name to find the space, or specify the key of the space you want to search in. Note that the key is case sensitive.

Content Type **(type)**	all	Specify the content type. The content types are: **page**, **comment**, **blogpost**, **attachment**, **userinfo** (the content of user profiles only) and **spacedesc** (the content of space descriptions only).
Last Modified **(lastModified)**	all	Specify a period of time in weeks, days, hours and/or minutes, to see the content modified within that time frame. These are the values you can use: • `w` = weeks • `d` = days • `h` = hours • `m` = minutes For example: • `2h 35m` • `3d 30m` Notes: • If no time category is specified, Confluence assumes minutes. • If you specify more than one time period (for example, weeks and days), you must separate the periods with a space. You can put them in any order. • The time categories are not case sensitive. For example, '`4d`' is the same as '`4D`'.
Restrict to this Username **(contributor)**	all	Specify the username of a Confluence user, to show only content created or updated by that user.

Notes

Permissions: When a user views the page containing the Search Results macro, the search results will show only pages and other content types for which the user has 'View' permission.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `search`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="search">
  <ac:parameter ac:name="lastModified">3w</ac:parameter>
  <ac:parameter ac:name="query">choc</ac:parameter>
  <ac:parameter ac:name="contributor">
    <ri:user ri:userkey="12356789123456789123456789912345"/>
  </ac:parameter>
  <ac:parameter ac:name="maxLimit">10</ac:parameter>
  <ac:parameter ac:name="type">page</ac:parameter>
  <ac:parameter ac:name="spacekey">
    <ri:space ri:space-key="ss"/>
  </ac:parameter>
</ac:structured-macro>
```

```
{search:lastModified=3w|query=choc|contributor=admin|maxLimit=10|type=page|spacekey
=ds}
```

Section Macro

Used with the Column macro to define columns on a page. See Working with page layouts and columns and sections.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `section`

Macro body: Rich text, consisting of one or more Column macros.

Parameter name	Required	Default	Parameter description and accepted values
`border`	No	False	If the value is `true`, the columns in this section will have a dashed border.

Storage format example

The following example shows a section and column macro together.

```
<ac:structured-macro ac:name="section">
  <ac:parameter ac:name="border">true</ac:parameter>
  <ac:rich-text-body>
    <ac:structured-macro ac:name="column">
      <ac:parameter ac:name="width">100px</ac:parameter>
      <ac:rich-text-body>
        <p>This is the content of <strong>column 1</strong>.</p>
      </ac:rich-text-body>
    </ac:structured-macro>
    <ac:structured-macro ac:name="column">
      <ac:rich-text-body>
        <p>This is the content of <strong>column 2</strong>.</p>
      </ac:rich-text-body>
    </ac:structured-macro>
  </ac:rich-text-body>
</ac:structured-macro>
```

Wiki markup example

```
{section:border=true}
{column:width=100px}
This is the content of *column 1*.
{column}
{column}
This is the content of *column 2*.
{column}
{section}
```

Space Attachments Macro

The Space Attachments macro displays a list of all files attached to pages in a space. It shows details of the file and the includes a link to the page a file is attached to.

Filters allow you to show only files with a particular label or file extension.

Screenshot: The Space Attachments macro

Using the Space Attachments macro

To add the Space Attachments macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

The Space Attachments macro displays a list of all the attachments in a space. You can choose to show attachments from the current space, or another space.

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Space	(none)	Selects the Confluence space to display attachments for. If you do not specify a space, the current space will be used.
Show Filter Controls (`showFilter`)	true	Determines whether or not the filter panel is shown. If you select this option, people viewing the page will be able to filter the list of attachments by file type (extension) and by label.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `space-attachments`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="space-attachments">
    <ac:parameter ac:name="showFilter">false</ac:parameter>
    <ac:parameter ac:name="space">
      <ri:space ri:space-key="ds"/>
    </ac:parameter>
  </ac:structured-macro>
```

Wiki markup example

```
{space-attachments:showFilter=false|space=ds}
```

Space Details Macro

The Space Details macro displays the details of a Confluence space, including the space name, description, and more.

Using the Space Details macro

To add the Space Details macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description

Width of Table **(width)**	100%	The width of the space details table, specified as a percentage (%) of the page width.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `space-details`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="space-details">
    <ac:parameter ac:name="width">50%</ac:parameter>
  </ac:structured-macro>
```

Wiki markup example

```
{space-details:width=50%}
```

Space Jump Macro

This page describes the Space Jump macro, which is provided with the Documentation theme. You can use space jumping to link from a page in one space to a page with the same name in another space, without knowing the name of the page when you create the link.

> The Documentation theme is going away in Confluence 6.0.
> Read more about this change.

The Space Jump macro inserts a link onto the page. When the reader clicks the link, they will jump to a page that has the same name in another space. You specify the space when you insert the Space Jump macro.

Example: We use the Space Jump macro to put a standard message at the top of our archive spaces, telling people that they're reading an old version of the documentation and letting them jump quickly to the same page in the latest documentation. See the words '**this page in the current documentation**' in the screenshot below.

Screenshot: One way of using the Space Jump macro

Using the Space Jump macro

To add the Space Jump macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**

2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Space	None	The space key that contains a page with the same page title that you want the link to jump to. Space keys are case-sensitive.
Alias	None	The text to display as the hyperlinked text. If you do not supply an alias, Confluence will display the page name.

Notes

What happens if there is no page with the same name in the target space?

For the space jump to work, the target space must contain a page with the same name as the page that renders the Space Jump macro. If the target space does not contain such a page, you will see a broken link. Confluence handles this in its usual manner: the link is coloured red. If you click the link, Confluence offers to create the page for you.

Can I use the Space Jump macro in any space?

Yes. You can use the Space Jump macro in any space, even if that space is not currently using the Documentation theme. Provided that the Documentation theme plugin and its components are installed and enabled on your Confluence site, the Space Jump macro is available in any space.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `spacejump`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="spacejump">
    <ac:parameter ac:name="alias">My Space</ac:parameter>
    <ac:parameter ac:name="space">
      <ri:space ri:space-key="SS"/>
    </ac:parameter>
  </ac:structured-macro>
```

Wiki markup example

```
{spacejump:alias=my link|space=ds}
```

Spaces List Macro

Use the Spaces List macro to display a list of spaces on a page.

Using the Spaces List macro

To add the Spaces List macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Required?	Default	Description
Scope of spaces	no	all	The view from which spaces are listed. Available options are: • **(blank) -** All spaces in your site, with tabs. • **all** – All spaces in your Confluence site. • **category** – Spaces grouped according to space categories. • **favourite** – Spaces which you have added to My Spaces. • **new** – spaces created within the last 7 days. This parameter is unnamed in wiki markup and storage format.
Width of List (`width`)	no	100%	The width of the spaces list, specified as a percentage (%) of the window width.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `spaces`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="spaces">
    <ac:parameter ac:name="width">80%</ac:parameter>
    <ac:parameter ac:name="">favourite</ac:parameter>
  </ac:structured-macro>
```

Wiki markup example

```
{spaces:favourite|width=80%}
```

Status Macro

The Status macro displays a coloured lozenge (a rounded box) that is useful for reporting project status. You can choose the colour of the lozenge and the text that appears inside the lozenge. The macro also displays its current status in the editor, and you can change the status directly in the editor.

See the Status macro put to excellent use in How to build a release planning page in Confluence.

ALL GOOD!

Screenshot: Click on the Status macro lozenge to change the status.

Using the Status macro

To add the Status macro to a page using the Macro Browser:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Colour `(colour)`	Grey	The colour of the lozenge. The following colours are available: **Grey** , **Red**, **Yellow**, **Green** and **Blue**.
Title `(title)`	The colour that you select.	The text that will appear inside the lozenge. If you do not specify any text, the title will be the colour of the lozenge, that is 'Grey', 'Red', 'Yellow', 'Green' or 'Blue'.
Use outline style `(subtle)`	False	The style of the lozenge and its border. The default style lozenge is a solid background colour with white text. The outline style lozenge is white with a coloured border and coloured text as shown here **ALL GOOD!**

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `status`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="status">
    <ac:parameter ac:name="colour">Green</ac:parameter>
    <ac:parameter ac:name="title">On track</ac:parameter>
    <ac:parameter ac:name="subtle">true</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{status:colour=Green|title=On track|subtle=true}
```

Table of Contents Macro

The Table of Contents macro scans the headings on the current Confluence page to create a table of contents based on those headings. This helps readers find their way around lengthy pages, by summarising the content structure and providing links to headings.

Want to see the Table of Contents macro in action? Check out how it's used in Creating insightful customer interview pages.

Using the Table of Contents macro

Hint: For quick access from the editor toolbar, choose **Insert** > **Table of Contents**.

To add the Table of Contents macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Macro parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Output Type (`type`)	list	• **list** — produces a typical list-type table of contents. • **flat** — produces a horizontal menu-type series of links.
Display Section Numbering (`outline`)	clear	Select the check box to apply outline numbering to your headings, for example: 1.1, 1.2, 1.3.

List Style **(style)**	disc	Select the style of bullet point for each list item. You can use any valid CSS style. For example: `none` — no list style is displayed `circle` — the list style is a circle `disc` — the list style is a filled circle. This is the typical bullet list, and is used for this example list. `square` — the list style is a square `decimal` — the list is numbered (1, 2, 3, 4, 5) `lower-alpha` — the list is lower-case, alphabetised (a, b, c, d, e) `lower-roman` — the list style is lower roman numerals (i, ii, iii, iv, v, vi) `upper-roman` — the list style is upper roman numerals (I, II, III, IV, V, VI)
Heading Indent **(indent)**		Sets the indent for a **list** according to CSS quantities. Entering `10px` will successively indent heading groups by 10px. For example, level 1 headings will be indented 10px and level 2 headings will be indented an additional 10px.
Separator **(separator)**	brackets	This parameter applies to **flat** lists only. You can enter any of the following values: `brackets` — Each item is enclosed by square brackets: [ ]. `braces` — Each item is enclosed by braces: { }. `parens` — Each item is enclosed by parentheses: ( ). `pipe` — Each item is separated by a pipe: *anything* — Each item is separated by the value you enter. You can enter any text as a separator, for example "***". If using a custom separator, be aware that text displays exactly as entered, with no additional white space to further separate the characters.
Minimum Heading Level **(minLevel)**	1	Select the highest heading level to start your TOC list. For example, entering 2 will include levels 2, and lower, headings, but will not include level 1 headings.
Maximum Heading Level **(maxLevel)**	7	Select the lowest heading level to include. For example, entering 2 will include levels 1 and 2, but will not include level 3 headings and below.
Include Headings **(include)**		Filter headings to include according to specific criteria. You can use wildcard characters. See Sun's Regex documentation for examples of constructing regular expression strings.
Exclude Headings **(exclude)**		Filter headings to enclude according to specific criteria. You can use wildcard characters. See Sun's Regex documentation for examples of constructing regular expression strings.
Printable **(printable)**	checked	By default, the TOC is set to print. If you clear the check box, the TOC will not be visible when you print the page.
CSS Class Name **(class)**		If you have custom TOC styles in your CSS style sheet, use this parameter to output the TOC inside `<div>` tags with the specified class attribute.
Absolute URL **(absoluteURL)**		By default, the links in the TOC are relative URLs pointing to the current page. If checked, the links in the TOC will be full URLs. This setting is useful when you are including a page with a Table of Contents in another page, and want to control where the links should take the user.

Examples

The examples below are based on this table of contents:

- Places
    - Favourite Places
    - Unknown Places
- Foods
    - Favourite Foods
    - Unknown Foods
- Things
    - Favourite Foods
    - Unknown Foods

Filtered Table of Contents

This example filters the headings to include those that contain 'Favourite', but excludes headings which end with 'Things'. The list is styled with Roman numerals.

Parameter	Value
List Style	`upper-roman`
Include Headings	`Favourite.*`
Exclude Headings	`.*Things`

The resulting table of contents is:

   I. Favourite Places
  II. Favourite Foods

Flat List

This example filters all headings to render a flat list of 'Unknowns' enclosed in square brackets (the default list style).

Parameter	Value
Output Type	**flat**
Maximum Heading Level	2
Include Headings	`Unknown.*`

The resulting table of contents is:

[Unknown Places] [Unknown Foods] [Unknown Foods]

Notes

- When you use a Table of Contents macro in a **template**, you will see an error when you preview the template itself. But the Table of Contents macro works on the pages that people create from the template – the table of contents shows up after they have saved the page. (This is probably because the template is not defined as a page, and the Table of Contents macro works for pages only.)
- Due to an outstanding issue in the Table of Contents macro (CONF-10619), the macro browser's **Refresh** function does not render any parameter modifications. Currently, the rendering of parameter value modifications to the Table of Contents macro occurs only after the page is saved.
- **Using HTML heading markup with the Table of Contents macro**
  The Table of Contents macro cannot handle HTML heading markup on its own. Hence, if you use the HTML and HTML Include macros to render HTML heading markup in a Confluence page, the Table of Contents macro will not create a contents list out of these headings. (For more information about this issue, please refer to TOC-93.)
  However, if you insert an HTML anchor into each HTML heading on your page (based on the following syntax), the Table of Contents macro will incorporate these headings into your contents list.

```
<h2><a name="pagename-headingname"></a>Heading Name</h2>
```

  The syntax for the anchor name is the page name and heading name separated by a hyphen. Remove all spaces and convert all text to lower case. Convert all punctuation marks to their URL-encoded equivalent.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `toc`

Macro body: None.

Storage format example

This example shows a list-type table of contents.

```
<ac:structured-macro ac:name="toc">
    <ac:parameter ac:name="printable">true</ac:parameter>
    <ac:parameter ac:name="style">square</ac:parameter>
    <ac:parameter ac:name="maxLevel">2</ac:parameter>
    <ac:parameter ac:name="indent">5px</ac:parameter>
    <ac:parameter ac:name="minLevel">2</ac:parameter>
    <ac:parameter ac:name="class">bigpink</ac:parameter>
    <ac:parameter ac:name="exclude">[1//2]</ac:parameter>
    <ac:parameter ac:name="type">list</ac:parameter>
    <ac:parameter ac:name="outline">true</ac:parameter>
    <ac:parameter ac:name="include">.*</ac:parameter>
  </ac:structured-macro>
```

This example shows a flat table of contents.

```
<ac:structured-macro ac:name="toc">
    <ac:parameter ac:name="maxLevel">2</ac:parameter>
    <ac:parameter ac:name="minLevel">2</ac:parameter>
    <ac:parameter ac:name="class">bigpink</ac:parameter>
    <ac:parameter ac:name="exclude">[1//2]</ac:parameter>
    <ac:parameter ac:name="type">flat</ac:parameter>
    <ac:parameter ac:name="outline">true</ac:parameter>
    <ac:parameter ac:name="separator">pipe</ac:parameter>
    <ac:parameter ac:name="include">.*</ac:parameter>
</ac:structured-macro>
```

Wiki markup example

This example shows a list-type table of contents.

```
{toc:printable=true|style=square|maxLevel=2|indent=5px|minLevel=2|class=bigpink|e
xclude=[1//2]|type=list|outline=true|include=.*}
```

This example shows a flat table of contents.

```
{toc:printable=true|maxLevel=2|minLevel=2|class=bigpink|exclude=[1//2]|type=flat|
outline=true|separator=pipe|include=.*}
```

Table of Content Zone Macro

You can use the Table of Content Zone macro to mark out a section (zone) within the page from which to create a table of contents. You can style the headings as a flat list, and place the list of links at the top and bottom of the section or page, to provide navigation bars similar to web navigation.

Using the Table of Content Zone macro

To add the Table of Content Zone macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
location (`location`)	both	Specifies where in the zone the output list is displayed: **top**, **bottom**, or **both**, which encloses the page zone content.

Output Type **(type)**	list	Specifies the layout for the table of contents: • **list** – produces a vertical list, typical of a TOC. • **flat** – produces a horizontal menu-type series of links, for example: [Heading 1] [Heading 2] [Heading 3].
Display Section Numbering **(outline)**	false	Select to apply outline numbering to your headings, for example: 1.1, 1.2, 1.3.
List Style **(style)**	none	Specifies the style of bullet point for each list item. You can use any valid CSS style. For example: • `none` — no list style is displayed • `circle` --- the list style is a circle • `disc` — the list style is a filled circle. This is the typical bullet list, and is the one we're using for this example list • `square` — the list style is a square • `decimal` — the list is numbered (1, 2, 3, 4, 5) • `lower-alpha` — the list is lower-case, alphabetised (a, b, c, d, e) • `lower-roman` — the list style is lower roman numerals (i, ii, iii, iv, v, vi) • `upper-roman` — the list style is upper roman numerals (I, II, III, IV, V, VI)
Heading Indent **(indent)**		Sets the indent for a **list** output type, according to CSS quantities. Entering "10px" will successively indent list heading levels by 10px. For example, h1 headings will be indented 10px and h2 headings will be indented an additional 10px.
Separator **(separator)**	brackets	Only applies to the **flat** output type**.** Specifies the display style of the links. You can enter any of the following values: • `brackets` — Each item is enclosed by square brackets: [ ]. • `braces` — Each item is enclosed by braces: { }. • `parens` — Each item is enclosed by parentheses: ( ). • `pipe` — Each item is separated by a pipe: • *anything* — Each is separated by the value you enter. You can enter any text as a separator, for example **'***'**. If using a custom separator, be aware that text displays exactly as entered, with no additional white space to further separate the characters.
Minimum Heading Level **(minLevel)**	1	Select the largest heading level to start your TOC list. For example, 2 will list h2, h3, and h4 headings, but will not include h1 headings.
Max Heading Level **(maxLevel)**	7	Select the smallest heading level to include in the TOC. For example, 2 will list h1 and h2, but will not include h3 and below.
Include Headings **(include)**		Filter the included headings according to specific criteria. You can use wildcard characters. See Sun's Regex documentation for examples of constructing regular expression strings.
Exclude Headings **(exclude)**		Exclude headings according to specific criteria. You can use wildcard characters. See Sun's Regex documentation for examples of constructing regular expression strings.

Printable (`printable`)	true	By default, the TOC is set to print. If you clear this parameter, the TOC will not be visible when you print the page.
CSS Class Name (`class`)		If you have a custom TOC in your CSS style sheet, you can use this parameter to output the TOC with the specified "class" attribute.

Examples

The examples are based on a page with the following headings:

- Places
    - Favourite Places
    - Unknown Places
- Foods
    - Favourite Foods
    - Unknown Foods
- Things
    - Favourite Foods
    - Unknown Foods

Filtered Table of Contents

This example will filter all headings to include those that contain "Favourite", but will exclude any heading which ends with the word "Things". The list is styled with upper-case Roman numerals.

Parameter	Value
Output Type	**list**
List Style	`upper-roman`
Include Headings	`Favourite.*`
Exclude Headings	`.*Things`

Screenshot: Filtered TOC 'zone' headings

> I. Favourite Places
> II. Favourite Foods

Flat List

This example will filter all headings to render a flat list of "Unknowns" enclosed in square brackets.

Parameter	Value
Output Type	**flat**
Separator	brackets
Max Heading Level	2

Include Headings	Unknown.*

Screenshot: Filtered TOC 'zone' headings displayed as a flat list

[Unknown Places] [Unknown Foods] [Unknown Foods]

Notes

- Due to an outstanding issue in the Table of Content Zone macro (CONF-10619), the Macro Browser's **Refresh** function does not render any parameter modifications. Currently, the rendering of parameter value modifications to the Table of Content Zone macro occurs only after the page is saved.
- Using HTML heading markup with the Table of Content Zone macro — The Table of Content Zone macro cannot handle HTML heading markup on its own. Hence, if you used the HTML and HTML Include macros to render HTML heading markup in a Confluence page, the Table of Content Zone macro will not create a contents list out of these headings. (For more information on about this issue, please refer to TOC-93.)

  However, if you insert an HTML anchor into each HTML heading on your page (based on the following syntax), the Table of Content Zone macro will incorporate these headings into your contents list.

  ```
  <h2><a name="pagename-headingname"></a>Heading Name</h2>
  ```

  The syntax for the anchor name is the page name and heading name separated by a hyphen. Remove all spaces and convert all text to lower case. Convert all punctuation marks to their URL-encoded equivalent.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: toc-zone

Macro body: Accepts rich text.

Storage format example

```
<ac:structured-macro ac:name="toc-zone">
  <ac:parameter ac:name="printable">false</ac:parameter>
  <ac:parameter ac:name="maxLevel">2</ac:parameter>
  <ac:parameter ac:name="minLevel">2</ac:parameter>
  <ac:parameter ac:name="location">top</ac:parameter>
  <ac:parameter ac:name="type">flat</ac:parameter>
  <ac:parameter ac:name="outline">true</ac:parameter>
  <ac:parameter ac:name="separator">pipe</ac:parameter>
  <ac:rich-text-body>
    <p>Only headings within this block are included in the table of contents.</p>
  </ac:rich-text-body>
</ac:structured-macro>
```

Wiki markup example

```
{toc-zone:printable=false|maxLevel=2|minLevel=2|location=top|type=flat|outline=true
|separator=pipe}
Only headings within this block are included in the table of contents.
{toc-zone}
```

User List Macro

The User List macro displays a list of Confluence users, based on their group membership. The macro can also indicate when users are online or offline.

Using the User List macro

To add the User List macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Group(s) (`groups`)	none	Specify the group name. Specify multiple groups separated by a comma, or use * to show all users in Confluence. See this knowledge base page for more information about controlling which users can see the details of other users.
Display Online/Offline Users (`online`)	All registered users	List online or offline users. Leave blank to show all users, irrespective of status. Accepted values: • Unspecified – The macro will show all registered users. • `true` – The macro will show only online users. • `false` – The macro will show only offline users. See below for information on how to configure this macro to display online / offline users.

Configuring the User List macro

In order to use the **Display Online / Offline Users** parameter to indicate whether users are currently logged in to Confluence, you will need to enable the User Log In Listener component in the add-on that provides this macro. You'll need Confluence Administrator permissions to do this.

To enable the Display Online/Offline Users filter in the User List macro:

1. Go to

   > **Add-ons**
2. Select **System** from the drop down and search for the **User Lister** add-on
3. Expand the add-on and enable the **User Log In Listener** module.
4. Restart Confluence for the change to take effect.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `userlister`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="userlister">
    <ac:parameter ac:name="groups">confluence-users</ac:parameter>
    <ac:parameter ac:name="online">false</ac:parameter>
  </ac:structured-macro>
```

Wiki markup example

```
{userlister:groups=confluence-users|online=false}
```

User Profile Macro

The User Profile macro displays a short summary of a given Confluence user's profile. This is the same summary that appears in a Hover Profile, which appears whenever you mouse over a user's name in the Confluence interface.

Screenshot: Example of the User Profile macro

Using the User Profile macro

To add the User Profile macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Default	Description
Username (`user`)	none	The username of the Confluence user whose profile summary you wish to show.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `profile`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="profile">
  <ac:parameter ac:name="user">
    <ri:user ri:userkey="12345678912345678912345678912345"/>
  </ac:parameter>
</ac:structured-macro>
```

Wiki markup example

```
{profile:user=admin}
```

View File Macro

> **The View File macros are going away in a future release**
>
> Read more about this change or find out about a better way to insert Office and PDF files
> .

The View File macros allow you to embed an Office or PDF document on a page. First attach the document to a page and then use one of the View File macros to display the document's content.

When people view the page, they will see the content of the Office or PDF document. They do not need to have Office installed in order to see the content of the file.

For a full list of Office Connector prerequisites and limitations, please refer to:

- Office Connector Prerequisites
- Office Connector Limitations and Known Issue

Displaying an Office or PDF Document in Confluence

To display an Office or PDF document in a page, use one of the following View File macros in the macro browser:

- Office Excel

- Office PowerPoint
- Office Word
- PDF

To add one of the View File macros to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters for the Office and PDF macros

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Macro	Parameter	Default	Description
All View File macros	**Page Name**	The page which contains the macro	Enter a page name, if you wish to display a document which is attached to another Confluence page.
	File Name	none	The file name of the Office or PDF document to be displayed. The document must be attached to a page on your Confluence site.
Office Excel	**Show Grid?**	true	Select to show grid lines around each cell of the Excel spreadsheet. Clear to hide these grid lines.
	Worksheet Name	Last worksheet viewed in the spreadsheet	The name of the worksheet that you want displayed.
	Last Row	Last row with content	The number of the last row you want displayed, starting from '0' as the first row.
	Last Column	Last column with content	The number of the last column you want displayed, starting from '0' as the first column. **Hint for reducing the size of the spreadsheet:** Use the **Last Column** and **Last Row** parameters to reduce the size of the spreadsheet displayed on the wiki page. This is especially useful to prevent the display from showing empty cells. This will also help to prevent 'out of memory' errors.
Office PowerPoint	**Height**		Specify the height of the display, in pixels (default) or as a percentage of the window's height.
	Slide Number	none	Specify the number of the slide that you want displayed on the Confluence page, where the first slide is numbered zero. Instead of a slide show, the page will display just the single slide, represented as a JPEG image. If not specified, all slides display as a slideshow.
	Width		Specify the width of the display, in pixels (default) or as a percentage of the window's width.

Editing an Office document

You can launch your Office application and edit Office documents displaying in the view file macros directly from your Confluence page.

- **Word** and **Excel** - choose the **Edit Document** link above the content
- **PowerPoint** - choose the edit icon on the viewer.

You will find more information and other methods for editing attached Office documents in Edit Office Files.

Troubleshooting

Problems? Please refer to our guide to the Office Connector limitations and known issues.

Widget Connector Macro

Embed online videos, slideshows, photostreams and more directly into your page with the Widget Connector macro.

The macro currently supports content from these sites:

- YouTube
- Vimeo
- MySpace Video
- Flickr
- Twitter
- Slide Rocket
- Google Calendar

On this page:

• Using the Widget Connector
• Parameters
• Examples
  • YouTube
  • Vimeo
  • MySpace Video
  • Flickr
  • Twitter
  • SlideRocket
  • Google Calendar
• Troubleshooting
• Code examples
  • Storage format example

Using the Widget Connector

To add the Widget Connector to a page:

1. In the Confluence editor, choose **Insert** > **Oth er Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Typ e { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplet e for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Description
Web Site's Widget URL (`url`)	This is the external site's URL. In some sites this will be the URL shown in the address bar of your browser, and in other sites you may need to click a Share or Link button to get the URL.

Pixel Height (Value Only) **(height)**	The height of the display, in pixels.
Pixel Width (Value Only) **(width)**	The width of the display, in pixels.

Examples

Every site is a little different, so we've put together some info on what you'll need to do to embed each site's content on a page.

YouTube

The fastest way to embed a YouTube video is to paste the URL into the editor. Confluence will autoconvert the link and insert the macro for you, like magic.   Autoconvert works with both long and short YouTube URLs.

If you're pasting the URL into the Widget Connector macro URL field manually, you'll need to use the long URL (from the address bar). Long URLs look something like this `https://www.yo utube.com/watch?v=k6lK5hlB1nQ`.

- If you're not able to see the video in some browsers, try using https rather than http in your link.
- Links that contain a parameter to start a video at a particular time won't autoconvert or work in the Widget Connector macro, like this link: `https://www.yout ube.com/watch?t=15&v=LhHKkodOPFo`. Paste in the short sharing URL to be sure it works.

Vimeo

The fastest way to embed a Vimeo video is to paste the URL into the editor. Confluence will autoconvert the link and insert the macro for you.

You can use the URL from the address bar in your browser or the **Share** button in Vimeo.

MySpace Video

We don't autoconvert MySpace videos, so you'll need to add the Widget Connector macro to your page first, and then paste the video's URL into the URL field.

You can use the URL from the address bar in your browser or the **Sharing options** in MySpace.

Flickr

You can embed a single photo, a set, an entire user's photo stream or collection of photos with a particular tag.

The fastest way is to paste a URL into the editor. This works for sets and tags. For other URLs you'll need to add the Widget Connector macro to the page first and then paste your link into the URL field.

Use the URL from the address bar in your browser.

Twitter

We don't autoconvert Twitter, so you'll need to add the Widget Connector macro to your page first, and then paste the URL into the URL field.

To embed a single tweet you'll need to click the **Details** link on the tweet so you can grab the URL for just that tweet. The URL will look something like `https://twitter.com/atlassian/status/346976521250037760`.

To embed a dynamic list of tweets you'll need to create a Widget in Twitter first:

1. In Twitter, go to **Settings** > **Widgets** and create a widget to display the tweets you want to embed (for example, a user timeline, list of tweets or hashtag search).
2. Save the widget, then copy the page URL. The URL should look something like this `https://twitter.com/settings/widgets/354381809263472640` (without `/edit` on the end).
3. In Confluence, paste the URL into the Widget Connector.
4. The list of tweets will display, like this one below.

SlideRocket

We don't autoconvert SlideRocket presentations, so you'll need to add the Widget Connector macro to your page first, and then paste the video's URL into the URL field.

You can get the URL from the **Sharing** button on the SlideRocket player.

Google Calendar

We don't autoconvert Google Calendars, so you'll need to add the Widget Connector macro to your page first, and then paste the calender address into the URL field.

You can only embed public calendars. To get your calendar's URL, in Google Calendar go to **Calendar Settings** > **Calendar Address** and click the **HTML** button.

The URL will look something like this: `https://www.google.com/calendar/embed?src=en.austral ian%23holiday%40group.v.calendar.google.com&ctz=Australia/Sydney`

Troubleshooting

If the Widget Connector can't display content from the external site, the macro will look like this:

example.com

We rely on the external website's APIs to display content in the Widget Connector macro. APIs do change from time to time and this can cause the Widget Connector macro to stop rendering content.

If you experience problems, you can raise an issue about it to let us know.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `widget`

Macro body: None.

Storage format example

```
<ac:structured-macro ac:name="widget">
    <ac:parameter ac:name="height">480</ac:parameter>
    <ac:parameter ac:name="width">640</ac:parameter>
    <ac:parameter ac:name="url">
      <ri:url ri:value="http://youtube.com/watch?v=23pLByj_q5U"/>
    </ac:parameter>
  </ac:structured-macro>
```

Task Report Macro

Use the Task Report macro to display a list of tasks on a page. Filter the tasks by space, page, user, label, created date and more.

See Add, Assign, and View Tasks for more information on creating and assigning tasks. You can also use the Task Report blueprint, which will create a page and add this macro for you.

On this page:

- Using the Task Report macro
- Parameters
- Code examples

Related pages:

- Add, Assign, and View Tasks

Using the Task Report macro

To add the Content Report Table macro to a page:

1. In the Confluence editor, choose **Insert** > **Other Macros**
2. Find and select the required macro

Speeding up macro entry with autocomplete: Type { and the beginning of the macro name, to see a list of suggested macros. Details are in Autocomplete for links, files, macros and mentions.

To edit an existing macro: Click the macro placeholder and choose **Edit**. A macro dialog window will open, where you can edit the parameters of the macro.
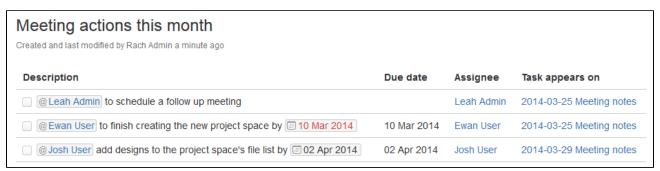
Task Report macro showing incomplete tasks, on pages with the label 'meeting-notes'.

Parameters

Parameters are options that you can set to control the content or format of the macro output. Where the parameter name used in Confluence storage format or wikimarkup is different to the label used in the macro browser, it will be listed below in brackets (`example`).

Parameter	Required	Default	Description

Space(s) and Page(s) (`spaceAndPage`)	No	None	Filter by the task location. The macro will only display tasks in the pages or spaces specified. You can enter a combination of spaces and pages.
Label(s) (`labels`)	No	None	Filter by Label. The macro will only display tasks on pages with this label. You can enter multiple labels, separated by a comma.
Assigned to (`assignee`)	No	None	Filter by Assignee. The macro will only display tasks assigned to the users specified.
Created by (`creator`)	No	None	Filter by Creator. The macro will only display tasks created by the users specified.
Created after (`createddateFrom`)	No	None	Filter by created date. The macro will only display tasks created on or after the date specified. Date must be entered as `dd-mm-yyyy`.
Task status (`status`)	Yes	Incomplete	Show complete or incomplete tasks.
Number of tasks to display (`pageSize`)	No	20	The number of tasks to display on each page of results in the table. Choose from 10, 20 or 40.
Display columns (`columns`)	No	description,duedate,assignee,location	Columns to include in the table. Available columns include description, duedate, assignee, location, completedate and labels.
Sort by (`sortBy`)	No	Due date	Sort tasks by due date, assignee or page title. Select the Reverse Sort check box to sort the table in reverse order.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `tasks-report-macro`

Macro body: None.

Storage format example

Note: The `spaceAndPage` multi-picker parameter populates the `spaces` and `pages` parameters.

```
<ac:structured-macro ac:name="tasks-report-macro">
         <ac:parameter ac:name="spaces">BSP</ac:parameter>
         <ac:parameter ac:name="status">complete</ac:parameter>
         <ac:parameter ac:name="assignees">
           <ri:user ri:userkey="2c9682654373c47f014373caa68a0003"/>
         </ac:parameter>
         <ac:parameter
ac:name="spaceAndPage">space:BSP,page:2555905</ac:parameter>
         <ac:parameter ac:name="pages">2555905</ac:parameter>
         <ac:parameter ac:name="labels">meeting-notes</ac:parameter>
         <ac:parameter ac:name="pageSize">10</ac:parameter>
      </ac:structured-macro>
```

Wiki markup

Wiki markup is not available for this macro. You cannot add this macro via wiki markup.

JIRA Chart Macro

If your Confluence site is connected to a JIRA application, you can display issues as charts on Confluence pages. You can connect Confluence to any JIRA application, including JIRA Software and JIRA Service Desk.

Before you can use this macro, your Confluence and JIRA applications must be connected via Application Links. People viewing the page will see charts for publicly accessible issues. If your JIRA application has restricted viewing (that is, people need permission to view issues) then they'll need to authenticate before seeing the charts. This macro is compatible with JIRA 5.x and later.

On this page:

- Adding the JIRA Chart macro to a page
- Pie chart
- Created vs Resolved chart
- Two Dimensional Chart
- Disabling the JIRA Chart macro
- Notes
- Code examples

Adding the JIRA Chart macro to a page

To add a JIRA chart to your page:

1. In the editor choose **Insert** > **Other macros** > **JIRA Chart**.
2. Choose the type of report you want to create (for example Pie, Created vs Resolved)
3. Select your JIRA server.
   If you have multiple JIRA servers linked to Confluence the drop down will default to the primary application link.
4. Search for issues - you can enter the query in JQL or paste a JIRA URL directly into the search field.
5. Choose **Preview** to generate the chart.
6. Choose **Display Options** to further control how your chart appears.
7. Choose **Insert**.

To find out more about searching for issues see Displaying issues using JIRA Query Language (JQL).

Screenshot: The JIRA Chart Macro in the macro browser

Insert JIRA Chart

Pie Chart

More to come...

Project = 'CONF' AND fixVersion = '5.4'   JAC   Preview

Search using any issue key, search URL, JIRA link, JQL or plain text

> Resolved = 19
> Awaiting Merge = 1
> Quality Review = 1

⌄ Display options

Chart by   Status

Width   90%

Enter a width in pixels, percent or leave blank to automatically resize

☐ Show border

☐ Show chart information

Select Macro   Insert   Cancel

Pie chart

Pie charts can be used to report on issue status, priority, assignee and more.

To further control how this chart appears on your page. Choose **Display options**:

- **Chart by** - select the field you want to segment the pie chart by such as:
  - Status
  - Fix version
  - Assignee name
  - Priority
  - Component
  - Issue type
- **Width** - define the total width of the chart area. You can enter values in pixels, percent or leave blank to auto fit.
- **Show border** - add a border around the chart area.
- **Show chart information** - include a text summary under the chart with the total issues count and chart by value.

> Resolved = 830
> Open = 32
> New = 23

Created vs Resolved chart

Created vs Resolved charts can be used to show the difference between the number of issues created versus the number of issues resolved over time.

To further control how this chart appears on your page choose **Display options**:

- **Period** - choose a time frame to report by (week, month, quarter etc).
- **Days previously** - the total number

of days to report on (counting back from today).
- **Cumulative totals** - choose to progressively add totals or report individual values for each period.
- **Show unresolved trend** - add a subplot showing unresolved issues over time.
- **Show versions** - indicate version release dates as a vertical line on the chart.
- **Width** - define the total width of the chart area. Enter values in pixels, percent or leave blank to auto fit.
- **Show border** - add a border around the chart area
- **Show chart information** - include a text summary under the chart with the total issues count and chart by value.

Two Dimensional Chart

Two Dimensional Charts can be used to show issue statistics in a matrix. You can configure the X and Y axes to display these issue fields:

- Status
- Priority
- Assignee
- Fix version
- Component
- Issue type.

Issue Type	Status					T:
	OPEN	RESOLVED	CLOSED	TECHNICAL REVIEW	NEW	
● Bug	13	121	41	1	0	**176**
↗ Improvement	3	105	30	0	3	**141**
⊞ New Feature	1	11	2	0	1	**15**
⊗ Sub-task	0	169	430	0	0	**599**
Total Unique Issues:	**28**	**836**	**652**	**1**	**22**	**1539**

Showing **4** of **6** statistics. Show more
View in JIRA

For example you could use the chart to show issue types by status (as shown above).

To configure the chart axes choose **Display options**:

- **X Axis** - the issue field to display on the X axis (columns).
- **Y Axis** - the issue field to display on the Y axis (rows).
- **Rows to display** - the maximum number to display in the chart.

Disabling the JIRA Chart macro

The functionality is provided by an add-on (plugin) called 'JIRA Macros'. This macro is also used for the JIRA Issues macro. To make the macro unavailable on your site, you can disable the add-on. See Disabling and enabling add-ons.

Notes

HTTPS: The JIRA Chart macro can access a JIRA site running under SSL provided the Confluence server is configured to accept the JIRA SSL certificate. See Connecting to LDAP or JIRA applications or Other Services via SSL.

Authentication: If the query includes issues that require authentication (issues that are not visible to anonymous users in JIRA), users will be prompted to authenticate to view charts on the Confluence page.

In order to search for issues in the macro browser you may need to authenticate. With JIRA 5.x you will be able to search for unrestricted issues as an anonymous user, however with JIRA 6.x you must be

authenticated to search for any issues.

Code examples

The following examples are provided for advanced users who want to inspect or edit the underlying markup for a Confluence page.

Macro name: `jirachart`

Macro body: None.

The following parameters are available in storage format. Note that some parameters are only available for specific chart types.

Parameter name	Chart type	Required	Default	Parameter description and accepted values
`chartType`	all	Yes	Pie	Type of chart to display. Chart types include: • `pie` • `createdvsresolved` • `twodimensional`
`showinfor`	all	No	False	Displays text information about the data below the chart. Includes the Total value and the Chart By value.
`jql`	all	Yes		JQL query for the chart to display.
`width`	all	600	(blank)	This is the total width of the chart area. Width can be entered in pixels, percent or left blank to fit to the available space.
`border`	all	No	False	Displays a border around the chart area.

statType	pie	Yes	Statuses	The JIRA field to segment the pie chart by: `statuses` - displays a breakdown of issues by **Status**. `allFixfor` - breakdown of issues by chart by all **Fix Versions** (useful if issues have more than one fix version). `assignees` - breakdown of issues by the **Assignee** name. `priorities` - breakdown of issues by **Priority**. `components` - breakdown of issues by **Component**. `issuetype` - breakdown of issues by **Issue Type**.
periodName	createdvsresolved	Yes		Time frame to report by. Options are `hourly`, `daily`, `weekly`, `monthly`, `quarterly` and `yearly`. There is a limit to the number of days the report can generate for each period. For example you can report hourly for a maximum of ten days, and report daily for a maximum of 300 days.
daysprevious	createdvsresolved	Yes	False	Number of days, counting backwards from today, to include in the report. Note that maximum days limits might apply depending on the Period you have chosen.
isCumulative	createdvsresolved	No	False	If `true`, the chart will progressively add data. If `false` the chart will show individual values for each period.
showUnresolvedTrend	createdvsresolved	No	False	Include a subplot showing unresolved issues trend.
versionLabel	createdvsresolved	No	None	Indicate version release dates with a vertical line. Values include `major`, `all` and `none`.

xstattype	twodimensional	Yes	Status	JIRA field for the X axis (columns) of the chart.
				These are the same issue fields that are available in the `statType` parameter (status, priority, assignee, etc).
ystattype	twodimensional	Yes	Assignee	JIRA field for the Y axis (rows) of the chart.
				These are the same issue fields that are available in the `statType` parameter (status, priority, assignee, etc).
numberToShow	twodimensional	No	5	Maximum number of rows to include in the chart.

Storage format example - Pie Chart

Example chart by Issue Type - note that Confluence will insert the `server` and `serverId` parameters, based on settings in Application Links:

```
<ac:structured-macro ac:name="jirachart">
  <ac:parameter ac:name="chartType">pie</ac:parameter>
  <ac:parameter ac:name="statType">issuetype</ac:parameter>
  <ac:parameter ac:name="showinfor">true</ac:parameter>
  <ac:parameter
ac:name="jql">project%20%3D%20CONF%20and%20fixVersion%20in%20('5.3')</ac:paramete
r>
  <ac:parameter ac:name="width">600</ac:parameter>
  <ac:parameter ac:name="server">JAC</ac:parameter>
  <ac:parameter
ac:name="serverId">144880e9-a353-312f-9412-e5028e8166fa</ac:parameter>
  <ac:parameter ac:name="border">true</ac:parameter>
</ac:structured-macro>
```

Storage format example - Created vs Resolved

Example chart - note that Confluence will insert the `server` and `serverId` parameters, based on settings in Application Links:

```
<p><ac:structured-macro ac:name="jirachart">
  <ac:parameter ac:name="isAuthenticated">true</ac:parameter>
  <ac:parameter ac:name="versionLabel">all</ac:parameter>
  <ac:parameter
ac:name="jql">project%20%3D%20confluence%20and%20component%20%3D%20documentation<
/ac:parameter>
  <ac:parameter ac:name="showUnresolvedTrend">true</ac:parameter>
  <ac:parameter ac:name="width">450</ac:parameter>
  <ac:parameter ac:name="daysprevious">30</ac:parameter>
  <ac:parameter ac:name="isCumulative">true</ac:parameter>
  <ac:parameter ac:name="chartType">createdvsresolved</ac:parameter>
  <ac:parameter ac:name="showinfor">false</ac:parameter>
  <ac:parameter ac:name="server">JAC</ac:parameter>
  <ac:parameter
ac:name="serverId">144880e9-a353-312f-9412-e5028e8166fa</ac:parameter>
  <ac:parameter ac:name="border">true</ac:parameter>
  <ac:parameter ac:name="periodName">daily</ac:parameter>
</ac:structured-macro></p>
```

Storage format example - Two Dimensional Chart

Example chart - note that Confluence will insert the `server` and `serverId` parameters, based on settings in Application Links:

```
<ac:structured-macro ac:name="jirachart">
  <ac:parameter ac:name="chartType">twodimensional</ac:parameter>
  <ac:parameter ac:name="ystattype">issuetype</ac:parameter>
  <ac:parameter ac:name="sortBy" />
  <ac:parameter ac:name="isAuthenticated">true</ac:parameter>
  <ac:parameter
ac:name="jql">project%20%3D%20CONF%20and%20Component%20%3D%20Documentation</ac:pa
rameter>
  <ac:parameter ac:name="width" /><ac:parameter
ac:name="server">JAC</ac:parameter>
  <ac:parameter ac:name="numberToShow">10</ac:parameter>
  <ac:parameter ac:name="sortDirection" />
  <ac:parameter
ac:name="serverId">144880e9-a353-312f-9412-ed528e8166fa</ac:parameter>
  <ac:parameter ac:name="xstattype">statuses</ac:parameter>
</ac:structured-macro>
```

Your Profile and Settings

Confluence is very flexible – not only in the many ways you can create and share content, but also in how you can tailor your own Confluence experience. Things like your profile picture, favourite spaces and pages, and your personal space can say a lot about you, and can also make navigating Confluence much quicker and easier. Even a simple thing like adding shortcut links to the sidebar of your personal space, can save you a lot of time in finding the things you use all the time.

Set up your personal space, and take a look at any of the pages below, to start making Confluence feel like home.

- Your User Profile
- Change Your Password
- Edit Your User Settings
- Set Your Profile Picture
- Choose Your Home Page

Related pages:

- Watch Pages, Spaces and Blogs
- Create a Personal Space
- Favourite Pages

- Favourite Pages
- View and Revoke OAuth Access Tokens

Your User Profile

Your user profile contains basic information about you, which other Confluence users can see. It's also displayed to other users when they click your name in the People Directory, if you haven't set up your personal space .

In your own profile, you can access account management features and update information about yourself, like your name, email address, and password. You can also view other users' profiles.

Find your user profile

To find your user profile:

Choose your profile picture at top right of the screen, then choose **Profile**, or choose the **Profile** link in the sidebar of your personal space.

To find someone else's user profile:

Hover your mouse pointer over a user's linked name or profile picture and choose the user's linked name to open their user profile. Alternatively, you can choose the **Profile** link in the sidebar of their personal space, or go directly to this URL:

```
http://your.confluence.site/users/viewuserprofile.action?username=USERNAME
```

Screenshot: User profile screen for the current user

From your user profile, you can access the following:

Profile	• View and edit your personal details, such as your name and email address details and optionally, your photograph and other personal information. Note that as a security precaution, in order to change your email address, you will be required to re-enter your password. • Upload a profile picture (optional). • Change your password.
Network	• View the recent activity of users that you are following via the Network view. • Follow other users from this view.
Favourites	• View a list of your favourite spaces.
Watches	• View a list of the pages and spaces you are currently watching.
Drafts	• Retrieve any pages you were in the process of editing. See Working with Drafts.
Settings	• Edit your General Settings (homepage, language and timezone). • Subscribe to email notifications. • View and revoke your OAuth access tokens.

Edit your user profile

1. Choose your profile picture at top right of the screen, then choose **Profile**
   Or, choose the **Profile** link in the sidebar of your personal space.
2. Choose **Edit Profile**
3. Enter details about yourself in the form displayed
4. Choose **Save**

Fields in your user profile:

Detail	Description
Full Name	Your name as you'd like it to appear in your profile.
Email	Your email address that will be used to send you mail notifications.
Phone	Your phone number.
IM	Your Instant Messenger (IM) details. To suit a variety of IM applications, this option accepts any string value. For example, you can enter IM details in the form of an email address, or a user ID, like '123456789'.
Website	Your website's URL.
About me	Information about yourself that other users can view (such as your professional information, hobbies, and other interests). You can use Confluence wiki markup in this field.
Position	Your title or position within your organisation.
Department	The name of your department or team.
Location	Your location. This can be your town, city, region or country.

> **Handy Hint**
> Confluence administrators can configure Confluence to mask email addresses (e.g. 'example at atlassian dot com'), protecting your email address from search engine spiders and the like.

Notes

The 'Administer User' link is visible to Confluence administrators only. The administrator can click this link to go directly to the user management screen in the Administration Console.

Change Your Password

There are two scenarios where you may want to change your Confluence password:

- You're logged in, but you want or need to change your password
- You've forgotten your password and can't log in, so you need to reset your password

Change your password when you're logged in:

1. Choose your profile picture at top right of the screen, then choose **Pr ofile**
2. On your **Profile** tab, click **Password** in the left-hand column
3. Enter your current password and your new password in the form displayed
4. Click **Submit**

Reset your password from the login page:

If you've forgotten your password and need to reset it, you can do so from the Confluence login page. Choose the 'Forgot your password?' link and Confluence will step you through the process to reset your password.

Edit Your User Settings

If you want to make Confluence fit *you*, like a well-worn pair of sneakers, you can set some preferences that will make you feel more at home:

- General preferences such as home page, language and time zone
- Editor settings
- Email settings for subscriptions to email reports.
- OAuth access tokens that you have granted from your Confluence user account.

General User Preferences

To edit your general user settings:

1. Choose **your profile picture** at top right of the screen, then choose **S ettings**
2. Choose **Edit** and update the settings
3. Choose **Submit**

Setting	Description
Site Homepage	Select the page that you would like to see whenever you log into Confluence.
Language	Select your language. See below.

Time zone	Select your time zone.
Use Keyboard Shortcuts	Enable keyboard shortcuts, other than for the editor.
Text select	Turn off the popup options panel when highlighting text.

Screenshot: Editing your user profile settings

More about Language

Setting your language preference in your user profile is described in the section above. This section gives more information about that setting and other settings that affect the language Confluence will use.

Individual users can choose the language that Confluence will use to display screen text and messages. Note that the list of supported languages depends on the language packs installed on your Confluence site.

The language used for your session will depend on the settings below, in the following order of priority from highest to lowest:

- The language preference defined in your user profile. Note that you need to be logged in for this setting to take effect.
- The language that you choose by clicking an option at the bottom of the Confluence login screen. Confluence stores this value in a cookie. When the cookie expires, the setting will expire too.
- The language set in your browser. The browser sends a header with a prioritised list of languages. Confluence will use the first supported language in that list. Your Confluence administrator can disable this option by setting a system property.
- The default language for your site, as defined by your Confluence site administrator.

Editor Preferences

You can set some options that determine the way the Confluence editor works. Note that these settings

affect only you. Other people using Confluence can enable or disable the settings on their user profiles independently.

To change your editor preferences:

1. Choose **your profile picture** at top right of the screen, then choose **Settings**
2. Click **Editor** under 'Your Settings' in the left-hand panel
3. Click **Edit** and make your changes
4. Click **Submit**

Setting	Description
Disable Autocomplete	Select to disable autocompletion when you press one of the trigger characters.
Disable Autoformatting	Select to disable autoformatting when you type wiki markup in the editor. Click **?** on the editor toolbar to learn more.

Screenshot: User settings for the editor

Set Your Profile Picture

Your profile picture is used as the icon for your personal space, to represent you in the People Directory, and to illustrate your comments. It also appears in various other places next to your name, such as in the list of recent updates on the dashboard.

When you upload your profile picture, you can resize and reposition it to make sure it looks great.

Upload and adjust your profile picture:

1. Choose your profile picture at top right of the screen, then choose **Pr ofile**
2. Choose **Picture** on the left
3. Choose **Upload image** > **Upload an image**
4. Locate and select the picture on your computer or file server
5. Adjust the size and position of your photo, then choose **Save**

Related pages:

- Your User Profile
- Create a Personal Space
- Your Profile and Settings

Screenshot: Choosing a profile picture

Screenshot: Resize your profile picture

Choose Your Home Page

The dashboard is the default landing page when you log into Confluence. It gives you easy access to what's happening in your site, and helps you get back to pages you recently viewed and worked on.

You can choose to personalise your experience. and use an existing space home page as your landing page.

To set your home page:

1. Choose **your profile picture** at top right of the screen, then choose **Settings**
2. Choose **Edit**
3. Choose an option from the **Site Homepage** drop down.
   Only spaces you're allowed to view will appear.
4. Choose **Submit**.

You'll be directed to your new home page the next time you log in. You can change your personal home page at any time.

Alternatively, if your Confluence administrator has set a space home page as the landing page for the whole site, you can choose **Dashboard** from the **Site Homepage** drop down to use the dashboard as your landing page.

Screenshot: Profile Settings

ℹ️ You can access the dashboard at any time using the dashboard URL. It'll look something like this: `https://yoursite.com/wiki/dashboard.action`.

Favourite Pages

Marking pages favourites means you'll be able to access them quickly from the dashboard or from your profile. It's a great way to get to the content you need to access regularly.

On this page:

- Add a page to your favourites
- View your favourite pages

Add a page to your favourites

To add a page as a favourite, hit the **Favourite** button at the top of the page.

The star icon will change to dark grey to indicate the page is in your favourites list. Hit the button again if you want to remove the page from your favourites.

View your favourite pages

To view your favourite pages:

- Choose **Favourites** on the dashboard sidebar.
- Choose **your profile picture** at top-right of the screen, then choose **Favourites** – there's a list of your favourite pages, and spaces that you've added to **My spaces**.

You can also use the Favourite Pages Macro to include a list of your favourite pages on any page.

Screenshot: Viewing and removing favourites from the dashboard

View and Revoke OAuth Access Tokens

OAuth access tokens allow you to use a Confluence gadget on an external web application or website (also known as the 'consumer') *and* grant this gadget access to Confluence data which is restricted or privy to your Confluence user account.

OAuth access tokens will only appear in your user profile if the following conditions have been met:

1. Your Confluence Administrator has established an OAuth relationship between your Confluence site and the consumer.
   Confluence Administrators should refer to Configuring OAuth for more information about establishing these OAuth relationships.
2. You have accessed a Confluence gadget on the consumer and have conducted the following tasks:
   a. Logged in to your Confluence user account via the gadget and then,
   b. Clicked the 'Approve Access' button to allow the gadget access to data that is privy to your Confluence user account.
      ℹ Confluence will then send the consumer an OAuth 'access token', which is specific to this gadget. You can view the details of this access token from your Confluence site's user account.

An OAuth access token acts as a type of 'key'. As long as the consumer is in possession of this access token, the Confluence gadget on the consumer will be able to access Confluence data that is both publicly available and privy to your Confluence user account. As a Confluence user, you can revoke this access token at any time. Furthermore, all access tokens expire after seven days. Once the access token is revoked or has expired, the Confluence gadget will only have access to publicly available Confluence data.

View your OAuth Access Tokens

To view all of your Confluence user account's OAuth access tokens:

1. Choose **your profile picture** at top right of the screen, then choose **Settings**
2. Click **View OAuth Access Tokens**. A view similar to screenshot below is displayed. Refer to OAuth Access Token Details below for information on interpreting this table.
   ℹ If no access tokens have been set, then 'None specified' is shown.

Screenshot: Viewing your OAuth Access Tokens

On this page:
- View your OAuth Access Tokens
- Revoke your OAuth Access Tokens

Related pages:
- Configuring OAuth for Confluence Admins

Authorised Applications
The following applications are using your account to access Confluence data:

Consumer	Consumer Description	Issued on	Actions
Atlassian JIRA Special Projects	Atlassian JIRA at https://atlassian.com/jira	Jan 23, 2013 expires in 1,801 days	Revoke Access

OAuth Access Token Details

Your list of OAuth access tokens is presented in a tabular format, with each access token presented in separate rows and each property of these tokens presented in a separate columns:

Column Name	Description
Consumer	The name of the Confluence gadget that was added on the consumer.
Consumer Description	A description of this consumer application. This information would have been obtained from the consumer's own OAuth settings when an OAuth relationship was established between Confluence and that consumer. If the consumer is another Atlassian application, this information is obtained from the Consumer Info tab's 'Description' field of the OAuth Administration settings. The application's administrator can customise this Consumer Info detail.
Issued On	The date on which the OAuth access token was issued to the consumer by Confluence. This would have occurred immediately after you approved this gadget access to your Confluence data (privy to your Confluence user account).
Expires On	The date when the OAuth access token expires. This is seven days after the 'Issued On' date. When this date is reached, the access token will be automatically removed from this list.
Actions	The functionality for revoking the access token.

Revoke your OAuth Access Tokens

To revoke one of your OAuth access tokens:

1. View your Confluence user account's OAuth access tokens (described above).
2. Locate the Confluence gadget whose OAuth access token you wish to revoke and click **Revoke OAuth Access Token** next to it.
   The gadget's access token is revoked and the Confluence gadget on the consumer will only have access to publicly available Confluence data.

Collaboration

Confluence is all about encouraging team collaboration to get the best results, so we've built in a number of ways you can notify other people about content that may be of interest to them.

You can:

- Share a link to a page or blog post via email
- Mention a user when you write a page, blog post, comment, or add a task
- Like a page, blog post or comment

Whenever you mention another user, they'll receive an email notification; if you like a page, blog post, or comment, the author will be notified that you *lik e* the content.

Related pages:

- Pages and Blogs
- Watch Pages, Spaces and Blogs
- Export Content to Word, PDF, HTML and XML

Other users can also find out about changes to content in Confluence by wat ching pages and spaces.

Another way to share Confluence content is by exporting it to other formats such as XML, HTML, Microsoft Word and PDF.

Network Overview

You can Create a *network* of users who are important to you, to make sure you're always up-to-date with their Confluence activity. You might want to follow your boss or teammates, to see what they're working on, or whoever creates the most entertaining blog posts.

When someone's part of your network, you'll be able to see when they:

- Add or edit pages or blog posts
- Comment on a page or blog post or edit existing comments
- Update their user profile

On this page:

- Follow another user
- Access your network view
- Notes

Related pages:

- Network Macro
- Subscribe to a Network RSS Feed
- Email Notification s

Follow another user

You can follow another user by using either their Hover Profile or your Network view.

To follow a user with their **Hover Profile**, hover your mouse over their profile picture when it appears in a page and choose **Follow**.

To follow a user from your Network view:

1. Choose **your profile picture** at top right of the screen, then choose **Network**
   Alternatively, choose **More** in the **Network** section of your profile sidebar.
2. Search for and select the user in the **Following** field
3. Choose **Follow**

If you now refresh or revisit your Network view, the profile picture(s) of the user(s) you just followed will appear within the **Following** list on the right. Their tracked activities will also start appearing in the **Recent Activity** list.

Access your network view

If you want to see what's been happening in your network, access your network view as described above.

You can access another user's Network view using the Hover Profile by choosing **More** > **Network Page**.

Screenshot: Example of the Network view

Notes

- **RSS feeds:** you can subscribe to any Confluence user's network RSS feed and receive summaries on the activities of other users they're following in their network. See Subscribe to a Network RSS Feed.
- **Email notifications:** you can request email notifications of any activity in your network. See Email Notifications.

Likes and Popular Content

Has someone written a good blog post or page on Confluence? Or made a comment you agree with? Click the **Like** button to them know.

👍 Like

When you like a page, blog post or comment, the author of the content receives a notification. If enough people like the content, it'll appear on the **P opular** tab of the dashboard.

Disabling the 'like' feature

The *like* functionality is provided by a plugin called the 'Confluence Like Plugin'. To remove the *like* functionality from your site, disable the plugin.

Disabling notifications when your content is 'liked'

There are two ways to turn the 'someone likes your page' notifications off.

Do either of the following:

- Open an email notification of a like, and click **Manage Notifications**
- Go to `<your confluence URL>/plugins/likes/view-notifications.action`

Mentions

Mentions (often known as @mentions) are a seamless way of drawing someone's attention to a page or comment, or assigning a task to them. When you mention a user, they'll receive a notification by email and in their workbox; if you mention them in a task, the task is assigned to them and appears in their tasks list.

There are two ways to mention someone: using autocomplete, or via the Insert menu in the editor toolbar.

On this page:

- Disabling the 'like' feature
- Disabling notifications when your content is 'liked'

Related pages:

- The dashboard
- Email Notifications
- Network Overview

Use autocomplete

To mention someone using autocomplete, type '**@**' in the editor then start typing their name. Choose the person you want to mention from the list of suggestions.

Screenshot: Autocomplete for mentions

I want to mention someone @ josh

User suggestions

Josh User

Use the Insert menu

If you'd rather use the **Insert** menu, choose **Insert** > **User Mention** then search for and select the user you want to mention.

Notes

- **Disable the user mention feature** – The functionality is provided by a plugin called the 'Confluence Mentions Plugin'. If you need to remove the user mention functionality from your site, you can disable the plugin. See Disabling or Enabling a Plugin.
- **Mentioning groups** – You can only mention individual users. There's a feature request to allow

  mentions for groups: **CONF-23015** - Extend 'Mentions' to work with groups as well **OPEN**
- **Link to a user profile** – You can use a square bracket '[' and a person's name to trigger Confluence autocomplete and link to a person's user profile or personal space. Confluence will send the person a notification just as if you had used @mention (unless the administrator has disabled the user mention feature).

Share a Page or Blog Post

You can use the '**Share**' option in Confluence to email anyone a link to a page or blog post. The recipients can be Confluence users or any email address(es).

To share a link to a Confluence page or blog post:

1. Go to the page or blog post you wish to share
2. Choose **Share**

   ⓘ The **Share** button will only visible if your Confluence administrator has configured an outgoing mail server for your Confluence site.
3. Enter a username, group or email address, and select the appropriate user, group or email address from the list of suggestions
   Repeat this process to add multiple recipients to the list.

4. Enter an optional message into the **Note** box
5. Choose **Share** to send the link via email

In addition to an email, Confluence users will also receive a notification in their Confluence workbox. See Workbox Notifications .

ⓘ To remove a recipient from the list, choose the delete icon to the right of the recipient.

Comment on pages and blog posts

Comments are a great way to bring others into the conversation about a page or blog post. They allow you to remark on content, add important information, ask questions, and generally drive collaboration and teamwork.

You can add a comment at the bottom of any page or blog post, or add an in line comment by highlighting specific text on the page.

Add a page or blog post comment

1. Type your comment in the comment field at the bottom of the page
2. Optionally, choose **Preview** to see how your comment will appear
3. By default, **Watch this page** is ticked (This means you'll start receiving notifications about the page. Uncheck it if you don't want to watch the page.)
4. Choose **Save** (Ctrl+S or +S)

Other users can reply and/or like your comment, and you or a space administrator can edit your comment(s).

Add an inline comment

1. Highlight the text you want to comment on
2. Choose the add comment

    button that appears above the highlighted text
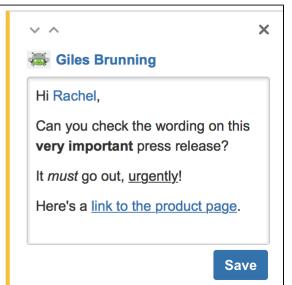3. Type your comment and choose **Save** (Ctrl+S or +S)

The selected text will appear with a yellow highlight indicating an inline comment; choose any highlighted text on the page to display the related comment(s).

Just like page and blog post comments, others can reply to, or like, your inline comments, and you'll be notified when they do.

Resolve inline comments

Hit **Resolve** to hide a set of inline comments once the conversation's finished. If you want to view resolved comments, choose

•••

> **Resolved comments**; to reopen a resolved comment, choose **Reopen** at the bottom left.

Rich comments

Inline and page comments might look simple, but they support rich text (like **bold**, <u>underline</u>, and *italics),* bulleted and numbered lists, links, and @mentions. You can also drop images into any comment, to really illustrate your point.

Link to a comment

You can link directly to a comment on a page. See Links for more information.

> If you don't see a popup when you highlight text, check that Text Select is enabled in your profile settings.

Comment permissions

- **Add a comment** – You need the 'Add Comments' permission in the space.
- **Edit a comment** – You need the 'Add Comments' permission. Space administrators can edit all comments within their space. The date on a comment always indicates the time the comment was last edited.
- **Delete a comment** – You need the 'Remove Comments' permission. Deleted comments cannot be restored. If you don't have the 'Remove Comments' permission, you can delete your own comments, but only if there are no replies to your comment.
- **Disable comments** – If you don't want comments in a particular space, remove the 'Add Comments' permission from the 'confluence-users' or 'users' group, anonymous users and all other users and groups. The option to add comments will no longer appear on pages or blog posts in that space.

See Space permissions for more information. There is no permission that controls comments across the entire site.

Notes

- Choose **Watch** at the top-right of the page to receive an email notification whenever anyone edits or adds a comment to the page.
- On blog posts only, an 'Author' lozenge will appear on any comments made by the original author of the post.
- It's not possible to delete all comments on a page simultaneously, or change the order of comments.

Watch Pages, Spaces and Blogs

You can 'watch' a Confluence page, blog post or space. Confluence will then send you a notification email whenever anyone updates your watched content.

You'll receive email notifications for:

- Edits (unless the author clears the 'Notify watchers' check box).
- Deletions.
- Attachments, including new versions or deletions of an existing attachment.
- Comments, including new comments, edits of existing comments or deletions of existing comments.

By default, Confluence will assign you as a watcher of any page or blog post that you create or edit. This behaviour is called 'autowatch'.

There's no daily digest for email notifications. You'll receive an email

notification every time someone makes a change.

You **will not** receive email notifications for content changes due to the output of a macro, because the page content itself hasn't been edited.

You need 'View' permission for the page, blog post or space to receive notifications.

Watching a page or blog post

To start watching a page or blog post:

1. Go to the page or blog post
2. Choose **Watch** and select the relevant check box

To stop watching the page or post, deselect the relevant check box.

Watching an entire space

You can choose to watch all the pages and blog posts in a particular space.

The quickest way is to use the **Watch** option on a page or blog post, as described above.

To stop watching the space, deselect the relevant check box.

Alternatively, choose **Pages** in the space sidebar, then choose **Watch this space** at the top right.

If your space uses the Documentation theme, choose **Browse** > **Space Operations**. Then choose **Watch this space** in the left-hand menu.

Watch for new blog posts in a space

You can choose to receive a notification whenever someone adds a blog post in the space. You will not receive notification of updates to or deletions of blog posts, nor of comments on the blog posts.

To watch for new posts:

1. Go to a blog post in the space
2. Choose **Watch** and select **Watch all blog posts in this space**

To stop watching for new blog posts, deselect the relevant check box. Alternatively, choose **Blog** in the space sidebar, then choose **Watch this blog** at the top right.

If your space uses the Documentation theme, choose **Browse** > **Space Operations**. Then choose **Watch this blog** in the left-hand menu.

Watch all spaces on the site

You can receive notifications about changes to the content of pages, blog posts and comments from all spaces on a Confluence site.

To start watching for content changes across the whole site:

1. Choose **your profile picture** at top right of the screen, then choose **Settings**
2. Choose **Email**
3. Choose **Edit** then choose **Subscribe to daily updates**
4. Choose **Submit**

Watching for all new blog posts on the site

You can choose to watch for all new blog posts in all spaces on the Confluence site. You will not receive notification of updates to or deletions of blog posts, nor of comments on the blog posts.

To start watching for all new blog posts:

1. Choose **your profile picture** at top right of the screen, then choose **Settings**
2. Choose **Email**
3. Choose **Edit** then choose **Subscribe to all blog posts**
4. Choose **Submit**

Manage watches from your user profile

The '**Watches**' page in your user profile displays a list of all pages and spaces you are currently watching.

To manage your watches:

1. Choose **your profile picture** at top right of the screen, then choose **Watches**
2. Choose **Stop Watching** for any unwanted spaces or pages

Manage watches from the email message

The email notifications that you receive from Confluence have some useful links at the bottom of the email message. The links in each message vary, depending on the context. In general, the links allow you to view the page online, reply to a comment, and so on.

In particular with respect to setting your notification preferences, you will see one or more of the following links:

- **Stop watching page** – Click this link to stop watching the page that triggered the email notification.
- **Stop watching space** – Click this link to stop watching the space that triggered the email notification.
- **Stop following this user** – Click this link to stop following the user whose update triggered the email notification.
- **Manage Notifications** – Click this link to go to the email settings page in your user profile.

Screenshot: Example email notification footer showing links

Autowatch and other notification options

By default, Confluence will assign you as a watcher of any page or blog post that you create or edit. This behaviour is called 'autowatch'. You can turn autowatch on or off, and set other notification options, in the email settings section of your user profile. See Edit Your User Settings.

Manage Watchers

If you're an administrator of a space, you can manage the watchers for all pages and blog posts in that space. This means that you can:

- View a list of the watchers of a page or blog post
- View a list of the watchers of the space that contains the page or blog post
- Add users as watchers of the page or blog post
- Remove existing watchers of the page or blog post
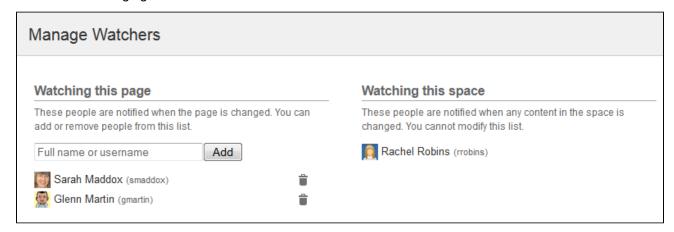
You can't remove watchers of the space.

Related pages:

- Watch Pages, Spaces and Blogs
- Email Notifications
- Watch Pages, Spaces and Blogs
- Your Profile and Settings

To manage the watchers of a page or blog post:

1. Go to the page or blog post for which you want to manage the watchers
2. Choose **Watch** > **Manage Watchers**
   The left-hand column of the 'Manage Watchers' screen shows the users watching the page or blog post. The right-hand column shows the users watching the space.
3. Do either of the following:
   - To remove an existing page watcher, choose the trash can icon next to the user's name.
   - To add a user as a watcher of the page, type their username and choose **Add**.

Screenshot: Managing watchers

Email Notifications

You can 'watch' a page, blog post or space. Confluence will then send you a notification by email whenever anyone adds or updates content on that page or space. You can also subscribe to daily email reports and other notifications of various updates, as described below.

You will only receive notifications for content that you have permission to view. Users that have been disabled by an administrator will not receive email notifications.

Subscribing to email notifications

You can subscribe to the following email summary reports:

- A daily report of updates to all spaces that you have permission to view.
- An immediate report of all blog posts added or changed, in all spaces that you have permission to view.
- An immediate report of all updates made by the people you are following, in all spaces that you have permission to view.
- An immediate report when someone follows you.
- A daily or weekly report of recommended updates, in all spaces that you have permission to view.

To edit your email notification settings:

1. Choose **your profile picture** at top right of the screen, then choose **S ettings**
2. Click **Email** in the left-hand panel
3. Click **Edit**

On this page:
- Subscribing to email notifications
- Notes for administrators

Related pages:
- Watch Pages, Spaces and Blogs
- Subscribe to RSS Feeds within Confluence
- Your Profile and Settings
- Edit Your User Settings

Below is an explanation of all the email settings is explained below.

Setting	Description	Content	Frequency
Autowatch	**Option:** Do you want Confluence to automatically add you as a watcher on each page or blog post that you add or update? If you are a watcher of a page or a post, you will receive notification of future changes.	• Pages and blog posts that you create, edit or comment on.	(Not applicable)
Subscribe to daily updates	Receive email reports showing changes to content in all spaces that you have permission to view. Note: *Daily email reports do not include information about attachments on a page or blog post that are added, edited or deleted.*	• Pages and blog posts that are added, edited or deleted. • Comments on a page or blog post that are added, edited or deleted. • Updates by users who have changed their personal profile.	Daily

Subscribe to all blog posts	Receive email notifications for changes to blogs in your Confluence installation that you have permission to view.	• Blog posts added, edited or deleted.	Immediately
Subscribe to network	Receive email notifications for changes to content by all users that you are following, which you have permission to view.	• Pages being added, edited or deleted. • Blog posts being added, edited or deleted. • Comments being added, edited or deleted.	Immediately
Subscribe to new follower notifications	Receive an email message when anyone chooses to follow you.		Immediately
Notify on my actions	**Option:** Do you want to receive email notifications for your own changes? Note: If you have not subscribed to any email notifications and are not watching any pages/spaces, then selecting 'Notify on my actions' will not do anything.	• All pages and spaces that you are watching. • This affects all subscriptions set.	(Not applicable)
Show changed content	**Option:** Do you want your notifications to include details of the changes made to the content? • If you do not select this option, your notifications will include only the title of the page, and any comment the author made when updating the page. • If you do select this option, your notifications will show the differences between the current and previous versions of the page. See Page History and Page Comparison Views.	• Edits to pages and blog posts.	(Not applicable)

Subscribe to recommended updates	Receive an email message showing the top content that is relevant to you from spaces that you have permission to view. How do you set the frequency of the mail message? A link in the email message allows you to choose daily or weekly notifications. How do you enable and disable the notification? You can turn off the notification by clicking a link in the email message. You can also turn the notification on or off by setting the 'Subscribe to recommended updates' option in your user profile.	Confluence chooses the content to display, based on: • Pages and blog posts that people have recently liked. • Pages and blog posts that people have recently commented on. • Pages and blog posts that have recently been created. 'Recent' means any activity that occurred since the last recommended updates message was sent to you. The activities are listed in order of popularity, with the most popular at the top. Likes, comments and content creations are scored equally. Activity that involves people in your network ranks higher than activity not involving your network. Content from My spaces also ranks higher than content in other spaces. The recommended updates summary does not include any content that you created yourself, and it gives a lower ranking to content that you have participated in, for example by adding a comment or updating the page. If there is no activity to report, Confluence will not send the email message.	Daily (weekdays at 1pm) or weekly

Notes for administrators

- To enable Confluence to send email notifications, a System Administrator must configure an email server. See Configuring a Server for Outgoing Mail.
- Confluence Administrators can set the default options for the recommended updates notification. Choose the **cog icon**

  , then choose **General Configuration** under Confluence Administration. Click **Recommended Updates Email** in the left-hand panel. See Configuring the Recommended Updates Email Notification.

Subscribe to RSS Feeds within Confluence

An RSS feed is a format for delivering summaries of regularly changing web content. Subscribing to an RSS feed allows you to stay informed of the latest content from sites that you are interested in.

RSS is not designed to be read in a regular web browser. Specialised RSS newsreader programs can check RSS files every so often, and tell you what's new on a site. Your reader may be on a website, an add-on to your browser, part of your email program, or a stand-alone program.

Confluence generates its own RSS feeds for tracking updates to content within Confluence. You will need an RSS reader which can grab the RSS feeds from Confluence and display them for you.

Confluence's RSS macro allows you to display the contents of an RSS feed on a Confluence page. The feeds may come from a Confluence feed generator or from external sites. In this way, Confluence can act as an RSS reader.

For a technical description of RSS, read Mark Pilgrim's "What is RSS?" article on XML.com.

On this page:
- Confluence RSS feeds
- RSS newsreaders
- Remove an RSS feed

Related pages:
- Watch Pages, Spaces and Blogs
- RSS Feed Macro

Confluence RSS feeds

RSS feeds allow you to track updates to content within Confluence. You will need an RSS newsreader to read a feed.

You can create a customised RSS feed using the RSS Feed Builder or subscribe to one of the pre-specified feeds generated by Confluence.

What would you like to do?

- Create and subscribe to customised RSS feeds using the RSS Feed Builder – Create a customised RSS feed. For example, you can filter your feed using a label, specify the number of items and days to include in your feed, and so on.
- Subscribe to pre-specified RSS feeds
  – Generate an RSS feed automatically in a minimal number of steps.
- Subscribe to a feed of any Confluence user's network – Track the activities of users the selected person is following.

RSS newsreaders

The following are some popular RSS readers for various operating systems. You can find a more comprehensive list on Google's open directory.

Windows

- SharpReader
- NewsGator
- Syndirella
- FeedDemon
- NewzCrawler

Mac OS X

- Safari
- NetNewsWire
- NewsFire
- Shrook

Multi-platform

- NewsMonster (Runs in the Mozilla web browser)
- Radio Userland (Windows and MacOS)
- AmphetaDesk (Windows, Unix, Mac OS X)

Remove an RSS feed

There is no need to try to delete or remove an RSS feed built by the Confluence RSS feed builder.

Explanation: The feeds generated by the RSS Feed Builder are dynamically generated via the parameters included in the feed URL (address). For example, take a look at the following feed URL:

```
http://confluence.atlassian.com/createrssfeed.action?types=page&sort=modified&showContent=true...
```

The above feed URL will generate a list of pages ('types=page'), sorted by the modification date and showing the page content. The feed is generated at the time when the URL is fetched and there is no RSS feed information stored on the database. For that reason, there is no need to remove anything.

Subscribe to pre-specified RSS feeds

This page tells you how to get hold of an RSS feed which Confluence has predefined for you.

To subscribe to predefined RSS feeds for a particular space:

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar

2. Choose **RSS Feeds**
3. Copy and paste the link for one of the feeds into your RSS newsreader

If your space uses the Documentation theme:

1. Choose **Browse** > **Space Operations**
2. Choose **RSS Feeds** in the space operations options
3. Copy and paste the link for one of the feeds into your RSS newsreader

Feeds include:

- Pages
- Blog
- Mail

- Comments
- Attachments
- All content

To subscribe to predefined RSS feeds for a particular page (where available):

Note that the word 'page' here means a part of the Confluence user interface, rather than a page that contains Confluence content. For example, your Network view offers an RSS feed.

1. Go to the page
2. Locate the following icon, which is available in the top-right corner of certain pages:
3. Copy and paste the icon's link into your RSS newsreader

Notes

The predefined RSS feed will return no more than 10 entries within the last 5 days, if you want to customise your Confluence RSS feed (for example, use a label to filter your feed), use the RSS Feed builder instead of the above instructions.

The RSS Feed Builder

Using the RSS feed builder, you can create customised RSS feeds to subscribe to changes within Confluence.

Wondering what an RSS feed is? See more information about RSS Feeds.

Build an RSS feed

Follow the steps below to build your feed, choosing the type of content and the time period you want to monitor.

To create a customised RSS feed:

1. Choose the help icon

   at top right of the screen, then choose **Feed Builder**
2. Select the content types you want in your feed

   Check **Mail** if you want to know when the email archive is updated. (See the overview of mail archives in Confluence.)
3. Select one or more spaces from the list
4. Click **Advanced Options** to set the following:

Option	Description
Feed Name	The default name is based on the name of your Confluence installation. For example, 'Extranet RSS Feed'.

With these labels	Enter one or more labels separated by spaces or commas. Confluence returns all content (of the selected types) that matches one or more of the labels. See the hint below about using labels to customise your feeds.
Exclude these spaces	Exclude specific spaces from those already selected.
Sorted by	Sort content by either the date or creation or the date they were last updated.
Limit to	Specify the number of items returned in your feed.
Within the last	Specify how old items returned can be.
Include content for pages	Specify whether the entire page is displayed in the feed.

5. Choose **Create RSS Feed**
6. Drag or copy the link into your RSS reader

Hints

- **Separate feeds.** Try building separate feeds, one for pages only and one that includes comments as well. This allows you to monitor only pages if you are short of time, and to read the comments when you have more time.
- **Labels to customise your feed.** You can use the RSS feed builder to track updates to labelled pages and comments on those pages. Here is an idea for customising your RSS feed by using your own personal label(s). This is useful if you want to track updates to specific pages or blog posts, and you do not want to deal with emails. You can use this method as an alternative to watching pages.
    - Build an RSS feed that returns pages, blog posts and comments labelled with a personal label, such as '**my:feed**'.
    - Each time you want to 'watch' a page, just label it with 'my:feed'.
    - All updates and comments will automatically come through your RSS feed.

Notes

- **Removing an RSS feed:**

    There is no need to try to delete or remove an RSS feed built by the Confluence RSS feed builder.

    Explanation: The feeds generated by the RSS Feed Builder are dynamically generated via the parameters included in the feed URL (address). For example, take a look at the following feed URL:

    ```
    http://confluence.atlassian.com/createrssfeed.action?types=page&sort=modifi
    ed&showContent=true...
    ```

    The above feed URL will generate a list of pages ('types=page'), sorted by the modification date and showing the page content. The feed is generated at the time when the URL is fetched and there is no RSS feed information stored on the database. For that reason, there is no need to remove anything.

- **Feed authentication options:** Confluence can offer you the option of an anonymous feed or a feed that requires authentication.
  - An **anonymous** feed will show only the content that is visible to anonymous users. The feed URL does not contain the `&os_authType` parameter mentioned below. This feed is useful only if your Confluence site allows anonymous access. If a feed is anonymous, you only get anonymously-viewable content in the feed regardless of whether you are a Confluence user or not.
  - An **authenticated** feed requires you to log in to Confluence before you can retrieve the content. The feed URL contains the following parameter: `&os_authType=basic`.
  - The option to choose between an anonymous and an authenticated feed is currently not available on the feed builder screen. The feed builder offers only **authenticated** feeds. See CONF-21601 for details and a workaround.

Subscribe to a Network RSS Feed

You can create an RSS Feed from any user's network view, allowing you to receive summaries on the activities of users they are following in their network. The types of activities tracked in these RSS feed summaries include:

- Additions or edits to pages or blog posts
- Comments added to a page or blog post or edits to existing comments
- Updates to a user's profile

On this page:

- Subscribe to a user's network feed
- Customise your network RSS feed
- Notes

Related pages:

- Network Overview
- Subscribe to RSS Feeds within Confluence
- Your Profile and Settings

Subscribe to a user's network feed

To subscribe to a user's network RSS feed:

1. Locate the RSS icon

   , which is available from the top-right of:
   - The 'Recent activity of the users you are following' section of your network page, or
   - The 'Activity of followed users' section of another user's network page.
2. Copy and paste the icon's link into your RSS newsreader

Customise your network RSS feed

Confluence does not provide a way of customising a network RSS feed via the user interface. However, you can modify the maximum number of results and type of content displayed in these feeds by directly editing the RSS feed link in your RSS newsreader.

To modify the maximum number of results displayed in your RSS feed:

1. Edit the RSS feed link in your RSS newsreader.
2. Change the value of the `max` parameter from its default value of `40` to a value of your choice. Example:
   http://confluence.atlassian.com/feeds/network.action?username=MYNAME&max=60&publicFeed=false&os_authType=basic&rssType=atom
3. Save the modified link in your RSS newsreader.

To modify the type of content displayed in your RSS feed:

1. Edit the RSS feed link in your RSS newsreader.
2. Append the parameter `contentType` to the end of the link, followed by an equals sign (`=`) and then add the appropriate content type value of your choice:
   - `PAGE` — restricts the RSS feed to page additions or updates.
   - `BLOG` — restricts the RSS feed to blog post additions or updates.
   - `ATTACHMENT` — restricts the RSS feed to attachment additions or updates.
   - `COMMENT` — restricts the RSS feed to comment additions or updates.

Content type values are case-sensitive. Ensure that each parameter is separated from the other by an ampersand (&).
Example:
http://confluence.atlassian.com/feeds/network.action?username=ggaskell&max=40&publicFeed=false&os_authType=basic&rssType=atom&contentType=BLOG

3. Save the modified link in your RSS newsreader.

Notes

It is not possible to filter for more than one type of content by adding multiple values to the `contentType` parameter.

Workbox Notifications

The Confluence workbox

displays all notifications collected from Confluence page watches, shares, mentions, and tasks. From your workbox you can reply to comments, like a comment or page, watch a page, or open the relevant page or blog post.

If your Confluence site is linked to a JIRA application such as JIRA Software or JIRA Service Desk, you'll also see notifications from your JIRA application in the workbox.

Looking to manage your notification email messages instead? See Email Notifications.

On this page:
- Manage your notifications
- Which notifications are included?
- Keyboard shortcuts
- Manage notifications with Confluence mobile
- Notes

Related pages:

- Configuring Workbox Notifications
- Email Notifications
- Watch Pages, Spaces and Blogs
- Likes and Popular Content

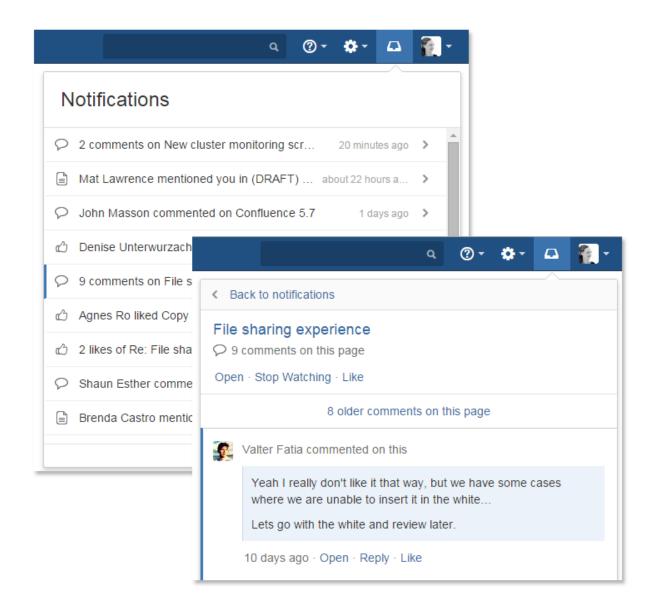**Manage your notifications**

1. Choose the workbox icon

   in the header.
   - A number will appear the workbox icon, to indicate the number of unread notifications waiting for your attention.
   - You can use the keyboard shortcut: Type **g** then **n**. (When in the Confluence editor, click outside the editor before pressing the keyboard shortcut keys.)
2. Choose a notification from the list, to see the notification details. You can then:
   - **Open** the related page, blog post, or comment.
   - **Like** or **Unlike** the page, blog post, or comment.
   - **Watch** or **Stop Watching** to receive notifications, or stop receiving notifications, about a page or blog post.
   - **Reply** a comment, without leaving the workbox.

Screenshot: Your Confluence notifications in the workbox

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.

Keyboard shortcuts

Shortcut	Action
g then n	Open the Confluence workbox.

j	Move down to the next entry in the notification list.
k	Move up to the previous entry in the notification list.
n	Move down to the next notification for a particular page or blog post.
p	Move up to the previous notification for a particular page or blog post.
Enter	Open the selected notification.
u	Return to the notification list after opening a particular notification.

Manage notifications with Confluence mobile

You can also view and respond to notifications on your phone or other mobile device. See Confluence Mobile for more about mobile platforms.

Notes

- Read notifications are automatically deleted after 2 weeks.
- Unread notifications are automatically deleted after 4 weeks.
- You cannot delete your notifications yourself.
- If a new notification arrives while you have workbox open, the count appears on the workbox icon

   but the notification is not added to the workbox. You need to close workbox and re-open it to see the new notification.

- The ability to receive notifications from JIRA or another Confluence site is available in **Confluence 4.3.3 and later.** To receive JIRA notifications, you need **JIRA 5.2 or later**.

- Administrators can enable and disable the workbox on your Confluence site. They can also connect a JIRA site or another Confluence site, so that notifications from those sites appear in your workbox too. See Configuring Workbox Notifications.
- The Confluence workbox is provided by a set of plugins. To remove the personal notifications and tasks functionality from your site, you can disable the following plugins. See Disabling or Enabling a Plugin for instructions. Disabling these plugins will disable the entire workbox . It is not possible to disable only tasks or only notifications:
  - Workbox - Common Plugin
  - Workbox - Host Plugin
  - Workbox - Confluence Provider Plugin
- If you want to re-enable the plugins, do so in the following order: Common Plugin, Host Plugin, Confluence Provider Plugin.
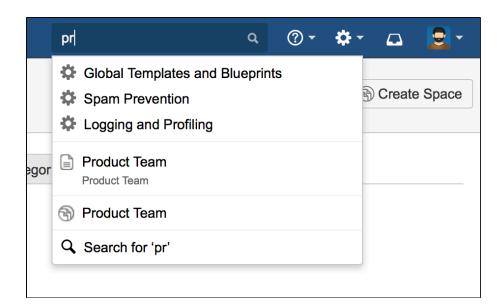- There is no option to disable the workbox for an individual user.

Search

Confluence gives you a few ways to find what you're looking for. Here's an overview of Confluence search, and a few tips to help you find things more easily.

Quick search

To perform a quick search in Confluence, choose the search field at the top-right of every page or type **/** on your keyboard to place your cursor in the search field. Type the name of a page, blog post, person, file/attachment, or space and choose from the list of options displayed.

If you don't immediately see what you need in the quick search, hit **Enter** or choose the **Search for** option at the bottom of the search results to do a full search.

pr|

⚙ Global Templates and Blueprints
⚙ Spam Prevention
⚙ Logging and Profiling

🗎 Product Team
Product Team

📦 Product Team

🔍 Search for 'pr'

Create Space

egor

On this page:
- Quick search
- Full search
- Filter your search results
- Filter with CQL
- Advanced search
- Search options for the Document ation theme
- Search Confluenc e from your browser's search field
- Useful plugins

Related pages:

- Confluenc e Search Syntax
- Confluenc e Search Fields
- Recently Viewed Pages and Blog Posts
- Search Results Macro
- Livesearch Macro

- Search the People Directory

You can also search for administrative options in the quick search. For example, type 'general' into the search field to go to the General Configuration screen.

More information about quick navigation:

- Matching items are grouped by type so that you can quickly find the type you want. Confluence shows a maximum of 3 administrative items, 6 pages and/or blog posts, 2 attachments, 3 people and 2 spaces
- Items are ordered with the most recently updated first
- Permissions determine the admin options that appear in the search results. You'll only see the options

you have permission to perform.

Full search

The full search allows you to search all content in all spaces (site and personal), mail, personal profiles, attachments and space descriptions, or filter the results to show only what you're after.

The search will also look at the content of the following types of attachments:
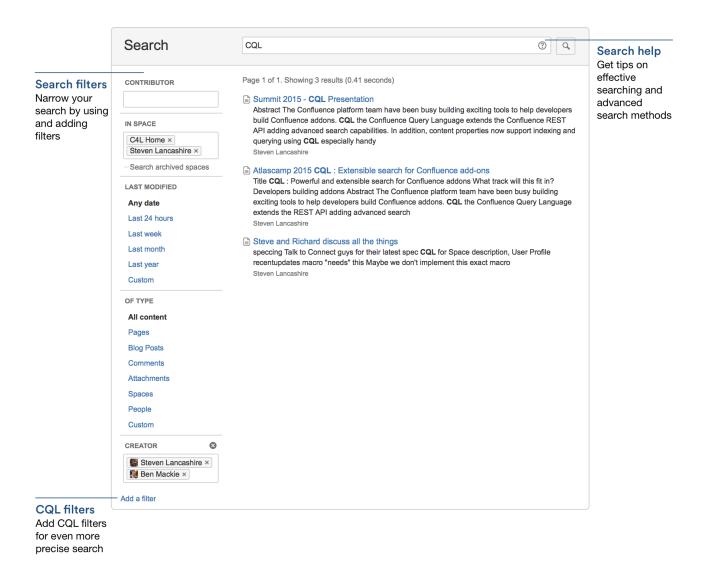
- Word
- Text
- PowerPoint
- Excel
- PDF
- HTML

To search the content of other attachment types, you will need to use an attachment content extractor plugin. For more information, take a look at the following:

- Existing extractor plugins which you can install on your Confluence site. See this search on Atlassian Marketplace.
- Guidelines on developing your own attachment content extractor plugin. See our developer documentation on extractor plugins.

To use the full search:

1. Type your query into the search field at the top-right corner of every screen (or at the top of the Search screen)
2. Hit the **Enter** key

Screenshot: Search results page

Search filters
Narrow your search by using and adding filters

CQL filters
Add CQL filters for even more precise search

Search | CQL | ⑦ | 🔍

Search help
Get tips on effective searching and advanced search methods

CONTRIBUTOR

IN SPACE

C4L Home ✕
Steven Lancashire ✕

‐ Search archived spaces

LAST MODIFIED

Any date
Last 24 hours
Last week
Last month
Last year
Custom

OF TYPE

All content
Pages
Blog Posts
Comments
Attachments
Spaces
People
Custom

CREATOR ⊗

👤 Steven Lancashire ✕
👤 Ben Mackie ✕

Add a filter

Page 1 of 1. Showing 3 results (0.41 seconds)

📄 Summit 2015 - **CQL** Presentation
Abstract The Confluence platform team have been busy building exciting tools to help developers build Confluence addons. **CQL** the Confluence Query Language extends the Confluence REST API adding advanced search capabilities. In addition, content properties now support indexing and querying using **CQL** especially handy
Steven Lancashire

📄 Atlascamp 2015 **CQL** : Extensible search for Confluence add-ons
Title **CQL** : Powerful and extensible search for Confluence addons What track will this fit in? Developers building addons Abstract The Confluence platform team have been busy building exciting tools to help developers build Confluence addons. **CQL** the Confluence Query Language extends the REST API adding advanced search
Steven Lancashire

📄 Steve and Richard discuss all the things
speccing Talk to Connect guys for their latest spec **CQL** for Space description, User Profile recentupdates macro "needs" this Maybe we don't implement this exact macro
Steven Lancashire

⚠️ We recommend you don't use special characters in page or attachment names, as the page or attachment may not be found by Confluence search, and may cause some Confluence functions to behave unexpectedly.

Filter your search results

On the left of the search results page are options that allow you to filter the search results.

You'll can filter using the following fields:

- **Contributor** – Restrict your search to content that's been modified (created, edited, or commented on) by a particular person. Start typing the person's username or part of their name and Confluence will offer you a list of possible matches.

  > You need to enter at least two letters. For example, if you enter just 'john s', the filter will look for users called 'john' and will ignore the 's'.

- **In space** – Only search in a particular space or list of spaces. There are also preset groups of spaces – favourite (my spaces), site, and personal spaces – which you can choose from the list of suggested spaces. Tick the checkbox to include archived spaces if you need to.
- **Last modified** – Choose or enter a date range to only show content updated within a particular period of time.
- **Of type** – Only return content of a certain type, like pages, blog posts, or comments.

Filter with CQL

Confluence search supports adding CQL filters to refine your results.

⌄ What's CQL, and how do I use it?

CQL (Confluence Query Language) is a query language developed for Confluence, which you can use in some macros and the Confluence search. Confluence search and CQL-powered macros allow you to add filters to build up a search query, adding as many filters as you need to narrow down the search results.

Use the **Add a filter** link to add more filters to your query.

⌄ Use AND, OR, and NOT operators

- For an OR search, specify multiple values in the same field.
So to show pages with 'label-a', 'label-b' or both you'd put 'label-a' and 'label-b' in the same Label field, like this:

Label *

| label-a × | label-b × |

- For an AND search, add more than one filter and specify a single value in each.
To show only pages with label-a and label-b you'd put 'label-a' in one label field, then add a second Label field to the macro, and put 'label-b' in the second one, like this:

Label *

| label-a × |

Label   ⊗

| label-b × |

Put simply, OR values are entered in the same filter, AND values are entered in different filter. Only some filters support AND. If the filter doesn't support the AND operator, you won't be able to add that filter more than once.
- For a NOT search, enter a minus sign (-) before the label. This'll exclude everything with that label.

You can use the following CQL filters to build your query:

Filter	Description	Operators
Label*	Include pages, blog posts or attachments with these labels.	OR (multiple values in the same filter) AND (multiple Label filters)
With ancestor	Include pages that are children of this page. This allows you to restrict the macro to a single page tree.	OR (multiple values in the same filter)
Contributor*	Include pages or blog posts that were created or edited by these people.	OR (multiple values in the same filter)
Creator	Include items created by these people.	OR (multiple values in the same filter)
Mentioning user	Include pages and blog posts that @mention these people.	OR (multiple values in the same filter)

With parent	Include only direct children of this page (further sub-pages won't be included)	EQUALS (one page only)
In space**	Include items from these spaces.	OR (multiple values in the same filter)
Including text**	Include items that contain this text.	CONTAINS (single word or phrase)
With title	Include items that contain this text in the title.	CONTAINS (single word or phrase)
Of type**	Include only pages, blogs or attachments.	OR (multiple values in the same filter)

* This field is required in CQL-powered macros.

** You can add these filters in CQL-powered macros but in search they're part of the standard search filters, so they don't appear in the **Add a filter** menu.
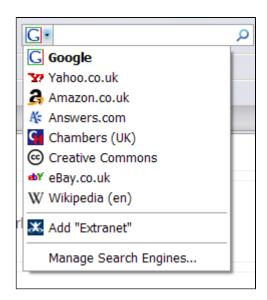
Advanced search

Still need more ways to search? For more advanced search syntax, check out Confluence Search Syntax for more ways to refine the text you enter in the search field. You can also read Confluence Search Fields for special parameters you can use to search various types of content metadata.

Search options for the Documentation theme

When using the Documentation theme, the Confluence search offers a few options as described below.

Using the search box at the top right of the page:

- By default, the main Confluence search is configured to search the entire Confluence site.
  - You will see the words '**Search Confluence**' in the search box at top right of the page.
  - The Confluence search will look for matches in the entire Confluence site. This is the default behaviour for other themes too.
- A space administrator can configure the Documentation theme to restrict the search to the current space.
  - You will see the words '**Search this space**' in the search box at top right of the page.
  - The search will return results from the current space only.
  - You can override the search restriction. Enter '**all:**' and your search term to search the entire site. For example, enter the following into the search box at top right of the page to search the entire site for 'technical writing':

```
all: technical writing
```

Using the search box in the left-hand panel:

- By default, the Documentation theme's left-hand panel includes a search box. Enter your search term there, to restrict the search to the current space. Specifically, this will search only the pages that are children of the space's home page.
- If your administrator has restricted the main search to the current space, there will not be a search box in the left-hand panel.

For screenshots and a full description, see The Documentation Theme.

Search Confluence from your browser's search field

If you're using Firefox or Internet Explorer, you can add your Confluence site as a search provider using the dropdown menu next to the browser's search field.

The example to the right shows the 'Extranet' Confluence site offered for inclusion as a search engine in the browser's search field.

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.
- Your Confluence Administrator can enable or disable the Open Search feature using the Confluence Administration Console.

Useful plugins

*Before installing an add-on (also called a plugin) into your Confluence site, please check the add-on's information page to see whether it is supported by Atlassian, by another vendor, or not at all. See our guidelines on add-on support.*This plugin extends the Confluence search, so that you can search for macro parameter names and values: the Confluence Macro Indexer plugin.

Confluence Search Syntax

This page describes the special words and punctuation marks you can use to refine your search.

Matched phrase search

Use double quotes to search for content that contains the phrase 'cheese one', or a phrase where 'cheese' and 'one' are the major words:

```
"cheese one"
```

Note: Confluence will ignore common words (stop words), including 'and', 'the', 'or', and more, even if they are included within double quotes. See the default list of stop words used by Confluence's search engine, Lucene, in the Lucene documentation.

For example:

1. Searching for "cheese one" returns only pages in which 'one' appears as the first word after 'cheese'.
2. Searching for "the one" returns all pages containing 'one' because 'the' is a stop word.

If you'd like to override Lucene's tokenisation and stemming, cast your vote on this improvement request:

> **CONF-14910** - Provide ability to override Lucene tokenisation and stemming and search for exact text **OPEN**

OR search

To search for content that contains one of the terms, 'chalk' or 'cheese', use the operator OR in capital letters :

```
chalk OR cheese
```

AND search

To search for content that contains both the terms 'chalk' and 'cheese', use the operator AND in capital letters:

```
chalk AND cheese
```

NOT search

To search for content that contains 'chalk' but NOT 'cheese', use the operator NOT in capital letters:

```
chalk NOT cheese
```

Excluded term search

To search for content that contains 'chalk' and 'butter' but not 'cheese':

```
chalk butter -cheese
```

Group search terms

To search for content that must contain 'chalk' but can contain either 'cheese' or 'butter', use brackets to group the search terms:

```
(cheese OR butter) AND chalk
```

Title search

To search for pages and posts with certain words in the title, use the **Add a filter** option at the bottom left of the page and choose **With title**. Enter the words you're looking for, in any order, and hit enter.

Confluence search filters are powered by CQL.

˅ What's CQL, and how do I use it?

CQL (Confluence Query Language) is a query language developed for Confluence, which you can use in some macros and the Confluence search. Confluence search and CQL-powered macros allow you to add filters to build up a search query, adding as many filters as you need to narrow down the search results.

Use the **Add a filter** link to add more filters to your query.

˅ Use AND, OR, and NOT operators

- For an OR search, specify multiple values in the same field.
  So to show pages with 'label-a', 'label-b' or both you'd put 'label-a' and 'label-b' in the same Label field, like this:

- For an AND search, add more than one filter and specify a single value in each.
  To show only pages with label-a and label-b you'd put 'label-a' in one label field, then add a second Label field to the macro, and put 'label-b' in the second one, like this:

Put simply, OR values are entered in the same filter, AND values are entered in different filter. Only some filters support AND. If the filter doesn't support the AND operator, you won't be able to add that filter more than once.
- For a NOT search, enter a minus sign (-) before the label. This'll exclude everything with that label.

You can use the following CQL filters to build your query:

Filter	Description	Operators
Label*	Include pages, blog posts or attachments with these labels.	OR (multiple values in the same filter)
		AND (multiple Label filters)
With ancestor	Include pages that are children of this page. This allows you to restrict the macro to a single page tree.	OR (multiple values in the same filter)
Contributor**	Include pages or blog posts that were created or edited by these people.	OR (multiple values in the same filter)
Creator	Include items created by these people.	OR (multiple values in the same filter)
Mentioning user	Include pages and blog posts that @mention these people.	OR (multiple values in the same filter)
With parent	Include only direct children of this page (further sub-pages won't be included)	EQUALS (one page only)
In space**	Include items from these spaces.	OR (multiple values in the same filter)
Including text**	Include items that contain this text.	CONTAINS (single word or phrase)
With title	Include items that contain this text in the title.	CONTAINS (single word or phrase)
Of type**	Include only pages, blogs or attachments.	OR (multiple values in the same filter)

* This field is required in CQL-powered macros.

** You can add these filters in CQL-powered macros but in search they're part of the standard search filters, so they don't appear in the **Add a filter** menu.

Date range search

To search for content modified within a certain date range, pick a timeframe in the **Last modified** section on the left. If you're looking for something created within a particular date range, use the **Add a filter** option at the bottom left of the page and choose **Created**. For either option, you can pick from some predefined options, like last 24 hours and last week.

Confluence search filters are powered by CQL.
⌄ What's CQL, and how do I use it?
   CQL (Confluence Query Language) is a query language developed for Confluence, which you can use in

some macros and the Confluence search. Confluence search and CQL-powered macros allow you to add filters to build up a search query, adding as many filters as you need to narrow down the search results.

Use the **Add a filter** link to add more filters to your query.

⌄ Use AND, OR, and NOT operators

- For an OR search, specify multiple values in the same field.
  So to show pages with 'label-a', 'label-b' or both you'd put 'label-a' and 'label-b' in the same Label field, like this:

- For an AND search, add more than one filter and specify a single value in each.
  To show only pages with label-a and label-b you'd put 'label-a' in one label field, then add a second Label field to the macro, and put 'label-b' in the second one, like this:

  Put simply, OR values are entered in the same filter, AND values are entered in different filter. Only some filters support AND. If the filter doesn't support the AND operator, you won't be able to add that filter more than once.
- For a NOT search, enter a minus sign (-) before the label. This'll exclude everything with that label.

You can use the following CQL filters to build your query:

Filter	Description	Operators
Label*	Include pages, blog posts or attachments with these labels.	OR (multiple values in the same filter) AND (multiple Label filters)
With ancestor	Include pages that are children of this page. This allows you to restrict the macro to a single page tree.	OR (multiple values in the same filter)
Contributor**	Include pages or blog posts that were created or edited by these people.	OR (multiple values in the same filter)
Creator	Include items created by these people.	OR (multiple values in the same filter)
Mentioning user	Include pages and blog posts that @mention these people.	OR (multiple values in the same filter)
With parent	Include only direct children of this page (further sub-pages won't be included)	EQUALS (one page only)
In space**	Include items from these spaces.	OR (multiple values in the same filter)

Including text**	Include items that contain this text.	CONTAINS (single word or phrase)
With title	Include items that contain this text in the title.	CONTAINS (single word or phrase)
Of type**	Include only pages, blogs or attachments.	OR (multiple values in the same filter)

* This field is required in CQL-powered macros.

** You can add these filters in CQL-powered macros but in search they're part of the standard search filters, so they don't appear in the **Add a filter** menu.

Wildcard searches

You can use one or more wildcard characters in your search and place them anywhere in the search string, **e xcept at the very beginning**. So, you could search for `http*.atlassian.*`, but you can't search for `*.a tlassian.*` or `?ttps://confluence.atlassian.*`, as they begin with a wildcard.

Wildcards can either replace a single character in your search, or multiple characters.

Single character

To replace a single character in your search, use a question mark (?) as a wildcard, For example, to search for 'butter', 'bitter', 'better', or 'batter'.

```
b?tter
```

Multiple characters

To replace multiple characters in your search, use an asterisk (*) as a wildcard. For example, to search for 'chicken' or 'chickpea':

```
chick*
```

Multiple wildcards

Use multiple wildcards in your search. The following query will search find 'chick', 'coconut', or 'chickpea':

```
c*c*
```

You can also combine wildcard characters in one search. For example, the search term below will return 'chick' but not 'chickpea':

```
c*c?
```

Note: Confluence doesn't support leading wildcards. This means searching for `*heese` will not return

cheese.

Proximity searches

Use a tilde character followed by a number, to find two words within a certain number of words of each other.

For example, the following search will return 'Octagon blog post', but not 'Octagon team blog post':

```
"octagon post"~1
```

The following search isn't valid, because you can't search for two words within zero words of each other. If you think the words are next to each other, use the matched phrase search.

```
"octagon post"~0
```

Range search

Use the operator 'TO', in capital letters, to search for names that fall alphabetically within a specified range:

```
[adam TO ben]
```

Note: You can't use the AND keyword inside this statement.

Fuzzy search

Use a tilde (~) character to find words spelled similarly.

If you want to search for octagon, but you're not sure how it's been spelt, type the word followed by a tilde:

```
octogan~
```

Combined search

You can also combine various search terms together:

```
o?tag* AND past~ AND ("blog" AND "post")
```

Searching for macros

You can search Confluence content to find where a macro is used. Start your search string with `macroName` ﹕ and type the macro name after the colon. For example, to search for all excerpt-include macros:

```
macroName:excerpt-include*
```

For more information about `macroName` and other search fields, see Confluence Search Fields.

Searching for labels

To search for pages and posts with a particular label or labels, use the **Add a filter** option at the bottom left of the page and choose **Label**.

Confluence search filters are powered by CQL.
  ⌄ What's CQL, and how do I use it?

CQL (Confluence Query Language) is a query language developed for Confluence, which you can use in some macros and the Confluence search. Confluence search and CQL-powered macros allow you to add filters to build up a search query, adding as many filters as you need to narrow down the search results.

Use the **Add a filter** link to add more filters to your query.

  ⌄ Use AND, OR, and NOT operators
  - For an OR search, specify multiple values in the same field.
    So to show pages with 'label-a', 'label-b' or both you'd put 'label-a' and 'label-b' in the same Label field, like this:

    Label *

    | label-a × | | label-b × |

  - For an AND search, add more than one filter and specify a single value in each.
    To show only pages with label-a and label-b you'd put 'label-a' in one label field, then add a second Label field to the macro, and put 'label-b' in the second one, like this:

    Label *

    | label-a × |

    Label   ⊗

    | label-b × |

    Put simply, OR values are entered in the same filter, AND values are entered in different filter. Only some filters support AND. If the filter doesn't support the AND operator, you won't be able to add that filter more than once.
  - For a NOT search, enter a minus sign (-) before the label. This'll exclude everything with that label.

You can use the following CQL filters to build your query:

Filter	Description	Operators
Label*	Include pages, blog posts or attachments with these labels.	OR (multiple values in the same filter) AND (multiple Label filters)
With ancestor	Include pages that are children of this page. This allows you to restrict the macro to a single page tree.	OR (multiple values in the same filter)
Contributor**	Include pages or blog posts that were created or edited by these people.	OR (multiple values in the same filter)
Creator	Include items created by these people.	OR (multiple values in the same filter)
Mentioning user	Include pages and blog posts that @mention these people.	OR (multiple values in the same filter)

With parent	Include only direct children of this page (further sub-pages won't be included)	EQUALS (one page only)
In space**	Include items from these spaces.	OR (multiple values in the same filter)
Including text**	Include items that contain this text.	CONTAINS (single word or phrase)
With title	Include items that contain this text in the title.	CONTAINS (single word or phrase)
Of type**	Include only pages, blogs or attachments.	OR (multiple values in the same filter)

* This field is required in CQL-powered macros.

** You can add these filters in CQL-powered macros but in search they're part of the standard search filters, so they don't appear in the **Add a filter** menu.

Confluence Search Fields

This page gives an overview of the Apache Lucene search fields used in Confluence.

Filter with CQL

Before you dive into learning more about Lucene fields, you may want to learn about the powerful search filtering offered by Confluence Query Language (CQL).

⌄ What's CQL, and how do I use it?

CQL (Confluence Query Language) is a query language developed for Confluence, which you can use in some macros and the Confluence search. Confluence search and CQL-powered macros allow you to add filters to build up a search query, adding as many filters as you need to narrow down the search results.

Use the **Add a filter** link to add more filters to your query.

⌄ Use AND, OR, and NOT operators
  * For an OR search, specify multiple values in the same field. So to show pages with 'label-a', 'label-b' or both you'd put 'label-a' and 'label-b' in the same Label field, like this:

    Label *

    label-a ×   label-b ×

  * For an AND search, add more than one filter and specify a single value in each.
    To show only pages with label-a and label-b you'd put 'label-a' in one label field, then add a second Label field to the macro, and put 'label-b' in the second one, like this:

Put simply, OR values are entered in the same filter, AND values are entered in different filter.
Only some filters support AND. If the filter doesn't support the AND operator, you won't be able to add that filter more than once.

- For a NOT search, enter a minus sign (-) before the label. This'll exclude everything with that label.

You can use the following CQL filters to build your query:

Filter	Description	Operators
Label*	Include pages, blog posts or attachments with these labels.	OR (multiple values in the same filter) AND (multiple Label filters)
With ancestor	Include pages that are children of this page. This allows you to restrict the macro to a single page tree.	OR (multiple values in the same filter)
Contributor**	Include pages or blog posts that were created or edited by these people.	OR (multiple values in the same filter)
Creator	Include items created by these people.	OR (multiple values in the same filter)
Mentioning user	Include pages and blog posts that @mention these people.	OR (multiple values in the same filter)
With parent	Include only direct children of this page (further sub-pages won't be included)	EQUALS (one page only)
In space**	Include items from these spaces.	OR (multiple values in the same filter)
Including text**	Include items that contain this text.	CONTAINS (single word or phrase)

With title	Include items that contain this text in the title.	CONTAINS (single word or phrase)
Of type**	Include only pages, blogs or attachments.	OR (multiple values in the same filter)

* This field is required in CQL-powered macros.

** You can add these filters in CQL-powered macros but in search they're part of the standard search filters, so they don't appear in the **Add a filter** menu.

Searching for content in specific fields

Confluence data is stored in fields which can be specified in the search. To search a specific field, type the name of the field followed by a colon ':' and then the term you are looking for.

Preliminary examples:

```
title:"Some Title"
```

```
labelText:chalk
```

The field specification applies only to the term directly preceding the colon. For example, the query below will look for "Some" in the title field and will search for "Heading" in the default fields.

```
title:Some Heading
```

Confluence search fields

Below are the fields which can be searched, listed by content type.

Personal Information

Name	Indexed	Stored	Tokenised	Notes
handle	true	true	false	
type	true	true	false	
urlPath	true	true	false	
fullName	true	true	true	
title	true	true	false	
labelText	true	true	true	
modified	true	true	false	
created	true	true	false	

contentBody	true	true	true

Pages

Name	Indexed	Stored	Tokenised	Notes
handle	true	true	false	
type	true	true	false	
urlPath	true	true	false	
title	true	true	true	
spacekey	true	true	false	
labelText	true	true	true	
modified	true	true	false	
created	true	true	false	
contentBody	true	true	true	
macroName	true	true	false	The name of a macro used on the page

Blog

Name	Indexed	Stored	Tokenised	Notes
handle	true	true	false	
type	true	true	false	
urlPath	true	true	false	
title	true	true	true	
spacekey	true	true	false	
labelText	true	true	true	
modified	true	true	false	
created	true	true	false	
contentBody	true	true	true	
macroName	true	true	false	The name of a macro used in the blog

Attachments

Name	Indexed	Stored	Tokenised	Notes
handle	true	true	false	
type	true	true	false	
urlPath	true	true	false	
filename	true	true	true	
title	true	true	false	
comment	true	true	true	

spacekey	true	true	false	
modified	true	true	false	
created	true	true	false	
contentBody	true	true	true	

Mail items

Name	Indexed	Stored	Tokenised	Notes
handle	true	true	false	
type	true	true	false	
urlPath	true	true	false	
title	true	true	true	
spacekey	true	true	false	
messageid	true	true	false	
inreplyto	true	true	false	
recipients	true	true	true	
labelText	true	true	true	
modified	true	true	false	
created	true	true	false	
contentBody	true	true	true	

Useful plugins

Before installing an add-on (also called a plugin) into your Confluence site, please check the add-on's information page to see whether it is supported by Atlassian, by another vendor, or not at all. See our guidelines on add-on support.

This plugin extends the Confluence search, so that you can search for macro parameter names and values: the Confluence Macro Indexer plugin.

Notes

To find out the version of Lucene Confluence is using go to `<installation directory>/confluence/WEB-INF/lib` and locate the Lucene jar files. The Lucene version number will be part of the filename.

Search the People Directory

The people directory displays a list of people who are authorised to log in to your Confluence site.

The people directory includes anybody who has logged into Confluence or who has had a user account created for them in Confluence.

The people directory does not include users who can log into Confluence using external user management if they have never yet logged in.

View the people directory

Choose **People** at the top of the screen.

Search for people

To search for a particular person, type their first name and/or last name into the search box and choose **Search**.

- To see everyone who uses your Confluence site, choose **All People**.
- To see just those people who have set up a personal space, choose **People with Personal Spaces**.

Follow people's activities

Confluence's network features allow you to 'follow' (that is, keep track of) other people's activities in your Confluence site. For more information, please refer to Network Overview. You can use the hover profile feature in the people directory to start following other people.

- To start following someone, move your mouse over their name or profile picture and choose **Follow** in their profile popup.
- To stop following someone, move your mouse over their name or profile picture and choose **Stop Following** in their profile popup.

Once you start following another person, their activities will start appearing in your network view.

Screenshot: The people directory

People Directory

All People	All People	Search Clear
People with Personal Spaces		

All People

- Josh User
  sample@email.com.au
- Ewan User
  sample@email.com.au
- Sophie Staunton
  sample@email.com.au
- Rach Admin
  sample@email.com.au

Notes

- The people directory uses the hCard microformat for simple integration with a variety of microformat-enabled tools. hCard is an open data format for representing people, companies, organisations, and places. Read more about microformats and hCard.
- By default, deactivated users (disabled user accounts) are excluded from the people directory. You can include them by adding the `showDeactivatedUsers` parameter to the URL. For example:

```
http://my.confluence.com/dopeopledirectorysearch.action?showDeac
tivatedUsers=true
```

- By default, externally deleted users (for example, users deleted from an LDAP repository) are excluded from the people directory. You can include them by adding the `showExternallyDeletedU sers` parameter to the URL. For example:

```
http://my.confluence.com/dopeopledirectorysearch.action?showExte
rnallyDeletedUsers=true
```

- The Confluence administrator can hide the people directory. If it is hidden, you will not see the **People Directory** option.

Recently Viewed Pages and Blog Posts

The **Recently Viewed** list in Confluence keeps track of pages and blog posts you've recently visited, and allows you to easily navigate back to them.

Related pages:
- Create and Edit Pages
- Your Profile and Settings

To view your recently viewed content:

1. Choose **your profile picture** at top right of the screen, then choose **Recently Viewed**
2. Choose the title of the page you want to revisit

To filter the list, type part of a page title or user's name in the **Filter** field.

Your last ten recently viewed pages also appear when you click in Confluence's **Search** field before you start typing a search query.

Recently Viewed Pages

Filter

TODAY

- Status update - Josh User
- Josh User's Home - Josh User
- Angry Nerds Home - Angry Nerds
- Sophie User's Home - Sophie User
- Blogs are a great way to communicate - Sophie User
- Test page 1 - Sophie User

OLDER

- Test Space Home - Sample Space
- Step 3 - share your favourite nerd - Angry Nerds
- Step 2 - Play - Angry Nerds
- Step 1 - download the game - Angry Nerds

Hint: type "g" and then "r" anywhere to quickly open this menu   Close

Permissions and Restrictions

As a tool for communication and collaboration, we believe Confluence is at its best when everyone can participate fully. Confluence keeps a history of all changes to pages and other content, so it's easy to see who has changed what, and reverse any changes if you need to.

Confluence does, however, give you the choice to make your site, spaces, and pages as open or closed as you want to.

Levels of permission

There are three levels of permissions in Confluence: Global permissions, space permissions, and page restrictions.

Global permissions

Global permissions are site-wide permissions, and can be assigned by either a System or Confluence administrator.

These permissions are pretty broad, and don't really interact with space permissions or page restrictions.

On this page:
- Levels of permission
  - Global permissions
  - Space permissions
  - Page restrictions
- How do permissions and restrictions interact?

Related pages:

- Confluence Groups
- Global Permissions Overview
- Space Permissions Overview
- Page Restrictions
- Configuring Confluence Security
- Confluence Security Overview and Advisories

For full details, check out the overview of global permissions in the *Administrator's Guide*.

Space permissions

Every space has its own independent set of permissions, managed by the space admin(s), which determine the access settings for different users and groups.

They can be used to grant or revoke permission to view, add, edit, and delete content within that space, and can be applied to groups, users, and even to anonymous users (users who aren't logged in) if need be.

> One thing to watch out for is where a user is a member of multiple groups. You may have revoked permission for that individual user to add pages, for example, but if they're a member of a groups that *is* allowed to add pages, they'll still be able to create new pages in the space.
>
> If you can't get the result you want from space permissions, or you're not sure, check with one of your Confluence administrators to determine what permissions you should apply to individuals and groups.

Page restrictions

Page restrictions work a little differently to global and space permissions. Pages are open to viewing or editing by default, but you can restrict either viewing or editing to certain users or groups if you need to.

Don't forget, every page in Confluence lives within a space, and space permissions allow the space admin to revoke permission to view content for the whole space. Even the ability to apply restrictions to pages is controlled by the 'restrict pages' space permission.

How do permissions and restrictions interact?

You can restrict viewing of a page or blog post to certain users or groups, so that even if someone has the 'view' permission for the space, they won't be able to view the content of the page or blog post.

If someone's a space admin and you've used page restrictions to prevent them viewing a page, they won't be able to see the page when they navigate to it. As a space admin though, they can see a list of restricted pages in the space and remove the restrictions.

What about links?

Space permissions and page restrictions affect how links between Confluence pages are displayed.

- If someone doesn't have 'View' space permission, links to pages in that space won't be shown at all.
- If someone has the "View" space permission, but the page has view restrictions, the link will be visible but they'll get an "access denied" message when they click the link.

Links to attachments are also affected. If the visitor doesn't have permission to view the page the attachment lives on, the link won't be rendered.

Confluence Groups

For Confluence administrators, grouping users in Confluence is a great way to cut down the work required when managing permissions and restrictions. Groups are also very useful, however, to anyone who's a space admin, or can apply page restrictions.

If you're a space admin, you can assign a set of space permissions to a group rather than to each individual user. And as a page creator with 'Add/Delete Restrictions' permission, you can also add and remove page restrictions for groups.

Default Confluence groups

There are some default groups in every Confluence instance but, beyond that, Confluence administrators are free to set up and edit groups in any way they see fit.

The two special groups in Confluence are:

- **confluence-administrators** – Can perform most of the Confluence administrative functions, like assign permissions to other users, but they can't perform any functions that could compromise the security of the Confluence system. They can also access the Confluence Admin console.
- **confluence-users** - this is the default group into which all new users are assigned. Permissions defined for this group will be assigned to all new Confluence users.

Anonymous Users

All users who don't log in when they access Confluence are know as 'anonymous' users. By default, anonymous users don't have access to view or change any content in your Confluence instance, but Confluence admins can assign permissions to this group if it's required.

> **Overlapping group and user permissions**
> When a user is assigned more than one permission, the more powerful permission will prevail.
>
> Further explanation:
>
> - A user may be assigned a permission specifically to their username. They may also be assigned a permission by belonging to a group, or even several groups.
> - The user will then be able to perform all functions assigned to them.
> - So if a user is allowed to do something over and above what the group can do, the user will be able to do it. And if the group is allowed to do something over and above the specific permissions granted to the user, the user will still be able to do it.
> - If anonymous users are allowed to do something over and above what the user or group can do, the user will be able to do it, (even while logged in).

Unlicensed users from linked applications

If you're using Confluence as a knowledge base for JIRA Service Desk, your JIRA Service Desk administrator can choose to allow all active users and customers (that is logged in users who do not have a Confluence license) to view specific spaces.

These users have very limited access, and cannot be granted permissions in the same was as an individual or group. However, it's important to note that this permission overrides all existing space permissions, so any logged in Confluence user will also be able to see the space (regardless of their group membership). This is due to the way Confluence inherits permissions.

Add-ons and Integrations

Confluence has a wide range of features on its own, but you can also extend those features with add-ons, and by integrating Confluence with other applications. Integrating with JIRA applications and HipChat can really take your Confluence experience to the next level by improving the way your teams communicate and collaborate, track vital work, and plan and release new products.

If there's an extra piece of functionality you need, the Atlassian Marketplace is the place to look for useful Confluence add-ons (or add-ons for other Atlassian products). Whether you need to create diagrams, like the ones you can create with Gliffy, or you want to make awesome mockups and wireframes with Balsamiq, there are heaps of great add-ons in the marketplace. You may even find a really useful add-on you never knew you needed, but now can't live without.

In this section:

- Use JIRA applications and Confluence together
- Use HipChat and Confluence together
- Request Add-ons
- Use a WebDAV Client to Work with Pages
- Mail Archives
- Gadgets

Related pages:

- Use JIRA applications and Confluence together
- Use HipChat and Confluence together
- Use a WebDAV Client to Work with Pages
- Gadgets
- Mail Archives
- Request Add-ons

Use JIRA applications and Confluence together

Confluence and JIRA are like bacon and eggs; coffee and cake; Simon and Garfunkel. Separately, they're great, but together, they're amazing!

If your Confluence and JIRA sites are connected using Application Links, you can display and create JIRA issues and more from within Confluence.

What you can do with Confluence and JIRA depends on the JIRA application and version you have. Find out about the required applications and versions later in this page.

For every project or team

Display issues on a page

You can display JIRA issues on a Confluence page using the JIRA Issues macro. Display a single issue, a list of issues, or show the total number of issues.

The simplest way to add a JIRA issue to Confluence is to paste a JIRA URL on a Confluence page.
˅ Here's some examples...

On this page:

- For every project or team
  - Display issues on a page
  - Create reports and charts
  - Create issues from inside Confluence
  - Move between JIRA and Confluence
- For software teams
  - Define your requirements
  - Manage your sprints
- For service desk teams
  - Provide self help resources for your customers
  - Create knowledge base articles
  - Allow any active user to see knowledge base spaces
- JIRA applications required

- `<yourjirasite.com>/browse/CONF-1234` will insert the JIRA Issues macro and display a single issue.
- `<yourjirasite.com>/issues/?filter=56789` will insert the JIRA Issues macro and display a list of issues matching the saved filter.
- `<yourjirasite.com>/issues/?jql=project%20%3D%20CONF` will insert the JIRA Issues macro and display a list of issues matching the JIRA search.

Alternatively, you can add the JIRA Issues Macro to the page and search for issues directly:

T	Key	Summary	Status
◘	CONF-26791	PDF Export Table Hidden by Margin	RESOLVED
⚡	CONF-35848	Roadmap time period only showing months	RESOLVED
➕	CONF-25846	Automated row numbering in tables	RESOLVED

3 issues ↻ Refresh

1. In the editor choose **Insert** > **JIRA Issue**.
2. Follow the prompts in the macro browser to choose a project and search for an issue – you can even use JIRA Query Language (JQL).

Once you've added the macro, you can customise how the issue or list of issues appears on the page, including how much information to display, how many issues, and more.

Create reports and charts

Reporting on information stored in JIRA is simple in Confluence. In addition to the JIRA Issues Macro, you can use the JIRA Report blueprint or JIRA Chart macro to show information from your JIRA application visually. It's the best way to give your stakeholders a snapshot of your team or project's progress.

You can:

- Use the JIRA Report blueprint to create a Change Log or Status report.
- Use the JIRA Chart Macro to display data as a chart, including pie charts, created vs resolved, and two dimensional charts.
- Use JIRA Gadgets to display detailed JIRA reports and charts on pages.

Create issues from inside Confluence

You can create issues while viewing a page or from the within the editor. This is really useful if you use Confluence for planning and gathering requirements.

To create an issue when viewing a page:

1. Highlight some text on your page and choose the
   ⚙
   Create JIRA issue icon that appears above the highlighted text.
2. Enter your **server** (if you have multiple JIRA sites connected to Confluence), **project**, **issue type** and **description**. Your highlighted text will populate the issue **summary** automatically.
3. Choose **Create**.

The issue will be created in JIRA and added to your page. If your text is in a table, you'll have the option to create multiple issues using text from the same column.

If you don't see a popup when you highlight text, check that Text Select is enabled in your profile settings.

To create an issue in the editor:

1. In the editor choose **Insert** > **JIRA Issue** > **Create new issue**.
2. Enter your **server** (if you have multiple JIRA sites connected to Confluence), **project**, **issue type**, **summary**, and **description**.
3. Choose **Insert**.

The issue will be created in JIRA and added to your page.

There are some limitations when creating JIRA issues from Confluence. The JIRA Issues macro or Create JIRA Issue dialog will notify you if it's unable to create an issue in the selected project. You can find out more in the JIRA Issues Macro page.
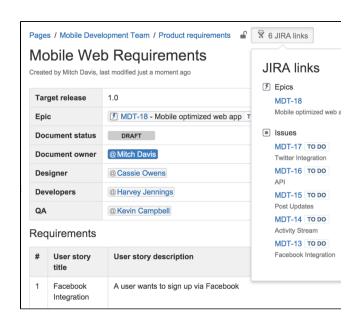
Move between JIRA and Confluence

Whenever you add a link to JIRA issues in Confluence, or link to a Confluence page from your JIRA application, the **JIRA Links** button appears at the top of the Confluence page. This makes it really easy to jump from Confluence to JIRA and vice versa, speeding up your workflow.

The number on the JIRA Links button indicates the total number of issues, epics, and sprints connected to that page, regardless of whether you have permission to view them. The dropdown, however, will only show details of issues, epics, and sprints that you have JIRA permissions to view.

⌄ There are a few exceptions...
- The JIRA Links button only appears in the default theme; it's not available in the Documentation theme.
- The button doesn't detect links from issues displayed in the JIRA Issues macro in table format.

For software teams

Here's some suggestions to help you get the most out of Confluence and JIRA Software and unleash the potential in your agile development team.

Define your requirements

Confluence is the perfect place to start defining your requirements. You can use the Product Requirements Blueprint to capture your requirements, then create your JIRA epic and other issues right from the requirements page in Confluence.

Here's how it works:

1. Create a Confluence page using the Product Requirements Blueprint.
2. Choose the placeholder text '*Link to JIRA epic or feature*' and choose **Create new issue** to create your epic in JIRA.
3. Collaborate with your team to define your stories and save the page.
4. Highlight text on your requirements page and choose the **Create JIRA issue**
   ✗
   link to create stories in JIRA, and automatically link them to your epic.
5. Track the progress of the stories from the Confluence page or from within JIRA.

The tight integration between Confluence and JIRA Software means you can easily access issues from the Confluence page and see their status at a glance, and from within JIRA Software you can see links to related Confluence pages. All the information you need is right there.

Manage your sprints

There's often a lot of material in Confluence that provides useful context for your team during a sprint. These might be requirements documents, designs, tech specs, customer research and more. By linking these pages to epics, you make them easy for your team to find during the sprint.

Here's how you can use Confluence to support your sprint from within JIRA Agile:

- In JIRA Software, create a Confluence page to plan your sprint. The page is created using the Meeting Notes Blueprint – a handy template that helps capture the details you need – and is automatically linked to the sprint.
- In an epic, link to useful Confluence pages, including requirements, designs, and more.
- Report on your progress to stakeholders using the JIRA Reports blueprint in Confluence.
- Use the Retrospective Blueprint in Confluence at the end of your sprint to take stock of what went well and not so well.

For people who work mostly in JIRA Software, the integration means that useful Confluence pages are only a click away.

For service desk teams

Provide self help resources for your customers

If you use JIRA Service Desk, you can help your customers resolve their issues without creating a request by connecting your Service Desk project to a knowledge base in Confluence.

In Service Desk, head to **Administration** > **Confluence KB** (or **Settings** > **Knowledge Base** if you're using JIRA Service Desk 2.x or earlier) to connect or create a Confluence space.

When Service Desk customers search in the Customer Portal, pages in the linked knowledge base space will be suggested, allowing customers to help themselves.

Create knowledge base articles

The Knowledge Base space blueprint, along with templates for how-to and troubleshooting articles make creating new knowledge base articles super simple for your Service Desk agents.

The templates used in the how-to and troubleshooting blueprints are completely customisable too. Set up the template with all your standard information and let your agents take it from there.

Allow any active user to see knowledge base spaces

If your Confluence instance is not public, you can still make a knowledge base space available via the customer portal.

When you link your JIRA Service Desk project to a Confluence space, you can choose to allow all active users and customers to see pages in the linked space, even if they don't have a Confluence license. These people get very limited Confluence access.

Unlicensed users can:

- View pages via the JIRA Service Desk customer portal.
- Follow a URL to a page and then navigate within the linked space.

Unlicensed users can't:

- Like, comment on or edit pages (or be granted permission to do this).

- See the dashboard, user profiles, the people directory or space directory.
- Search the whole site.

This permission can only be enabled via JIRA Service Desk, but you can revoke access to the whole site or to particular spaces via Confluence's global permissions or space permissions.

ˇ A note about Confluence permissions...

Allowing all active users and customers to view a space will override all existing space permissions, so any logged in, licensed Confluence user will also be able to see the space (regardless of their group membership). This is due to the way Confluence inherits permissions.

JIRA applications required

As you've seen, Confluence has many integration points, some of which are only available in particular JIRA applications or versions.

This matrix outlines the specific JIRA applications you'll need for each feature. We've also included the minimum legacy JIRA Server version (plus any add-ons) that you'll need if you're not using the latest JIRA applications.

Feature	JIRA Core	JIRA Software	JIRA Service Desk	Minimum legacy version
Display issues using JIRA issues macro				JIRA 4.3
Display issue and project information using JIRA chart macro				JIRA 5.0
Display issue and project information using the JIRA Report blueprint				JIRA 5.0
Create an issue from JIRA issues macro				JIRA 4.3
Create issues by highlighting text on a Confluence page				JIRA 6.3.1
Create issue by highlighting text on a Confluence page and automatically link issues to an epic				JIRA 6.3.1 and JIRA Agile 6.3.5
Link and create Confluence pages from epics and sprints				JIRA 6.3.1 and JIRA Agile 6.3.5
View linked issues with the JIRA links button in Confluence				JIRA 6.3.1
Use a Confluence space as a knowledge base in the customer portal				JIRA 5.2 and JIRA Service Desk 1.0
Allow Service Desk customers to view knowledge base articles without a Confluence license				JIRA Service Desk Cloud only
Search for an existing Confluence page within the JIRA link dialog				JIRA 4.3
Delegate user management to JIRA				JIRA 4.3
Add a JIRA gadget to a Confluence page				JIRA 4.3

That's it? Time to jump into Confluence and give some of these great features a try with your team or project.Want to find out more about how to connect your JIRA application to Confluence? Check out Integrating JIRA and Confluence.

Use HipChat and Confluence together

HipChat is group and video chat for teams.

Connect Confluence to HipChat to send real time notifications to your team or project's HipChat room

Don't have HipChat yet? You can sign up for free a nd start using HipChat right away.

> **On this page:**
>
> - Connect Confluence and HipChat
> - Invite users
> - Send space notifications to HipChat
> - See who's online

Connect Confluence and HipChat

First your organisation's HipChat account needs to be connected to Confluence. It only takes a minute.

You'll need administrator permissions for your HipChat group to do this.

If you're a Confluence admin go to

⚙ ▾

 > **General Configuration** > **HipChat Integration** and click **Connect HipChat** to get started.

If you're a Space Admin you can go to **Space Tools** > **Integration** > **HipChat** to get started.
 ⌄ Using the Documentation theme?
   If your space uses the Documentation theme go to **Browse** > **Space Admin** > **HipChat**.

You'll need to be logged in to HipChat as a Group Admin to complete the integration.

Invite users

If you're a Confluence administrator you can send an email to invite users to join HipChat directly from the Integration screen (

⚙ ▾

 > **General Configuration** > **HipChat Integration)**. You'll need at least one space integrated with a room to see the invite users link.

Send space notifications to HipChat

Keep your team in the know by sending notifications about space activities, like new pages and blogs, to your team's HipChat room. Notifications appear in real-time, and one click takes you straight to Confluence.

To set up space notifications go to **Space Tools** > **Integration** > **HipChat** and add a room to the list.
 ⌄ Using the Documentation theme?
   If your space uses the Documentation theme go to **Browse** > **Space Admin** > **HipChat**.

You'll need Space Admin permissions, and if you are connecting to private HipChat rooms, you will need to log in to HipChat on the integration screen.

See who's online

When you hover over a user mention or a byline it will show if the user is available in HipChat. Green, yellow and red icons indicate when someone is available, away or doesn't want to be disturbed.

Request Add-ons

The Atlassian Marketplace website offers hundreds of add-ons that the administrator of your Atlassian application can install to enhance and extend Confluence. If the add-on request feature is enabled for your Confluence instance, you can submit requests for add-ons from the Marketplace to your Confluence administrator.

The 'Atlassian Marketplace for Confluence' page provides an integrated view of the Atlassian Marketplace from within your Confluence instance. The page offers the same features as the Marketplace website, such as searching and category filtering, but tailors the browsing experience to Confluence.

This in-product view of the Marketplace gives day-to-day users of the Atlassian applications, not just administrators, an easy way to discover the add-ons that can help them work. When you find an add-on of interest, you can submit a request with just a few clicks.

Submit an add-on request

To browse for add-ons in the Atlassian Marketplace, follow these steps:

1. Choose **your profile picture** at top right of the screen, then choose **Atlassian Marketplace**.
2. In the Atlassian Marketplace page, use the search box to find add-ons or use the category menus to browse or filter by add-ons by type, popularity, price or other criteria. You can see what your fellow users have requested by choosing the **Most Requested** filter.
3. When you find an add-on that interests you, click **Request** to generate a request for your administrator.
4. Optionally, type a personal message to your administrators in the text box. This message is visible to administrators in the details view for the add-on.

5. When ready, click **Submit Request**.
6. Click **Close** to dismiss the 'Success!' message dialog box.

At this point, a notification appears in the interface your administrators use to administer add-ons. Also your request message will appear in the add-on details view, visible from the administrator's 'Find New Add-ons' page. From there, your administrator can purchase the add-on, try it out or dismiss requests.

Update an add-on request

After submitting the request, you can update your message at any time. Click the **Update Request** button next to the listing in the Atlassian Marketplace page to modify the message to your administrator.

The administrator is not notified of the update. However, your updated message will appear as you have modified it in the details view for the add-on immediately.

Use a WebDAV Client to Work with Pages

Create, move and delete pages and attachments in Confluence using a file manager like Finder (OS X), Explorer (Windows) or Dolphin (Linux) or other WebDav compatible local client like CyberDuck.

For example, if you need to delete a lot of pages you can bulk delete them in your local file manager (like Finder or Explorer), rather than one by one in your browser.

Access to Confluence through a native client is provided by the WebDav plugin. Your administrator may have disabled the WebDav plugin, or may have restricted the actions that you can perform using a local client. See Co nfiguring a WebDAV client for Confluence for more information on how to set it up.

Manage pages and files in a native client

Accessing Confluence through a native client is useful for performing bulk actions. Before you can start creating and moving things around, it's useful to understand how the content is organised.

The hierarchy in the file system looks like this:

- Type of space (global or personal)
  - Space (folder name is the spacekey)
    - Homepage (and other top level pages)
      - Child pages (folder name is the name of the page)
        - Attachments (filename of the attachment)

Essentially the file structure is the same as the page tree in your space. Here's how the Confluence demonstration space looks in Finder.

Here's some things you might choose to do in a local client, rather than in your browser:

- **Move pages to another space**
  Select the page folders, and drag them into the other space's folder (drag them from Space A to Space B)
- **Delete multiple pages**
  Select all the page folders you want to delete and delete them.
- **Delete multiple attachments from a page**
  Navigate down to the page folder, select the attachments you want to delete and delete them.
- **Upload multiple attachments**
  Navigate to the page folder, and drag the files into the folder (note you can attach multiple files through the insert dialog as well).

Mail Archives

Confluence allows you to collect and archive mail within each space. It's useful for storing the email messages that relate to a particular project – you can put them in the same Confluence space as the content for that project.

You can download mail from one or more POP or IMAP accounts, or import mail from an mbox file on your local system or on the Confluence server.

You need space administration permissions to manage the mail archives.

Related pages:

- Spaces
- Add a Mail Account

Confluence mail archiving is an optional feature. This means that the 'Mail' options may be disabled and will therefore not appear in the Confluence user interface. Mail archiving features are contained in a bundled plugin. To activate mail archiving features in Confluence, enable the plugin – choose

 > **General Configuration** > **Manage Add-ons**. Then choose **System** in the drop down, and enable the **Confluence Mail Archiving Plugin**.

To see archived mail:

- Go to a space and choose **Space tools** > **Integrations** > **Mail**
  Or, if your space uses the Documentation theme, choose **Browse** > **Mail** in the header.
- Choose a message to see its contents, or choose **Next**, **Previous** and other options to navigate around the mail archives.

Manage mail archives:

- Add a Mail Account
- Delete and Restore Mail
- Import Mail from an mbox

Screenshot: Viewing a message in the mail archive

Notes

- Only **site spaces** – not personal spaces – can store mail archives. See Spaces for information on site and personal spaces.
- You can also search the mail messages and their attachments. See Search.
- Confluence mail archiving is an optional feature. This means that the 'Mail' options may be disabled and will therefore not appear in the Confluence user interface. Mail archiving features are contained in a bundled plugin. To activate mail archiving features in Confluence, enable the plugin – choose

  ⚙ ▾

  > **General Configuration** > **Manage Add-ons**. Then choose **System** in the drop down, and enable the **Confluence Mail Archiving Plugin**.

Add a Mail Account

When you add a mail account, you're configuring Confluence to download mail from that account and archive it within the space.

You need space administration permissions to add a mail account. See Spa ce Permissions Overview.

⚠ **Note:** Confluence will remove email messages from an email account when it transfers them to the mail archive. You must therefore configure Confluence to poll a **clone** email account rather than the actual account. For example, to archive the actual account `sales@company.com` to your Confluence Sales space, you must first create a clone account such as `con f-sales@company.com` that contains the same email content.

On this page:

- Add a mail account
- Fetching Mail
- Notes

Related pages:

- Mail Archives
- How do I check which spaces have email accounts?
- How do I disable automatic mail polling?

Add a mail account

Step 1. Create a clone email account on the mail server

1. Add a new email account on the mail server with the clone email address.
2. Copy all existing email messages from the actual account to the clone account.
3. Set up the actual account to bcc sent email messages to the clone account.
4. Set up the actual account to forward received email messages to the clone account.

Step 2. Configure Confluence to archive the clone account

1. Go to the 'Mail Accounts' view:
   - Go to the space and choose **Space tools** > **Integrations** from the bottom of the sidebar. Choose **Mail Accounts**.
   - If your space uses the Documentation theme: Choose **Browse** > **Space Admin** from the header

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

     Choose **Integration** > **Mail Accounts**.
2. Choose **Add mail account**
3. Enter configuration details for the account:
   - **Account Name:** Enter a name for this account by which it will be known in Confluence.
   - **Description:** Provide a description for this account (optional).
   - **Protocol:** Choose from POP, IMAP, POPS or IMAPS.
   - **Hostname:** Enter the host name of the mail server on which the account resides.
   - **Port:** Don't edit this field. The mail server's port number will be displayed by default.
   - **Username:** Enter a username that has permission to retrieve mail from this account.
   - **Password:** Enter the account's password.
4. Choose **Test Connection** to verify the details
5. Choose **Create** to add the account to Confluence

For each mail account you add, you can perform these actions in the **Mail Accounts** tab:

- **Edit:** Change the configuration settings for the mail account.
- **Remove:** Remove the account permanently.
- **Disable/Enable:** Temporarily disable the account, or enable a disabled account.

Fetching Mail

Confluence automatically fetches mail from the server once every 30 minutes. You can manually retrieve new mail from the configured mail accounts by selecting the **Mail** tab and choosing **Fetch new mail**.

ℹ️ You need to be a space administrator to manually retrieve mail. See Space Permissions.

Notes

- Only **site spaces** – not personal spaces – can store mail archives. See Spaces for information on site and personal spaces.
- Confluence mail archiving is an optional feature. This means that the 'Mail' options may be disabled and will therefore not appear in the Confluence user interface. Mail archiving features are contained in a bundled plugin. To activate mail archiving features in Confluence, enable the plugin – choose ⚙️▾
  > **General Configuration** > **Manage Add-ons**. Then choose **System** in the drop down, and enable the **Confluence Mail Archiving Plugin**.
- Once mail is fetched it will be removed from the server.

Delete and Restore Mail

To delete mail from a space, you need 'Delete Mail' permission.

ℹ Only a space administrator can delete all email messages in the space simultaneously.

Delete mail from a space:

1. Go to a space and choose **Space tools** > **Integrations** > **Mail**
   Or, if your space uses the Documentation theme, choose **Browse** > **Mail** in the header.
   A list of email messages in the space is displayed in reverse chronological order
2. Do either of the following:
   - Delete an individual email message by choosing the trash
     🗑
     icon beside it.
   - Delete all email messages within the space by choosing **Delet e All**.

⚠ Email messages deleted using the 'Delete All' option can't be restored.

Space administrators can restore deleted email messages, provided they were deleted individually.

Restore mail that was deleted:

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar
2. Choose **Trash**
   You'll see a list of email messages and other content deleted from the space.
3. Choose **Restore** beside the email message you want to restore

If your space uses the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **Trash** in the left-hand panel
   You will see a list of email messages and other content deleted from the space.
3. Choose **Restore** beside the email message you want to restore

Import Mail from an mbox

Confluence allows you to import mail from mbox files located either on your local system or in a specified location on the Confluence server. Confluence will store the imported email messages in the space's mail archive.

You need to be a space administrator to import mail for a space. See Space Permissions.

NB: You may need to enable the **Confluence Mail Archiving Plugin** as it is disabled by default.

To import mail from an mbox file:

1. Go to the space and choose **Space tools** > **Integrations** from the bottom of the sidebar.
2. Choose **Mailbox Import**.
   - To import from a location on your file system: Browse to the location of the mbox file, select the file and then choose **Import**.
   - To import from the Confluence server: Enter the location of the mbox file on the server, then choose **Import**.

If your space uses the Documentation theme:

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

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

On this page:

- Delete mail from a space:
- Restore mail that was deleted:

Related pages:

- Mail Archives
- Add a Mail Account

Related pages:

- Mail Archives
- Add a Mail Account

2. Choose **Mailbox Import** in the space administration options.
   - To import from a location on your file system: Browse to the location of the mbox file, select the file and then choose **Import**.
   - To import from the Confluence server: Enter the location of the mbox file on the server, then choose **Import**.

Notes

- Only **site spaces** can store mail archives. Personal spaces cannot. See Spaces for an explanation of site spaces and personal spaces.
- Confluence mail archiving is an optional feature. This means that the 'Mail' options may be disabled and will therefore not appear in the Confluence user interface. Mail archiving features are contained in a bundled plugin. To activate mail archiving features in Confluence, enable the plugin – choose

  ⚙ ▾

  > **General Configuration** > **Manage Add-ons**. Then choose **System** in the drop down, and enable the **Confluence Mail Archiving Plugin**.
- For security reasons mail can only be imported from a specified location in the Confluence server's file system. We recommend administrators create a folder in their Confluence home directory, add the system property `confluence.mbox.directory` and specify the location for mailboxes to be imported from . Mail cannot be imported from the server until this system property is set. See Configuring System Properties.

Gadgets

Gadgets allow you to add dynamic content to a Confluence page or JIRA application dashboard. Confluence can display gadgets that support the OpenSocial specification, including third party gadgets.

For more information about Atlassian gadgets, see the introduction to Atlassian gadgets and the big list of Atlassian gadgets.

To see a list of available gadgets in your Confluence site go to **Help** > **Available Gadgets**.

The following gadgets are bundled with Confluence.

Name	Description
Confluence Page Gadget	The Confluence page gadget allows you to show content from a page on your Confluence site in a gadget. You can optionally configure the gadget to display links to view and/or edit the page on your Confluence site. The page gadget can also be displayed in canvas view, so that it takes up all of the space provided by your dashboard.
Activity Stream Gadget	The activity stream gadget is similar to the recently updated macro and shows a list of the most recently changed content within your Confluence site.
Confluence News Gadget	The Confluence news gadget is an example of a 'news feed' gadget that shows a list of recent Confluence Product Blogs and events at Atlassian.

On this page:
- Add a Confluence gadget to a page
- Add a JIRA gadget to a page
- Add a Confluence gadget to your JIRA application dashboard

Related pages:
- Gadget Macro
- All Atlassian gadgets
- External Gadgets

Confluence Quicknav Gadget	The Confluence QuickNav gadget provides the Quick Navigation Aid f unctionality for searching the Confluence site. For more information on how to use this feature, refer to the Quick Navigation Aid section of Search.

Add a Confluence gadget to a page

Where a similar result can be achieved using standard Confluence macros (for example using the Include Page Macro instead of Confluence Page Gadget), we recommend using macros as they have better performance and are easier to configure than gadgets.

To add a Confluence gadget to a page:

- In the editor go to **Insert** > **Other Macros**.
- Select the gadget you wish to add, and use the preview area to configure the gadget.
- Choose **Insert**.

See Gadget Macro for more information on adding the macro to your page, configuring the gadget itself and display options.

Add a JIRA gadget to a page

For displaying basic JIRA information, such as issues and charts, we recommend using the JIRA Issues Macro and JIRA Chart Macro as these macros have better performance and are easier to configure than gadgets.

If the JIRA information you want to display is not available from either of these macros a gadget will likely do the trick.

To add a JIRA Gadget to a Confluence page:

- In the editor go to **Insert** > **Other Macros**.
- Select the gadget you wish to add, and use the preview area to configure the gadget.
- Choose **Insert**.

See Gadget Macro for more information on adding the macro to your page, configuring the gadget itself and display options.

> If you don't see any JIRA Gadgets in the macro browser, ask your Confluence administrator to add the JIRA Gadget urls to the list of authorised external gadgets in Confluence, and check that the application link between Confluence and your JIRA application is configured correctly.

Add a Confluence gadget to your JIRA application dashboard

To add a Confluence gadget to your JIRA dasboard:

1. Go to the dashboard by selecting the **Dashboard** link in the header.
2. On the dashboard, Click **Add Gadget**.
3. Use the gadget wizard to choose the gadgets you want to add.

> If you don't see any Confluence gadgets in the JIRA gadget directory, ask your JIRA administrator to add the gadget URLs as follows.
>
> To add a Confluence gadget to your JIRA application's gadget directory:
>
> 1. In Confluence, go to **Help** > **Available Gadgets** and copy the gadget URL for the gadget you want to make available in JIRA.

2. In JIRA, go to the dashboard and choose **Add Gadget**.
3. Choose **Manage Gadgets** or **Add Gadget to Directory** (depending on your JIRA application and version)
4. Paste in the Confluence gadget URL and choose **Add Gadget**.

The gadget will now be available from the JIRA Gadget Directory.

Activity Stream Gadget

The activity stream gadget is similar to the recently updated macro and shows a list of the most recently changed content within your Confluence site.

In addition to showing a list of most recently changed content, the activity stream gadget also groups activities by separate date, and provides an RSS feed link to its content in the top-right corner.

Activity Stream Gadget Properties

Properties are settings for Confluence gadgets that allow the user to control the content or presentation of data retrieved by the gadget. These are similar to a Confluence macro's parameters. The table below lists relevant properties for this gadget.

These properties are located in the **preview** panel in the macro browser.

Property	Required?	Default	Description
Title	Yes	None	Adds a title to the top of the Activity Stream.
Global filters	No	None	Allows you to add filters to the gadget including: • space • username • update date • JIRA issue key (if your Confluence instance is integrated with a JIRA application)
Available streams	Yes	All	If you have application links to other sites, JIRA or another Confluence site, you can choose to include activity from those streams also.
Display options: limit	No	10	Specify the maximum number of results to be displayed. A maximum of 10 results will be displayed by default. The maximum value that this property can accept is 100.
Display options: Refresh Interval	No	Never/false	Specify the time interval between each 'refresh' action undertaken by the activity stream gadget. A refresh makes the activity stream gadget reflect any new activity that has been conducted on the Confluence site.

Confluence News Gadget

The Confluence news gadget is an example of a 'news feed' gadget that shows a list of recent Confluence Product Blogs and events at Atlassian.

Clicking an item in the Confluence news gadget takes you directly to the selected blog post of the Confluence Product Blogs feed or event details on the events at Atlassian page.

Clicking **More news** takes you to the Confluence Product Blogs feed page and clicking **More events** takes you to the events at Atlassian page.

Screenshot: The News Gadget

Confluence: News Gadget Properties

Properties are settings for Confluence gadgets that allow the user to control the content or presentation of data retrieved by the gadget. These are similar to a Confluence macro's parameters. The table below lists relevant properties for this gadget.

Property	Required?	Default	Description
Show News?	No	True	Shows a short list of the most recent Confluence Product Blogs.
Show Events?	No	True	Shows a short list of the most recent events at Atlassian.
Show Banners?	No	True	Shows any banner advertisements (if available).

Confluence Page Gadget

The Confluence page gadget allows you to show content from a page on your Confluence site in a gadget. You can optionally configure the gadget to display links to view and/or edit the page on your Confluence site. The page gadget can also be displayed in canvas view, so that it takes up all of the space provided by your dashboard.

On this page:

- Confluence Page Gadget Properties
- Working Macros

Macros that work with the page gadget
Please note, not all macros work with the page gadget. Please refer to the Working Macros section below for more information.

Screenshot: The Confluence page gadget displaying a sample page

Confluence Page Gadget Properties

Properties are settings for Confluence gadgets that allow the user to control the content or presentation of data retrieved by the gadget. These are similar to a Confluence macro's parameters. The table below lists relevant properties for this gadget.

Property	Required?	Default	Description
Space	No	None	Specify the space that your desired page is located in. Suggestions will display in a dropdown when you start typing. (Note, this property is only used to make searching for pages easier. It is not required.)
Page	Yes	None	Specify the page that you want to display in your gadget. Suggestions will display in a dropdown when you start typing.
Show View Link	No	Yes	Select whether to display a link to view the page on your Confluence site. Clicking the link will open the page in Confluence.
Show Edit Link	No	No	Select whether to display a link to edit the page on your Confluence site. Clicking the link will open the page for editing in Confluence.
Refresh Interval	No	Never/false	Specify the time interval between each 'refresh' action undertaken by the page gadget. A refresh makes the activity stream gadget reflect any new activity that has been conducted on the Confluence site.

Working Macros

The Confluence page gadget will only render a subset of the macros that are used in Confluence correctly. Refer to the table below for the list of macros that work and do not work with the page gadget and known limitations.

> Some of the issues with macros in the page gadget can be worked around, if you are comfortable developing in Confluence. Please see Troubleshooting Macros in the Page Gadget for more information.

Key:
✅ *Works with the page gadget*
✅ ** Partially works with the page gadget*
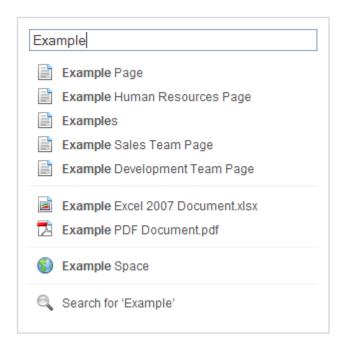❌ *Does not work with the page gadget*

Macro	Works with page gadget?	Limitations
Activity Stream		You cannot have another gadget embedded within the Confluence Page Gadget
Anchor (within a page)	*	Opens in a new page
Attachments		N/A
Blog Posts		N/A
Chart		N/A
Children Display		N/A
Content By Label		N/A
Content By User		N/A
Excerpt		N/A
Gallery		N/A
Include Page		N/A
Info		N/A
Labels List		N/A
Livesearch		N/A
Note		N/A
Metadata		N/A
Metadata Summary		N/A
Pagetree Search		N/A
Pagetree		N/A
Panel		N/A
Quick Nav		You cannot have another gadget embedded within the Confluence Page Gadget
Recently Updated		N/A
RSS Feed		N/A
Section & Column		N/A

Spaces List	✅	N/A
Table of Contents	*	Works, however links will be opened in a new browser window when clicked.
Team Calendars for Confluence		See here for the Improvement Request: **TEAMCAL-862** - Make Team Calendars display in a Confluence Page Gadget, for use on a JIRA Dashboard `REOPENED`
View File (PDF or PPT)	*	Works, but you may need to refresh the gadget the first time (see CONF-19932).
Widget Connector	*	Only works for some content: • **Works:** blip.tv, Episodic, Flickr, Google Calendar, presentations on Google Docs, MySpace Video, Scribd, Skitch.com, SlideRocket, SlideShare, Viddler, Vimeo, YouTube, Dailymotion, Metacafe, FriendFeed, Yahoo Video, Wufoo HTML Form Builder • **Does not work:** FriendFeed, Google Gadgets, Google Video (consumer service discontinued), Twitter, Widgetbox, DabbleDB, BackType

Confluence Quicknav Gadget

The Confluence QuickNav gadget provides the Quick Navigation Aid functionality for searching the Confluence site. For more information on how to use this feature, refer to the Quick Navigation Aid section of Search.

Screenshot: Using the QuickNav Gadget

Confluence QuickNav Gadget Properties

This gadget has no properties and cannot be customised.

Confluence Use-Cases

This section describes some specific use cases for Confluence.

- **Using Confluence for technical documentation**

  A technical communicator's guide to using Confluence – see Develop Technical Documentation in Confluence.
- **Setting up a knowledge base**

  A support engineer's guide to using Confluence as a knowledge base

– see Use Confluence as a Knowledge Base.

- **Setting up an intranet**

  A quick guide to setting up an intranet wiki – see Use Confluence as your Intranet.
- **Confluence for software teams**

  A series of blog posts to help your agile team get the most out of Confluence. See Confluence for Software Teams.

You may also be interested in:

- **Confluence markup formats**

  The syntax and usage of Confluence wiki markup and the Confluence XHTML-based storage format – see Confluence Markup.
- **Gadgets**

  How to add gadgets to a wiki page, or use Confluence gadgets on other sites – see Gadgets.

Develop Technical Documentation in Confluence

Confluence is a flexible platform with a range of features and add-ons that can help you capture, distribute, and update your technical documentation. Below are some tips to help you get your technical documentation site started, and to save you time and effort managing your documentation's life cycle.

Create your Documentation Space

Creating spaces in Confluence is quick and easy. All you need to do to get your documentation started is choose **Spaces** > **Create space** from the Confluence header. To make things even easier, choose the 'Documentation Space' option in the create space dialog; it'll give you a custom home page with a search box (the livesearch macro) to search just your documentation space, a recently updated macro, and a few other goodies.

Give your space a name, and Confluence will automatically create the home page and space key for it (change the space key if you're not happy with the one Confluence chooses for you). Feel free to customise the home page at any time; what it looks like is completely up to you!

Choose the Documentation theme (optional)

The Documentation theme is specifically formatted for technical documentation, and provides a left-hand navigation bar and space-specific search. If necessary, you can configure the Documentation theme to add your own page header and footer, or customise the default left-hand navigation bar. You should note though, that the Documentation theme has been deprecated and isn't compatible with some newer features like the JIRA links button.

If you'd like to set the space's theme to 'Documentation', choose **Space tools** > **Look and feel** from the bottom of the sidebar, then choose **Themes** and select the Documentation theme.
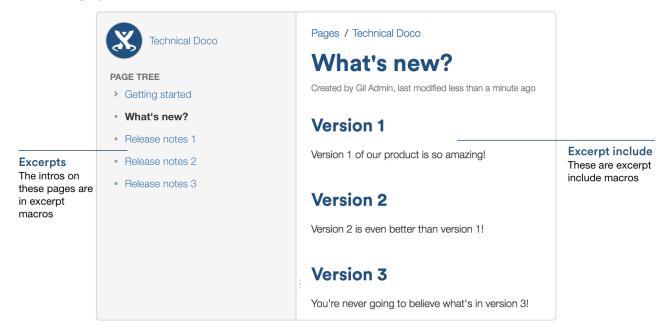
Save time by re-using content

If there's something you're going to use multiple times in your documentation space – whether it's a word, sentence or paragraph; an image; a product version number; or anything else – you can create it once and include it on as many pages as you like (or use it in the header and/or footer). Inclusions not only save you typing the same thing many times, they also make it easier when things change – it's much better to update the info in one place, than 47!

There are 3 macros that allow you to re-use content:

- The **Excerpt macro** to define a re-usable section, or 'excerpt', on a page – add content inside this macro, and you can reuse it on as many pages as you like.
- The **Excerpt Include macro** (`excerpt-include`) to include the contents of an excerpt on another page.
- The **Include Page macro** (`include`) to include the entire content of a page on another page.

For example, let's say you create release notes for each major release of your product, and you want to include the intro from each release notes page on a 'what's new' page. Place each release notes intro in an **E xcerpt macro**, then add an **Excerpt Include macro** for each set of release notes to the what's new page. Your intros will magically appear on the what's new page, and if you update the release notes it'll automatically update the what's new.

Another example is one of the ways we use the **Include Page macro**. Whenever the ellipsis (
●●●
) appears in our documentation – for example, go to
●●●
 > **Copy** – it's actually an **Include Page macro**. We have a page with just that image on it, so we can include it whenever we need an ellipsis.

Why do we do use an **Include Page macro** for one tiny image? Well, just in case that UI element is ever changed. If we attach the image to every page, there might be 50 pages we need to update when things change; if we use an **Include Page macro**, we update once and it's changed everywhere. Doing it this way also allows us to know how many pages we're using the image on. By going to
●●●
 > **Page Information**, we can see how many incoming links there are to this page, and that tells us how many pages use the image.
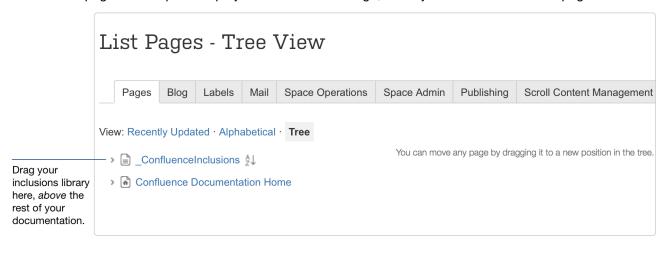
Create an inclusions library (optional)

You can include content from any Confluence page, but you may want to create an 'inclusions library' to hold content that's specifically for re-use. The inclusions library isn't a specific feature of Confluence; the pages in

the inclusions library are just like any other Confluence page. This is just a *technique* you can use if you want a place to store content that's specifically for re-use.

To create your inclusions library:

1. Choose **Create** and create a new page in your space
2. Enter a suitable title. We use '_ConfluenceInclusions' (the underscore before the title helps to let people know this page is special)
3. Enter some content and **save** the page
   We enter text explaining the purpose of the inclusions library and how to re-use the content
4. Choose **Space tools** > **Reorder pages** (or **Browse** > **Pages** if you're using the Documentation theme) and drag your new page *above* the space homepage
5. Go to your new inclusions page and choose **Create** to add child pages containing your re-usable content

Because you've moved the pages to the *root* of the space, they won't appear in the page tree in the sidebar. The pages will be picked up by other searches though, as they're normal Confluence pages.

List Pages - Tree View

| Pages | Blog | Labels | Mail | Space Operations | Space Admin | Publishing | Scroll Content Management |

View: Recently Updated · Alphabetical · **Tree**

Drag your inclusions library here, *above* the rest of your documentation.

› 📄 _ConfluenceInclusions ⇅

› 🏠 Confluence Documentation Home

You can move any page by dragging it to a new position in the tree.

Use page templates

Creating one or more page templates can be a real time-saver if you're creating a lot of pages with the same layout. If you're constantly adding the same macros, like panels and table of contents, save yourself from RSI and put them into a template – you can start with one, but make as many as you need to maximise your efficiency.

To create a page template that's available in all spaces:

1. Go to
   ⚙ ▾
   > **General Configuration**
2. Select **Global Templates and Blueprints** from the list on the left
3. Choose the **Add global page template** button at the top-right
4. Create your template page and choose **Save**

For detailed info on page templates, see Create a Template.

> To get to **Global Templates and Blueprints**, or any other admin page quickly, hit / on your keyboard and start typing the name of the admin page you're looking for.

Draft your work

When you're creating a new page in your documentation, you'll likely want to do it over time, saving as you

go, and have a select few people review it to provide feedback. A loose description of this workflow is 'draft, review, publish'.

You don't want any half-finished pages being seen by your users, and most documentation needs to be reviewed before it's finalised, so here's a technique for drafting pages and allowing for review:

1. Create a page and restrict its permissions
   For example, you might restrict viewing to a group of people such as your team, or a few select individuals. On a public site, you might restrict viewing to staff members, so that the general public can't see the page.
2. Write your page content
3. Share the page with your reviewers and ask them for feedback (make sure you haven't restricted them from seeing the page!)
   The reviewers can add comments to the bottom of the page or highlight text to add a comment inline. If you give them permission, they can also edit the page content directly.
4. Publish the page when ready, by doing the following:
   a. Delete any comments on the page
   b. Remove page restrictions so that your audience can see it

You've now published your page. The space permissions and site permissions now determine who can see and/or update the page.

Use links and anchors

Add links

In any documentation site, it's essential to be able to link from one page to another, and often to specific sections on a page. You can add any URL to a Confluence page and Confluence will automatically detect it and turn it into a link.

If you paste the URL for another page in your Confluence site, Confluence will display the link text as the page name and turn it into a relative link, meaning if the name of the page changes, Confluence will adjust the link so it doesn't break.

Add and link to anchors

The anchor macro allows you to create anchors in your documentation, which can be linked to from anywhere. I've added an anchor at the top of this page so you can click to go back to the top.

To add a macro and link to it from the same page:

1. Type {anchor in the editor, select the anchor macro and give your anchor a name (top in my example)
2. Select the text that'll link to the macro and hit Ctrl+K (Windows) or Cmd+K (Mac) (this opens the link dialog)
3. Choose **Advanced** from the options on the left and type # followed by your anchor name (#top in my example)

Check out our documentation for links and anchors to get the full rundown on linking to anchors on other pages and other anchor goodness.

Useful macros

Confluence ships with a great range of macros, and there are a few that are particularly useful in technical documentation. Here's a few:

Table of contents macro

The Table of Contents macro helps people navigate lengthy pages by summarising the content structure and

providing links to headings used on the page. The best part is, you don't need to do anything except add the macro; once you've added it, it'll automatically detect headings and add them to the table of contents.

Tip, Note, Info, Warning, and Panel macros

Often when creating documentation, there are elements of a page that you want to highlight or draw the the viewers' attention to. Confluence ships with the Tip, Info, Warning, Note and Panel macros, which will help you focus a viewer's attention on a particular part of your content.

> **Tip of the day**
> Use the tip macro to give your readers handy hints!

Keep track of page updates

In Confluence, it's quite usual for a number of different people to update a single page. Technical writers need to know what happens to our documents, both during review and after publication.

Watch pages or the space

So that you know when changes are made, it's a good idea to *watch* pages or even the entire space. That way, when changes are made to pages you're watching, or someone comments on them, you'll get an email notification letting you know who changed what.

Whenever you're on a page in your documentation space, choose the **Watch** button at the top-right of the page. From there, you can choose to watch just that page, or all pages in the space.

View page history

Confluence creates a new version of the page every time someone edits the page. The page history shows all the versions, with date, author, and any comments made on the update.

To view page history, go to the page and choose

•••
> **Page History**

On the page history view, you can:

- View the content of a specific version of the page.
- Revert to (restore) a specific version.
- Select any two versions and ask for a comparison, to see what has changed between those two versions.

Take a look at Page History and Page Comparison Views for a detailed explanation.

Show a list of contributors

If you want to see at a glance who's updated a page or pages, you can add the contributors macro. This macro displays a customisable list of people who've contributed by creating, editing, or, optionally, commenting on the page.

Customise PDF export

If you're planning to provide a PDF version of your documentation – whether it be for email, download, print, or any other form of delivery – you can customise the look of the PDF by adding a title page, header, and footer.

The process you take depends on whether you're trying to customise the PDF export for one space or for

your whole site, so, if you're keen to make these changes, take a look at our page on Customise Exports to PDF for more detailed instructions.

Other useful tools and add-ons

Confluence is already a great tool for technical documentation, but you can still add to it depending on your documentation and workflow needs. Here are some useful add-ons available on the Atlassian Marketplace, most of which we use ourselves, which can extend the functionality of Confluence.

Scroll Versions (supported)

Scroll Versions, by K15t, allows you to tie versions of your documentation to versions of your product, so that when a new version of your product ships you can publish that version of your documentation. Create as many versions of your documentation as you like, make the changes you need to, and keep them up your sleeve until release time. You can even publish different variations of your documentation – like if you have versions of your documentation for different operating systems – to different spaces or Confluence instances.

Copy Space (unsupported)

The Copy Space add-on does what its name suggests; it allows a space administrator to copy a space, including the pages within the space. Great for when you want a space template that you can copy to create other spaces.

This plugin is also useful when you need to archive a copy of a current space at a particular point in time, like when you're moving from one version of your product to the next – copy the space, give it a new name, and keep it wherever you like, all without losing the existing space.

ⓘ At this point this plugin won't copy page history, blog posts and email.

Scroll PDF Exporter (supported)

If you're going to produce a PDF of your documentation space, wouldn't you like it to be professionally formatted? The Scroll PDF Exporter, by K15t, lets you style single pages or whole spaces for export, using handy PDF templates.

Gliffy (supported)

Create diagrams, wireframes, flowcharts and more with Gliffy. Gliffy features a highly intuitive drag-and-drop interface, and allows you to export your diagrams in multiple formats, including: JPEG, PNG and SVG. Add Gliffy flowcharts, UI wireframes, and network diagrams directly to your Confluence pages to communicate your ideas visually, making them easy to understand and faster to spread through your team.

Lucidchart (supported)

Lucidchart is available in versions for Cloud and Server, and allows you to create and insert diagrams within your Confluence Cloud environment. Quickly draw flowcharts, wireframes, UML diagrams, mind maps, and more inside our feature-rich editor.

The server version also comes with a free Visio viewer, so you can view Microsoft Visio (.vsd) files, Visio stencils (.vss) and it also supports exporting back to Visio.

Use Confluence as a Knowledge Base

A knowledge base is a repository for how-to and troubleshooting information. Knowledge Bases are commonly used by IT Support teams, but can be useful for procedural and troubleshooting information in any organisation or team.

What do people want out of a knowledge base? Using an IT Support team as an example:

- Customers want fast access to a solution, and

relevant search results.
- Help desk staff want to be able to create new articles quickly.
- Help Desk team leads wants the space to be self curating, and do not want to spend a lot of time manually organising content.
- Everyone wants a way to be notified when articles they are interested in have been updated or important notices are added.

Create a knowledge base space

⚠ You'll need the Create Space global permission to do this.

To create your knowledge base space:

1. Choose **Spaces** > **Create space** > **Knowledge base space**
2. Choose **Space Tools** > **Permissions** to set permissions for the space, including anonymous access
3. Choose **Create** > **How-to** or **Troubleshooting** and follow the prompts to create your first knowledge base article

The knowledge base space blueprint includes everything you need to get started, including article templates, and a pre-configured homepage with Livesearch and Content By Label macros.

Page labels are essential in knowledge base spaces. These are used to add topics to your articles, and allows your knowledge base to become self-organising over time.

Users will generally find articles by searching, and using the topic navigation on the homepage and end of each article, rather than navigating through a tree-like page hierarchy.

When starting off your knowledge base space, it's a good idea to brainstorm a few topics to get started.

Customise your knowledge base space

> You'll need Space Admin permissions to do this.

To make it easy for your users to create knowledge base articles, such as your help desk or support team, we recommend customising the how-to and troubleshooting article templates to make them relevant for your organisation. The more guidance and structure you can put in your template, the faster it will be for your team to create great articles.

To edit the article templates:

1. Go to **Space Admin** > **Content Tools** > **Templates**.
2. Edit the **How-to** or **Troubleshooting** article templates.
3. Add headings and instructional text (choose **Template** > **Instructional Text**).

You can also add additional templates, such as a policy or procedure page templates.

We also recommend customising the look and feel of your space. Simple changes like a space logo and welcome message can make a huge difference.

To change the look and feel:

- Add a space logo and useful shortcuts to the sidebar (choose **Space Tools** > **Configure Sidebar**)
- Edit the homepage to add a custom welcome message.
- Edit the colour scheme (choose **Space Tools** > **Look and Feel** > **Colour Scheme**).

Provide communication and notification options

Channels of communication with your audience, internal or external, are essential in a good knowledge base. Here are some out-of-the-box options:

- **Blog** - blog updates and important notices, and encourage people to watch for new blogs in your

space.
- **Watch** - encourage people to watch pages that interest them, or watch the entire space.
- **Comments** - allow logged in users (or even anonymous users) to comment on knowledge base articles. This is a simple way to connect with your end users.
- **RSS** - create an RSS feed and add the link to your knowledge base homepage (choose **Help** > **Feed Builder**). Alternatively encourage users to create their own feed - useful if they want to keep up with particular topics (labels), rather than receive notifications for the whole space.

Integrating your knowledge base with other Atlassian products

If your Confluence site is connected to another Atlassian product (via an application link), you can make use of these great integration features:

- **If you use any JIRA application -** add a JIRA Issues macro to your troubleshooting article to provide quick access to known issues. This has the added advantage of automatically updating when an issue is resolved or its status changes. One simple way to do this would be to add some labels to JIRA to indicate the issue should appear in the knowledge base (for example 'printer-kb'), and then add a JIRA Issues macro with a query like 'label = 'printer-kb and status <> resolved'' on all articles with the printer topic.
- **If you use JIRA Service Desk** - link a Confluence space to be used as a knowledge base. Users (including those without a Confluence license) can search your knowledge base directly from within the Service Desk customer portal.
- **If you use Questions for Confluence** - add a Questions list macro to troubleshooting articles, to highlight the top questions with the same topic as the article, and an Ask a Question button to the knowledge base homepage.

Extending your knowledge base with third party add-ons

The Atlassian Marketplace has a large number of add-ons for Confluence. A common addition to Knowledge Base spaces is a survey or form tool, which enables you to get feedback on the usefulness or usability of your knowledge base articles.

Search for 'knowledge base' on Marketplace and see if there is an add-on that's right for your knowledge base.

Use Confluence as your Intranet

Your intranet is the hub of your organisation. When choosing your intranet platform, you need to ensure that the system is simple enough for non-technical users, information and content can be shared easily, and access is restricted to those within your organisation.

Confluence has a host of great out-of-the-box features that allow you to share and collaborate with your colleagues, while keeping your information secure. Share things like procedures, specifications and important files – or organise company events and functions – and get your teams working together. It's one place to share, find, and collaborate to get work done.

Create your community

It's quick and easy to add users to your Confluence site. Allow people to add themselves as users of the site; invite people to sign up by sending them an invitation link; add new users manually; or use an existing directory – like an LDAP directory – for authentication and to manage users and groups.

Whichever way you choose, you can quickly build a community of Confluence users and give them access to your intranet; you'll also have a ready-made people directory.
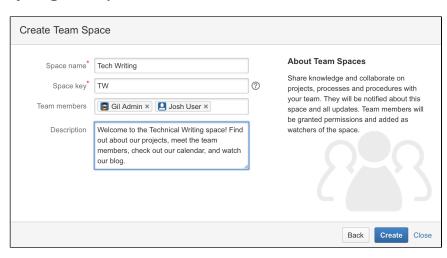
Match your company branding

Upload your company logo, and Confluence's auto look and feel will change the colour scheme to match. It'll make your intranet feel more familiar to your colleagues, and help with adoption.

A space for everything, and everything in its space

A Confluence space is essentially a container for a group of pages and blog posts with related content.

When you're starting out with Confluence, the easiest way to organise things is to create a space for each team or department within your organisation. Each team's space is then a place for them to create and share pages, blog posts, meeting notes, files, and much more – and becomes the *place to go* for team members to get the information they need.

Create Team Space	
Space name* Tech Writing	**About Team Spaces**
Space key* TW ⓘ	Share knowledge and collaborate on projects, processes and procedures with your team. They will be notified about this space and all updates. Team members will be granted permissions and added as watchers of the space.
Team members 👤 Gil Admin × 👤 Josh User ×	
Description Welcome to the Technical Writing space! Find out about our projects, meet the team members, check out our calendar, and watch our blog.	
	Back **Create** Close

Just choose **Spaces** > **Create space** from the header, and Confluence provides a list of space blueprints to help get you started.

Each space can have its own colour scheme and has a customisable home page, which you can edit to suit your purpose – like displaying and tracking team goals and displaying a list of team members. Use the built-in 'Team Space' template to automatically add all members of the team to the homepage, to help everyone get to know each other.

You can set permissions for each space, so if there's sensitive information that should only been seen by certain users or groups, it's easy to secure it with Confluence.

Don't feel restricted to creating spaces for teams though; you can also create spaces for projects (large or small), events, and anything else where you want to collect information under a common heading or permissions structure.

Once you have some spaces set up, create some pages and blog posts to give your colleagues an example of how Confluence can be used, then invite them to create their own pages and blogs.

Add a personal space

Every Confluence user, including you, can also create their own personal space; it can be a place to keep your own work, add shortcuts to your most used content, and you even get your own blog for sharing your ideas and opinions with the rest of your organisation (or just those that you want to see them).

Create pages, meeting notes and more

You can create pages for anything you want in Confluence - meeting notes, project plans, decisions, and more. Pages are editable so others can contribute and keep them up to date after you create them. Choose **Create** from the Confluence header and choose a blank page, or use a template to get you started.

Type your page, change its layout, add images and links, and do it all without any specialist skills or training. You can also attach files – allowing everyone in a team access to assets that are critical to the project – like mockups and requirements. You and your colleagues can like the page, and comment on it to start a conversation about the content.

Confluence also offers a series of useful built-in page blueprints, which help you with the content and formatting of the page. The meeting notes and decisions blueprints are two that can be really useful when others need to be in-the-know about what happened, and why it happened.

Avoid the reply-all and blog about it

Each space you create in Confluence has its own blog, where you and your teams can share news and

events, discuss important projects and developments, or congratulate a teammate for a special effort; blogging is a great way to foster company culture and celebrate achievements across your organisation.

You can watch any blog to make sure you get updated when there's a new post. Blog posts are automatically organised by date, and grouped by year and month, so they're also easy to find.
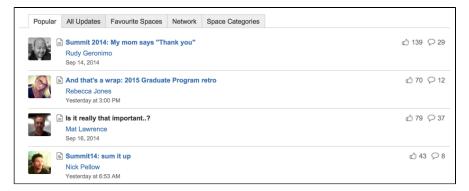
Share stuff that matters

If you need to be sure that the right people see a page or blog post, Confluence offers a range of ways to make sure you can get their attention. Type the @ symbol and the name of a Confluence user to mention them in a page, blog post, or comment. They'll get an email notification that you've mentioned them, with a link to the page, post or comment.

There's also a **Share** button at the top right of every page. Type the name or email address of a user or group and send them a short message with a link to the content you're sharing.

Watch and learn

Don't miss out on important updates. Watching spaces, pages, and blogs is a great way to stay up-to-date with what's happening in your own team, or any other team or person you need to keep up with. When you watch something, you'll get email updates when changes are made or a comment is added.

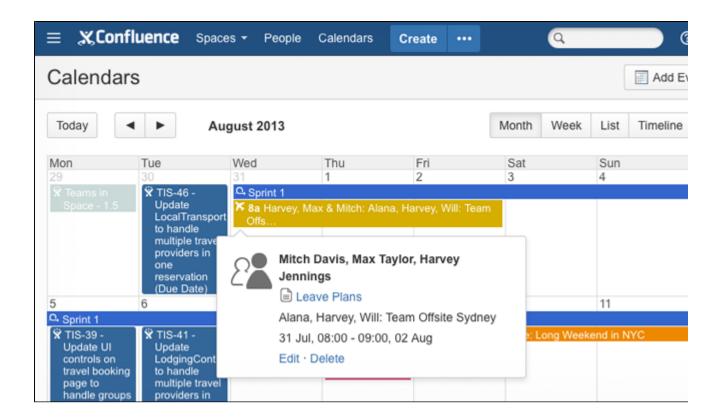The Confluence dashboard also has a recent activity feed, which allows you and your team to see what's trending throughout the company or in your network.

Let team collaboration take on a life of its own

If you want to communicate in-the-moment, HipChat – our private service for chat, video, and screen sharing – is built for teams. Share ideas and files in persistent group chat rooms, or chat 1-1 for that personal touch. HipChat integrates with Confluence, so you can get real-time notifications in HipChat whenever new information is shared in Confluence pages and blogs.

Got a question (that's not rhetorical)? Why not ask your team? Questions for Confluence is an add–on for Confluence that gives you knowledge sharing with your own Q&A service. Run company-wide polls to gauge reaction to a new marketing initiative, or let people vote on the venue for the Christmas party. Ask questions, get answers, and identify experts.

You can also try Team Calendars for Confluence for organising and sharing team events, leave, and other important appointments. Embed each team's calendar on their home page so that everyone knows what's happening, and when.

Confluence for Software Teams

Welcome to the Software Team's guide to using Confluence.

Check out the articles in this collection:

How to write product requirements

How to build a release planning page

Creating insightful customer interview pages

Create sprint retrospective

How to make better decisions

How to document releases

and demo pages (like a BOSS)   **as a development team**   **and share release notes**

Use blogs to share your development team's progress

How to create technical and onboarding documentation

Like what you see? Start creating these pages and more in Confluence!

Try it for free

Confluence Administrator's Guide

About the Confluence Administrator's Guide

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

For information on creating and administering spaces, See Spaces.

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

- Getting Started as Confluence Administrator
- Manage Users
  - Add and Invite Users
  - Delete or Disable Users
  - Restore Passwords To Recover Admin User Rights
  - Edit User Details
  - Change a Username
  - Managing Site-Wide Permissions and Groups
  - Configuring User Directories
- Managing Add-ons or Plugins
- Writing User Macros
  - User Macro Template Syntax
- Customising your Confluence Site
  - Changing the Look and Feel of Confluence
  - Changing the Default Behaviour and Content in Confluence
- Integrating Confluence with Other Applications
  - Linking to Another Application
  - Configuring Workbox Notifications
  - Integrating JIRA and Confluence
  - Registering External Gadgets

Downloads

Download the Confluence documentation in PDF format.

Other Resources

Confluence Installation and Upgrade Guide

Confluence Knowledge Base

Atlassian Answers
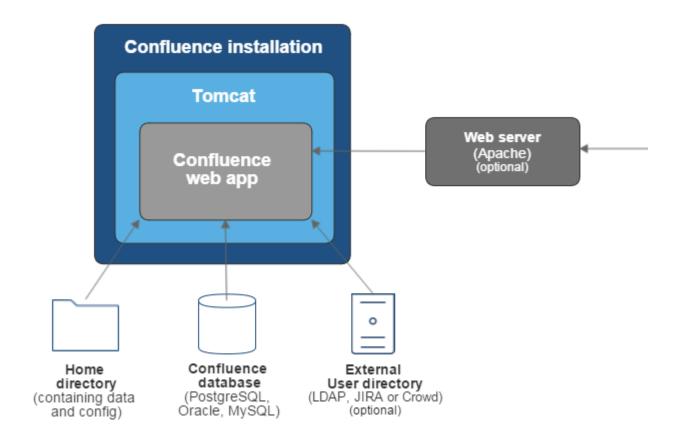
Getting Started as Confluence Administrator

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

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

On this page:
- Quick access to administrative functions via Confluence search
- How to administer and configure Confluence
- Getting started on a new Confluence site
- Getting to know an existing Confluence site
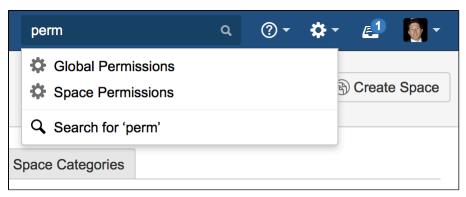
Diagram: A Confluence installation

Quick access to administrative functions via Confluence search

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

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

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

How to administer and configure Confluence

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

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

To access the Confluence Administration Console:

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

   ⚙️ ▾

   > **General Configuration** in the header

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

Getting started on a new Confluence site

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

- Decide whether you want to allow public (anonymous) access to your site.See Setting Up Public Access.
- Add a space and some content. See Create a Space then Pages and Blogs.
- Invite some users to your site. See Add and Invite Users.

- Decide whether you will manage your users in Confluence or hook up an external LDAP directory. See Configuring User Directories.
- Make sure you have set up an email server. The above task list will include this step, but it is worth mentioning it here again. Email notifications are an important part of collaborating on Confluence. See Configuring a Server for Outgoing Mail.

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

Getting to know an existing Confluence site

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

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

Manage Users

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

Confluence user management

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

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

On this page:

- Confluence user management
- Authentication
    - Seraph
    - XML-RPC and SOAP authentication
    - Password authentication
- Earlier user management frameworks

Related pages:

- Configuring Confluence Security

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

Authentication

Seraph

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

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

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

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

XML-RPC and SOAP authentication

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

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

Password authentication

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

Earlier user management frameworks

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

Add and Invite Users

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

- **By user signup**: If user signup is enabled on your Confluence site,

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

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

Allow user signup

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

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

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

You need Confluence Administrator or System Administrator permissions to change the signup options.

To set the user signup options:

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

Manage user signup notifications

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

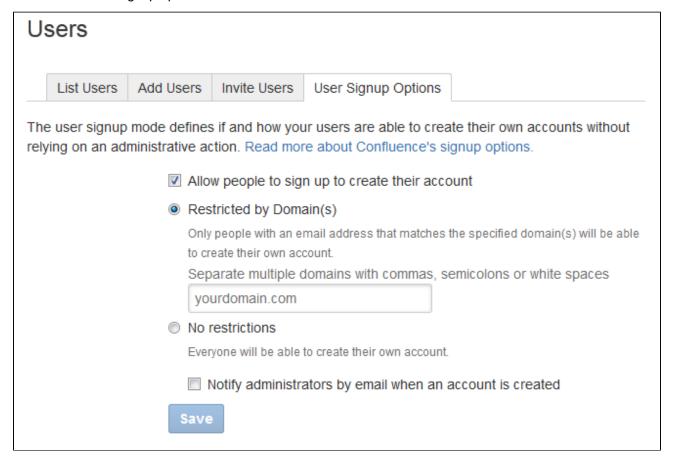link or by clicking the invitation URL sent by an administrator.

To disable this notification:

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

Screenshot: User signup options

Invite people to sign up

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

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

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

To invite people to sign up:

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

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

Reset the invitation link

The invitation link includes a security token, like this:

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

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

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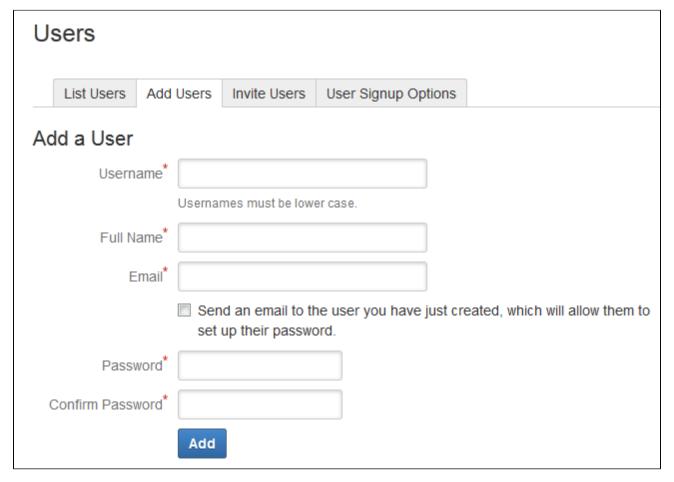
Screenshot: Inviting users

Add users manually

To add a new user:

1. Choose

    ⚙ ▾

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

Screenshot: Adding users

Notes

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

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

    See Managing Multiple Directories.

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

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

  ⚙ ▾

  > **General Configuration**> **Security Configuration**. See Disabling the Built-In User Management.

Delete or Disable Users

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

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

Related pages:

- Manage Users
- Configurin g User Directories

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

To delete or disable a user:

1. Choose

   ⚙ ▾

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

Screenshot: Administering a user

View User: ewan

« Back to Users

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

View Profile   Edit Groups   Edit Details   Set Password   Delete   Disable

Notes

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

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

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

Restore Passwords To Recover Admin User Rights

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

These instructions will not work for you if:

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

On this page:
- Before you Start
- Step 1. Identify Administrator
- Step 2. Replace Administrator Password
- Step 3. Put the Internal Directory in First Position
- Step 4. Clean Up

Before you Start

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

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

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

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

Get access to your database

If you're using the embedded H2 database, you can find the files containing your database in `<confluence -home-directory>/database`. See Embedded H2 Database for information on how to connect.

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

Step 1. Identify Administrator

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

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

If there are multiple results, choose one ID/username combination to use for the following steps.
If there are no results, skip down to If No Local Administrator Exists.

> It's important to make sure that the "active" field contains a value of "T". Without this flag, trying to authenticate with this user is a non starter.
>
> To set active to true run the following query replacing "<user_name>" with the username from the previous query
>
> ```
> UPDATE cwd_user
> SET active = 'T'
> WHERE user_name ='<user_name>';
> ```

If No Local Administrator Exists

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

1. Add a new admin user by running:

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


    insert into user_mapping values ('2c9681954172cf560000000000000001',
    'admin', 'admin');
    ```

2. Add new groups by running:

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

3. Add group memberships into cwd_membership:

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

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

Step 2. Replace Administrator Password

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

```
x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1ONxp3Y3
pB37A==
```

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

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

```
update cwd_user set credential =
'x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1
ONxp3Y3pB37A=='
where id=<id from Stage 1>;
```

Step 3. Put the Internal Directory in First Position

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

To put your Internal Directory in first position:

1. Find the directory names and their order:

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

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

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

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

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

5. If necessary, activate the directory:

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

Step 4. Clean Up

To tidy up:

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

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

4. If you created a new user in Stage 2, create a new admin via the UI and delete the admin you created in Stage 2
5. If you followed Stage Three, go to
   ⚙ ▾
   > **General Configuration** > **User Directories** and rearrange your directories so they're correctly configured again.

Edit User Details

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

Edit a user's details

1. Choose
   ⚙ ▾
   > **User management**
2. Do either of the following:

   - Choose **Show all users** to list everyone in the 'confluence-users' or 'users' group

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

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

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

View User: juser

« Back to Users

User	juser
Full Name	Josh User
Email	josh@user.com
Directory	Confluence Internal Directory
Created	Feb 24, 2015 10:13
Last Updated	Feb 24, 2015 10:13
Login	**Last Login:** Jun 22, 2015 10:56 **Current Failed Login Count:** 0
Groups	confluence-users

View Profile   Edit Groups   Edit Details   Set Password   Delete   Disable

---

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.
> See Managing Multiple Directories.

Change a Username

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

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

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

On this page:
* Confluence-managed users
* Users managed in an external directory
* Notes

Confluence-managed users

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

To change a username:

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

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

Users managed in an external directory

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

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

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

Notes

Some important things to note about changing usernames:

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

Managing Site-Wide Permissions and Groups

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

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

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

Related pages:

- Confluence Security Overview and Advisories
- Global Permissions Overview

Confluence Groups for Administrators

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

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

Special groups

There are three special default groups in Confluence:

1. **system-administrators (sys admin)** – The user that sets up a Confluence instance will be the first sys admin; they can also appoint other sys admins by granting them the 'System Administrator' global

permission. They can perform all Confluence administrative functions, including assigning permissions to other users.

2. **confluence-administrators**: This is a group of 'super-users' who can access the Confluence administration screens ('administration console') and perform site-wide administration. Members of this group can also see all spaces in the Confluence site. Any user who is a member of this group has site-wide administration powers, regardless of any other setting. The settings on the global permissions screen do **not** affect the powers allowed to members of this group.

3. **confluence-users**: This is the default group for all new users. Permissions you assign to this group will be assigned to all newly created users.

On this page:

- Special groups
- Anonymous users
- Add or delete groups
- Confluence Administrator permission vs confluence-admin group comparison
- Notes

Related pages:

- Manage Users
- Global Permissions Overview

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

View the comparison table.

Anonymous users

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

Add or delete groups

To add a new group:

1. Choose the **cog icon**

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

You're now ready to start adding users to the group.

To delete a group:

1. Choose the **cog icon**

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

Confluence Administrator permission vs confluence-admin group comparison

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

⚙ ▾

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

˅ Click to view the comparison table

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

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

Atlassian Support Tools	❌	✅
Clustering		

Notes

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

  Here is a summary of how the directory order affects the processing:
  - The order of the directories is the order in which they will be searched for users and groups.
  - Changes to users and groups will be made only in the first directory where the application has permission to make changes.
  See Managing Multiple Directories.

Adding or Removing Users in Groups

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

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

You can edit group membership in two places:

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

On this page:

- Add and remove members via group management
- Edit group membership via user management

Related pages:

- Manage Users
- Global Permissions Overview

Add and remove members via group management

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

To add members to a group:

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

To remove members from a group:

1. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
2. Choose **Groups** in the left-hand panel
3. Choose the group from which you want to remove the user
4. Choose the **Delete user from group** icon next to the user whose group membership you want to remove

Screenshot: Adding members

Group Members: techwriters

Groups Cancel

techwriters

Add Members | connie, jack | 🔍 Enter a comma separated list of user names to add users to this group

Add   Cancel

Edit group membership via user management

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

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

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

     ⚙ ▾

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

Screenshot: Editing a user's groups

Edit User Groups : connie

« Return to view user: connie
Select group memberships for **connie**
☐ confluence-administrators
☑ confluence-users
☐ developers
☐ techwriters
Save · Select All Groups · Deselect All Groups

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

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

See Managing Multiple Directories.

Global Permissions Overview

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

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

On this page:

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

Overview of global permissions

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

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

Comparing the System Administrator permission with the Confluence Administrator permission

Confluence recognises two levels of administrator:

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

- **Confluence Administrator** – Users with this permission can perform most, but not all, of the Confluence administrative functions.

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

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

Administration Screen	Excluded from Confluence Administrator permission
General Configuration	The following functionality is disallowed: - Server Base URL - Public Signup - Connection Timeouts
Further Configuration	The following functionality is disallowed: - Remote API plugin
Security Configuration	The following functionality is disallowed: - External user management - Append wildcards to user and group searches - Anti XSS Mode - Enable Custom Stylesheets for Spaces - Show system information on the 500 page - Maximum RSS Items - XSRF Protection
Plugins	The following functionality is disallowed: - Upgrade - Install - Confluence Upgrade Check
Daily Backup Admin	This function is disallowed entirely.
Mail Servers	This function is disallowed entirely.
User Macros	This function is disallowed entirely.
Attachment Storage	This function is disallowed entirely.
Layouts	This function is disallowed entirely.
Custom HTML	This function is disallowed entirely.
Backup & Restore	This function is disallowed entirely.
Logging and Profiling	This function is disallowed entirely.
Cluster Configuration	This function is disallowed entirely.
Scheduled Jobs	This function is disallowed entirely.

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

Comparing the confluence-administrators group with the administrator permissions

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

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

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

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

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

Updating global permissions

To view the global permissions for a group or user:

1. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
2. Choose **Global Permissions** in the left-hand panel. The 'View Global Permissions' screen appears.

Add or edit group and user permissions as follows:

To add permissions for a group:

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

To add permissions for a specific user:
(Consider adding the user to a group and then assigning the permissions to the group, as described above, instead of assigning permissions to the specific user.)
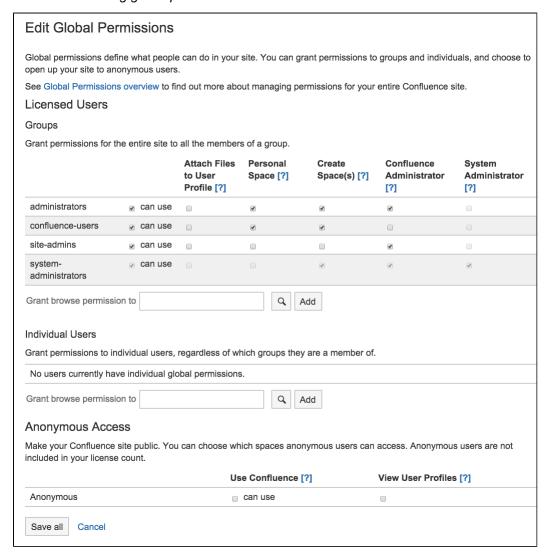
1. First add the user to Confluence, if you have not already done so.
2. Choose **Edit Permissions**. The 'Edit Global Permissions' screen appears.

3. Enter the username in the **Grant browse permission to** box in the 'Individual Users' section. You can search for the username.
4. Choose **Add**.
5. The username will appear in the list and you can now edit its permissions.

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

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

Screenshot: Editing global permissions

Edit Global Permissions

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

See Global Permissions overview to find out more about managing permissions for your entire Confluence site.

Licensed Users

Groups

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

		Attach Files to User Profile [?]	Personal Space [?]	Create Space(s) [?]	Confluence Administrator [?]	System Administrator [?]
administrators	☑ can use	☐	☑	☑	☑	☐
confluence-users	☑ can use	☐	☑	☑	☐	☐
site-admins	☑ can use	☐	☐	☐	☑	☐
system-administrators	☑ can use	☐	☐	☑	☑	☑

Grant browse permission to [   ] 🔍 Add

Individual Users

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

No users currently have individual global permissions.

Grant browse permission to [   ] 🔍 Add

Anonymous Access

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

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

Save all   Cancel

Revoking access for unlicensed users from JIRA Service Desk

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

To revoke access for unlicensed users:

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

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

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

Error messages you may see

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

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

Setting Up Public Access

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

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

> **On this page:**
>
> - Enabling anonymous access to the site
> - Disabling anonymous access to the site
> - Granting public access to a space
> - Notes
>
> **Related pages:**
>
> - Configuring Captcha for Spam Prevention
> - Add and Invite Users
> - Global Permissions Overview

Enabling anonymous access to the site

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

To enable anonymous access to your site:

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

Disabling anonymous access to the site

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

Granting public access to a space

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

- The site-wide 'can use' permission, as described above.
- The relevant space permissions. If you want a space to be publicly accessible, the anonymous user

must have at least the 'View Space' permission. To set space permissions, choose **Browse** > **Space Admin** > **Permissions**.

Notes

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

Configuring User Directories

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

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

On this page:

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

Related pages

- Configuring the Internal Directory
- Connecting to an LDAP Directory
- Connecting to an Internal Directory with LDAP Authentication
- Connecting to Crowd or JIRA for User Management
- Managing Multiple Directories
- Managing Nested Groups
- Synchronising Data from External Directories
- Diagrams of Possible Configurations for User Management
- User Management Limitations and Recommendations
- Requesting Support for External User Management
- Disabling the Built-In User Management

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

Configuring User Directories in Confluence

To configure your Confluence user directories:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Click '**User Directories**' in the left-hand panel.

Connecting to a Directory

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

- Confluence's internal directory. See Configuring the Internal Directory.
- Microsoft Active Directory. See Connecting to an LDAP Directory.
- Various other LDAP directory servers. See Connecting to an LDAP Directory.
- An LDAP directory for delegated authentication. See Connecting to an Internal Directory with LDAP Authentication.
- Atlassian Crowd or JIRA 4.3 or later. See Connecting to Crowd or JIRA for User Management.

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

Updating Directories

Limitations when Editing Directories

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

This limitation applies to all directory types. For example:

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

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

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

Recommendations

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

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

Enabling, Disabling and Removing Directories

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

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

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.

On this page:

- Overview
- Diagram of Possible Configuration

Related pages:

- Configuring User Directories
- How to Reenable the Internal Directory (Knowledge base article)

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

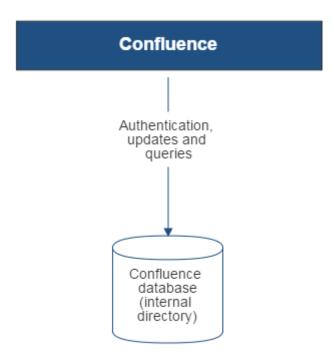**Diagram of Possible Configuration**

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

Connecting to an LDAP Directory

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

Overview

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

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

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

When to use this option: Connecting to an LDAP

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

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

Connecting to an LDAP Directory in Confluence

To connect Confluence to an LDAP directory:

1. Choose the **cog icon**

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

   For details see Managing Multiple Directories.

Server Settings

Setting	Description
Name	Enter a meaningful name to help you identify the LDAP directory server. Examples: • `Example Company Staff Directory` • `Example Company Corporate LDAP`
Directory Type	Select the type of LDAP directory that you will connect to. If you are adding a new LDAP connection, the value you select here will determine the default values for many of the options on the rest of screen. Examples: • `Microsoft Active Directory` • `OpenDS` • And more.
Hostname	The host name of your directory server. Examples: • `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 if the connection to the directory server is an SSL (Secure Sockets Layer) connection. Note that you will need to configure an SSL certificate in order to use this setting.

Username	The distinguished name of the user that the application will use when connecting to the directory server. Examples:
	`cn=administrator,cn=users,dc=ad,dc=example,dc=com` `cn=user,dc=domain,dc=name` `user@domain.name` ℹ️ Ensure that this is an administrator user for the LDAP engine. For example, in Active Directory the user will need to be a member of the built-in Administrators group. The specific privileges for the LDAP user that is used to connect to LDAP are bind and read (user info, group info, group membership, update sequence number, deleted objects). The need for admin privileges is because a normal user can't access uSNChanged attribute and deleted objects container, causing incremental sync to fail silently, this was reported in a bug here C WD-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.
Additional User DN	This value is used in addition to the base DN when searching and loading users. If no value is supplied, the subtree search will start from the base DN. Example: `ou=Users`
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`

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. ⓘ Note for Confluence users: Users from LDAP are added to groups maintained in Confluence's internal directory the first time they log in. This is only done once per user. There is a known issue with Read Only, with Local Groups in Confluence that may apply to you. See **CONF-28621** - User Loses all Local Group Memberships If LDAP Sync is Unable to find the User, but the User appears again in subsequent syncs ⎢OPEN⎥
Read/Write	LDAP users, groups and memberships are retrieved from your directory server. When you modify a user, group or membership via the application administration screens, the changes will be applied directly to your LDAP directory server. Please ensure that the LDAP user specified for the application has modification permissions on your LDAP directory server.

Adding Users to Groups Automatically

Setting	Description
Default Group Memberships	*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 behaviour allows users to be removed from automatically-added groups. In Confluence 3.5 and 3.5.1, they would be re-added upon next login. Please note that there is no validation of the group names. If you mis-type the group name, authorisation failures will result – users will not be able to access the applications or functionality based on the intended group name. Examples: • `confluence-users` • `confluence-users,jira-administrators,jira-core-users`

Advanced Settings

Setting	Description
Enable Nested Groups	Enable or disable support for nested groups. Some directory servers allow you to define a group as a member of another group. Groups in such a structure are called 'nested groups'. If you are using groups to manage permissions, you can create nested groups to allow inheritance of permissions from one group to its sub-groups.
Manage User Status Locally	If true, you can activate and deactivate users in Crowd independent of their status in the directory server.
Filter out expired users	If true, user accounts marked as expired in ActiveDirectory will be automatically removed. For cached directories, the removal of a user will occur during the first synchronisation after the account's expiration date.

Use Paged Results	Enable or disable the use of the LDAP control extension for simple paging of search results. If paging is enabled, the search will retrieve sets of data rather than all of the search results at once. Enter the desired page size – that is, the maximum number of search results to be returned per page when paged results are enabled. The default is 1000 results.
Follow Referrals	Choose whether to allow the directory server to redirect requests to other servers. This option uses the node referral (JNDI lookup `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 compare DNs to determine if they are equal. • If this checkbox is selected, the application will do a direct, case-insensitive, string comparison. This is the default and recommended setting for Active Directory, because Active Directory guarantees the format of DNs. • If this checkbox is not selected, the application will parse the DN and then check the parsed version.
Enable Incremental Synchronisation	Enable incremental synchronisation if you only want changes since the last synchronisation to be queried when synchronising a directory. Please be aware that when using this option, the user account configured for synchronisation must have read access to: • The `uSNChanged` attribute of all users and groups in the directory that need to be synchronised. • The objects and attributes in the Active Directory deleted objects container (see Microsoft's Knowledge Base Article No. 892806 for details). If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory".
Synchronisation Interval (minutes)	Synchronisation is the process by which the application updates its internal store of user data to agree with the data on the directory server. The application will send a request to your directory server every x minutes, where 'x' is the number specified here. The default value is 60 minutes.
Read Timeout (seconds)	The time, in seconds, to wait for a response to be received. If there is no response within the specified time period, the read attempt will be aborted. A value of 0 (zero) means there is no limit. The default value is 120 seconds.
Search Timeout (seconds)	The time, in seconds, to wait for a response from a search operation. A value of 0 (zero) means there is no limit. The default value is 60 seconds.
Connection Timeout (seconds)	This setting affects two actions. The default value is 0. • The time to wait when getting a connection from the connection pool. A value of 0 (zero) means there is no limit, so wait indefinitely. • The time, in seconds, to wait when opening new server connections. A value of 0 (zero) means that the TCP network timeout will be used, which may be several minutes.

User Schema Settings

Setting	Description
User Object Class	This is the name of the class used for the LDAP user object. Example: - `user`
User Object Filter	The filter to use when searching user objects. Example: - `(&(objectCategory=Person)(sAMAccountName=*))` More examples can be found here and here.
User Name Attribute	The attribute field to use when loading the username. Examples: - `cn` - `sAMAccountName` NB: In Active Directory, the 'sAMAccountName' is the 'User Logon Name (pre-Windows 2000)' field. The User Logon Name field is referenced by 'cn'.
User Name RDN Attribute	The RDN (relative distinguished name) to use when loading the username. The DN for each LDAP entry is composed of two parts: the RDN and the location within the LDAP directory where the record resides. The RDN is the portion of your DN that is not related to the directory tree structure. Example: - `cn`
User First Name Attribute	The attribute field to use when loading the user's first name. Example: - `givenName`
User Last Name Attribute	The attribute field to use when loading the user's last name. Example: - `sn`
User Display Name Attribute	The attribute field to use when loading the user's full name. Example: - `displayName`
User Email Attribute	The attribute field to use when loading the user's email address. Example: - `mail`
User Password Attribute	The attribute field to use when loading a user's password. Example: - `unicodePwd`
User Unique I D Attribute	The attribute used as a unique immutable identifier for user objects. This is used to track username changes and is optional. If this attribute is not set (or is set to an invalid value), user renames will not be detected — they will be interpreted as a user deletion then a new user addition. This should normally point to a UUID value. Standards-compliant LDAP servers will implement this as 'entryUUID' according to RFC 4530. This setting exists because it is known under different names on some servers, e.g. 'objectGUID' in Microsoft Active Directory.

Group Schema Settings

Setting	Description

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

Membership Schema Settings

Setting	Description
Group Members Attribute	The attribute field to use when loading the group's members. Example: • `member`
User Membership Attribute	The attribute field to use when loading the user's groups. Example: • `memberOf`
Use the User Membership Attribute, when finding the user's group membership	Check this if your directory server supports the group membership attribute on the user. (By default, this is the `'memberOf'` attribute.) • If this checkbox is selected, your application will use the group membership attribute on the user when **retrieving the list of groups to which a given user belongs**. This will result in a more efficient retrieval. • If this checkbox is not selected, your application will use the members attribute on the group (`'member'` by default) for the search. • If the **Enable Nested Groups** checkbox is seleced, 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.

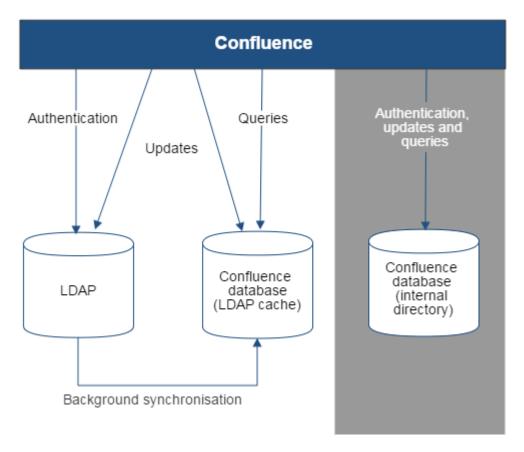***Diagrams of Some Possible Configurations***

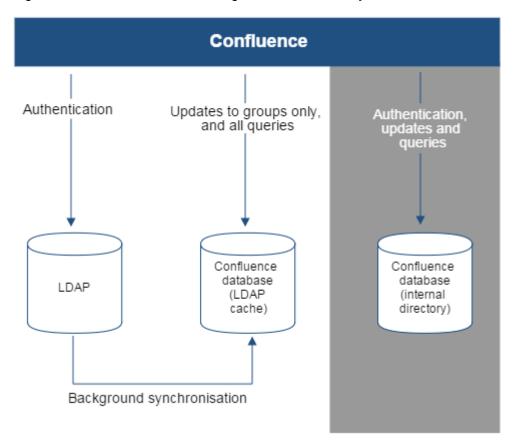*Diagram above: Confluence connecting to an LDAP directory.*

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.

To configure your LDAP connection pool:

1. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
2. Click '**User Directories**' in the left-hand panel.
3. Click '**LDAP Connection Pool Configuration**' in the 'Additional Configuration' section.

Setting	Description	Default Value
Initial Pool Size	The number of LDAP connections created when initially connecting to the pool.	`1`
Preferred Pool Size	The optimal pool size. LDAP will remove idle connections when the number of connections grows larger than this value. A value of `0` (zero) means that there is no preferred size, so the number of idle connections is unlimited.	`10`
Maximum Pool Size	The maximum number of connections. When the number of connections reaches this value, LDAP will refuse further connections. As a result, requests made by an application to the LDAP directory server will be blocked. A value of `0` (zero) means that the number of connections is unlimited.	`0`
Pool Timeout (seconds)	The length of time, in seconds, that a connection may remain idle before being removed from the pool. When the application is finished with a pooled connection, the connection is marked as idle, waiting to be reused. A value of `0` (zero) means that the idle time is unlimited, so connections will never be timed out.	`30`
Pool Protocol	Only these protocol types will be allowed to connect to the LDAP directory server. If you want to allow multiple protocols, enter the values separated by a space. Valid values are: - `plain` - `ssl`	`plain ssl` (Both plain and ssl)
Pool Authentication	Only these authentication types will be allowed to connect to the LDAP directory server. If you want to allow multiple authentication types, enter the values separated by a space. See RFC 2829 for details of LDAP authentication methods. Valid values are: - `none` - `simple` - `DIGEST-MD5`	`simple`

Notes:

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

Configuring an SSL Connection to Active Directory

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

> **On this page:**
> - Prerequisites
> - Step 1. Install the Active Directory Certificate Services
> - Step 2. Obtain the Server Certificate
> - Step 3. Import the Server Certificate

> There's a Confluence SSL plugin that facilitates this process.

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

Prerequisites

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

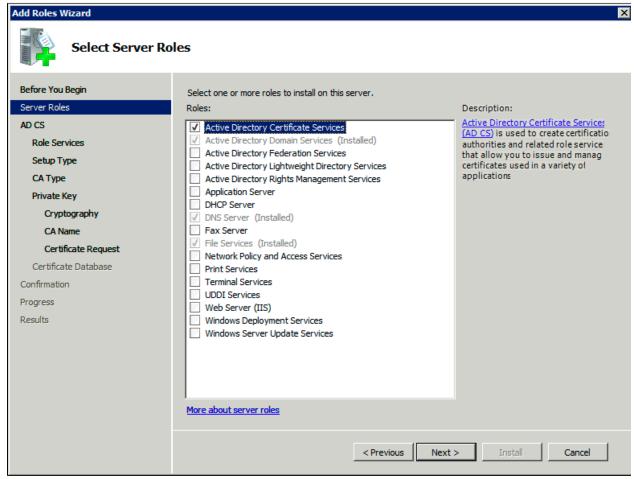| Required Component	Description
Internet Information Services (IIS)	This is required before you can install Windows Certificate Services.
Windows Certificate Services	This installs a certification authority (CA) which is used to issue certificates. Step 1, below, explains this process.
Windows 2000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

Step 1. Install the Active Directory Certificate Services

If Certificate Services are already installed, skip to step 2, below. The screenshots below are from Server 2008, but the process is similar for Server 2000 and 2003.
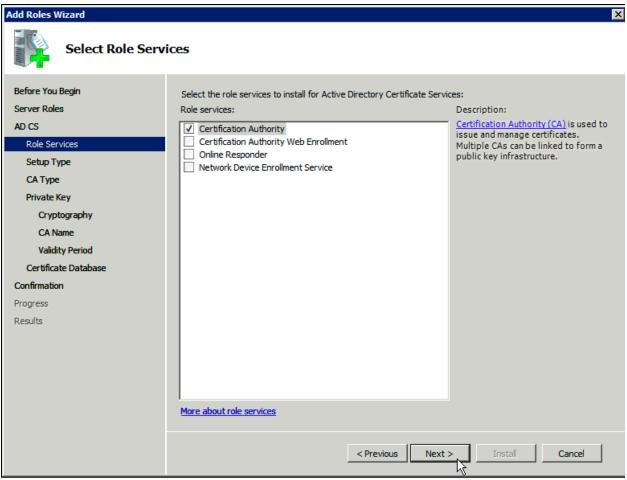
1. Log in to your Active Directory server as an administrator.
2. Click **Start**, point to **Administrative Tools**, and then click **Server Manager**.
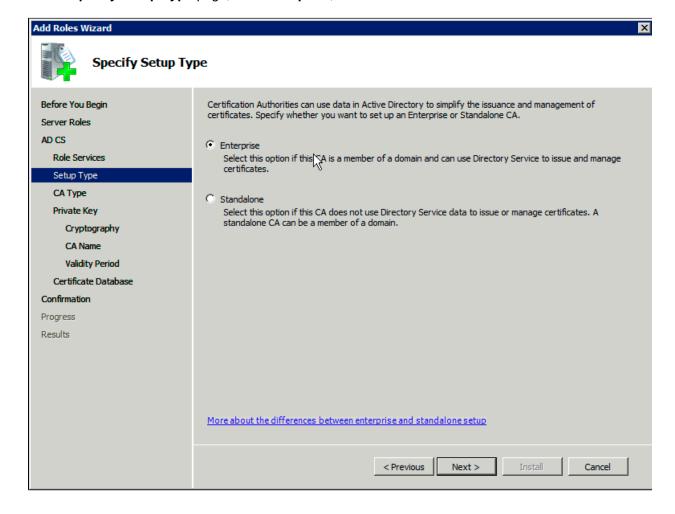3. In the **Roles Summary** section, click **Add Roles**.

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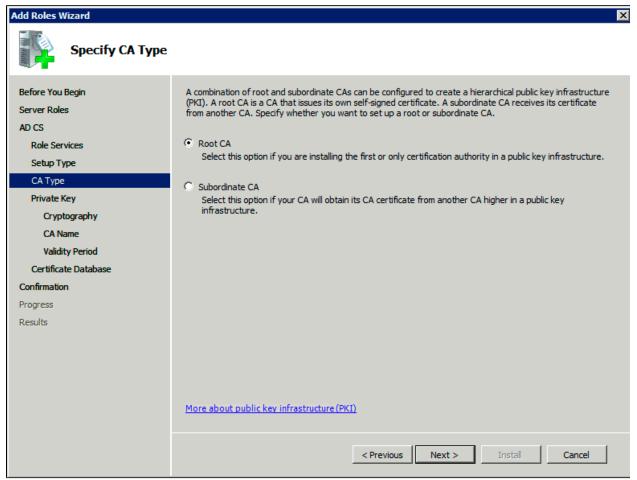

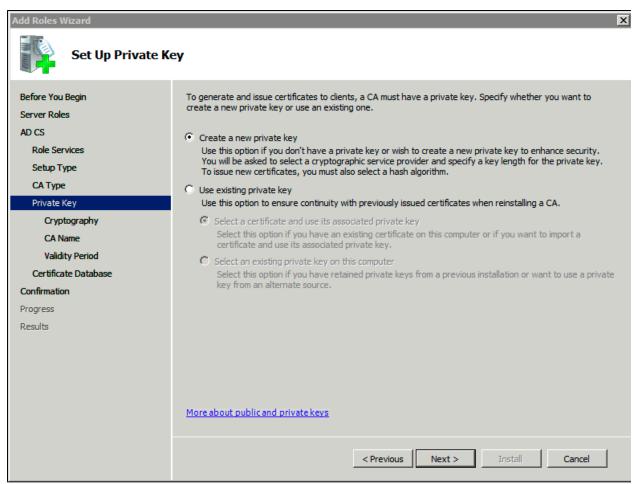5. On the **Select Role Services** page, select the **Certification Authority** check box, and then click **Next**.

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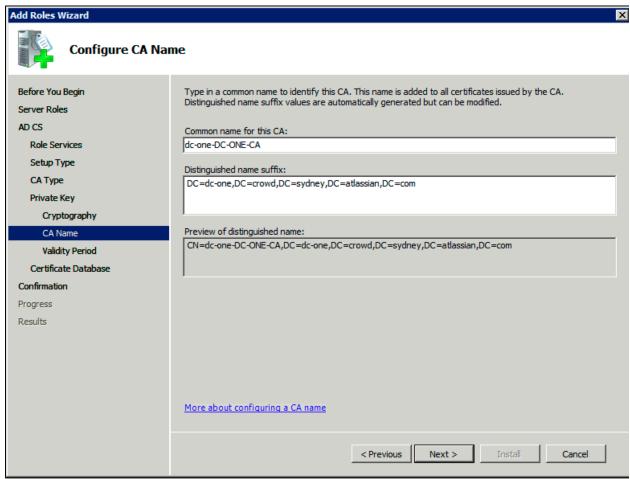
7. On the **Specify CA Type** page, click **Root CA**, and then click **Next**.

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

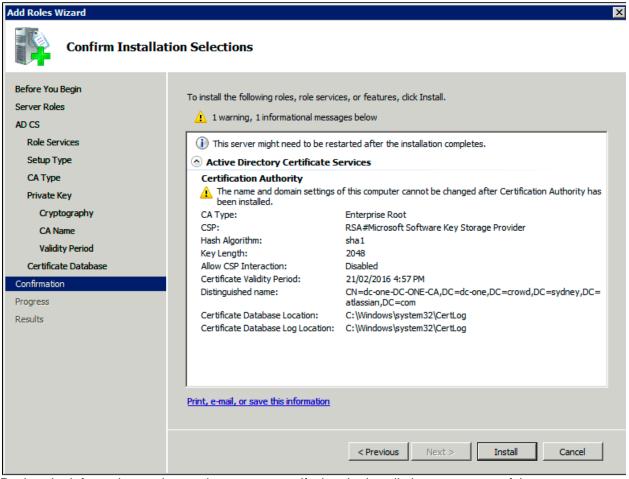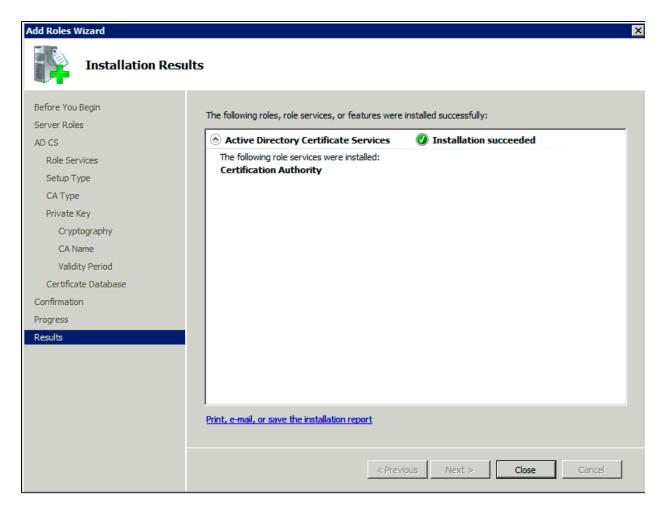The steps above describe how to install the certification authority (CA) on your Microsoft Active Directory server. Next, you will need to add the Microsoft Active Directory server's SSL certificate to the list of accepted certificates used by the JDK that runs your application server.

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

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

Step 3. Import the Server Certificate

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

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

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

pick up the changes.
Windows

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

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

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

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

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

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

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

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

```
cd $JAVA_HOME
```

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

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

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

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

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

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

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

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

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

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

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

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

RELATED TOPICS

Connecting to an LDAP Directory
Configuring User Directories

Connecting to an Internal Directory with LDAP Authentication

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

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

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.

Connecting Confluence to an Internal Directory with LDAP Authentication

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

1. Choose the **cog icon**

   ⚙▾

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

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

Server Settings

Setting	Description
Name	A descriptive name that will help you to identify the directory. Examples: • `Internal directory with LDAP Authentication` • `Corporate LDAP for Authentication Only`

Directory Type	Select the type of LDAP directory that you will connect to. If you are adding a new LDAP connection, the value you select here will determine the default values for some of the options on the rest of screen. Examples: • `Microsoft Active Directory` • `OpenDS` • And more.
Hostname	The host name of your directory server. Examples: • `ad.example.com` • `ldap.example.com` • `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.
Username	The distinguished name of the user that the application will use when connecting to the directory server. Examples: • `cn=administrator,cn=users,dc=ad,dc=example,dc=com` • `cn=user,dc=domain,dc=name` • `user@domain.name`
Password	The password of the user specified above.

Copying Users on Login

Setting	Description
Copy User on Login	This option affects what will happen when a user attempts to log in. If this box is checked, the user will be created automatically in the internal directory that is using LDAP for authentication when the user first logs in and their details will be synchronised on each subsequent log in. If this box is not checked, the user's login will fail if 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 • Synchronise Group Memberships • User Schema Settings (described in a separate section below)

Default Group Memberships	This field appears if you check the **Copy User on Login** box. If you would like users to be automatically added to a group or groups, enter the group name(s) here. To specify more than one group, separate the group names with commas. Each time a user logs in, their group memberships will be checked. If the user does not belong to the specified group(s), their username will be added to the group(s). If a group does not yet exist, it will be added to the internal directory that is using LDAP for authentication. Please note that there is no validation of the group names. If you mis-type the group name, authorisation failures will result – users will not be able to access the applications or functionality based on the intended group name. Examples: • `confluence-users` • `bamboo-users,jira-administrators,jira-core-users`
Synchronise Group Memberships	This field appears if you select the **Copy User on Login** checkbox. If this box is checked, group memberships specified on your LDAP server will be synchronised with the internal directory each time the user logs in. If you check this box the following additional fields will appear on the screen, both described in more detail below: • Group Schema Settings (described in a separate section below) • Membership Schema Settings (described in a separate section below)

Note: 'Copy Users on Login' must be enabled if you want to be able to change usernames.

Schema Settings

Setting	Description
Base DN	The root distinguished name (DN) to use when running queries against the directory server. Examples: • `o=example,c=com` • `cn=users,dc=ad,dc=example,dc=com` • For Microsoft Active Directory, specify the base DN in the following format: `dc=domain1,dc=local`. You will need to replace the `domain1` and `local` for your specific configuration. Microsoft Server provides a tool called `ldp.exe` which is useful for finding out and configuring the the LDAP structure of your server.
User Name Attribute	The attribute field to use when loading the username. Examples: • `cn` • `sAMAccountName`

Advanced Settings

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

Follow Referrals	Choose whether to allow the directory server to redirect requests to other servers. This option uses the node referral (JNDI lookup `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.

User Schema Settings

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

Setting	Description
Additional User DN	This value is used in addition to the base DN when searching and loading users. If no value is supplied, the subtree search will start from the base DN. Example: • `ou=Users`
User Object Class	This is the name of the class used for the LDAP user object. Example: • `user`
User Object Filter	The filter to use when searching user objects. Example: • `(&(objectCategory=Person)(sAMAccountName=*))`
User Name RDN Attribute	The RDN (relative distinguished name) to use when loading the username. The DN for each LDAP entry is composed of two parts: the RDN and the location within the LDAP directory where the record resides. The RDN is the portion of your DN that is not related to the directory tree structure. Example: • `cn`
User First Name Attribute	The attribute field to use when loading the user's first name. Example: • `givenName`
User Last Name Attribute	The attribute field to use when loading the user's last name. Example: • `sn`
User Display Name Attribute	The attribute field to use when loading the user's full name. Example: • `displayName`
User Email Attribute	The attribute field to use when loading the user's email address. Example: • `mail`

Group Schema Settings

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

Setting	Description
Additional Group DN	This value is used in addition to the base DN when searching and loading groups. If no value is supplied, the subtree search will start from the base DN. Example: • `ou=Groups`
Group Object Class	This is the name of the class used for the LDAP group object. Examples: • `groupOfUniqueNames` • `group`

Group Object Filter	The filter to use when searching group objects. Example: • `(objectCategory=Group)`
Group Name Attribute	The attribute field to use when loading the group's name. Example: • `cn`
Group Description Attribute	The attribute field to use when loading the group's description. Example: • `description`

Membership Schema Settings

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

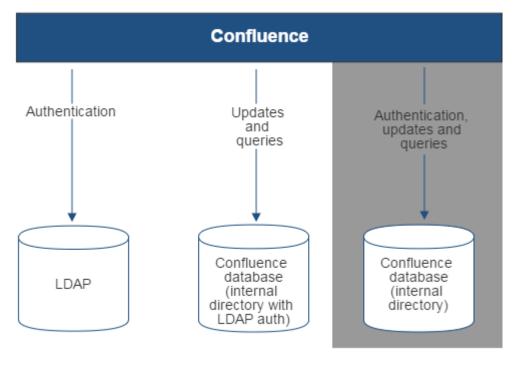| Setting	Description
Group Members Attribute	The attribute field to use when loading the group's members. Example: • `member`
User Membership Attribute	The attribute field to use when loading the user's groups. Example: • `memberOf`
Use the User Membership Attribute, when finding the user's group membership	Check this box if your directory server supports the group membership attribute on the user. (By default, this is the 'memberOf' attribute.) • If this box is checked, your application will use the group membership attribute on the user when **retrieving the members of a given group**. This will result in a more efficient retrieval. • If this box is not checked, your application will use the members attribute on the group ('member' by default) for the search.

Diagrams of Possible Configurations

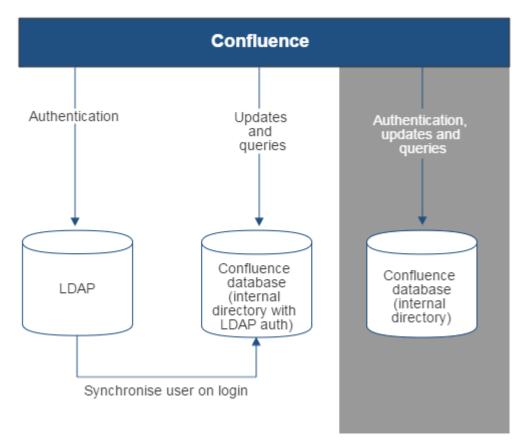*Diagram above: Confluence connecting to an LDAP directory for authentication only.*

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

Connecting to Crowd or JIRA for User Management

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

Connecting Confluence to Crowd for User Management

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

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

On this page:

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

Related pages:

- Configuring User Directories

To connect Confluence to Crowd:

1. Go to your **Crowd Administration Console** and define the Confluence application to Crowd. See the Crowd documentation: Adding an Application.

2. Choose the **cog icon**

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

   For details see Managing Multiple Directories.
7. If required, configure Confluence to use Crowd for single sign-on (SSO) too. See the Crowd documentation: Integrating Crowd with Atlassian Confluence.

Crowd Settings in Confluence

Setting	Description
Name	A meaningful name that will help you to identify this Crowd server amongst your list of directory servers. Examples: • `Crowd Server` • `Example Company Crowd`
Server URL	The web address of your Crowd console server. Examples: • `http://www.example.com:8095/crowd/` • `http://crowd.example.com`
Application Name	The name of your application, as recognized by your Crowd server. Note that you will need to define the application in Crowd too, using the Crowd administration Console. See the Crowd documentation on adding an application.
Application Password	The password which the application will use when it authenticates against the Crowd framework as a client. This must be the same as the password you have registered in Crowd for this application. See the Crowd documentation on adding an application.

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

CONF-33979 - New JIRA/Crowd password not saved after test **OPEN**   when configuring Confluence to use JIRA/Crowd as a external user directory.

Crowd Permissions

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

Advanced Crowd Settings

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

Connecting Confluence to JIRA applications for User Management

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

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

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

To connect Confluence to a JIRA application:

1. In your JIRA application go to

   > **User Management** > **JIRA User Server**.
   (For JIRA 6.4 and earlier go to your JIRA administration screen then **Users** > **JIRA User Server**)
   - Click **Add Application**.
   - Enter the **application name** and **password** that Confluence will use when accessing JIRA.
   - Enter the **IP address** or addresses of your Confluence server. Valid values are:
     - A full IP address, e.g. `192.168.10.12`.
     - A wildcard IP range, using CIDR notation, e.g. `192.168.10.1/16`. For more information, see the introduction to CIDR notation on Wikipedia and RFC 4632.
   - **Save** the new application.
2. Set up the JIRA user directory in Confluence:
   - Choose the **cog icon**

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

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

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

JIRA Settings in Confluence

Setting	Description
Name	A meaningful name that will help you to identify this JIRA server in the list of directory servers. Examples: - `JIRA Service Desk Server` - `My Company JIRA`
Server URL	The web address of your JIRA server. Examples: - `http://www.example.com:8080` - `http://jira.example.com`
Application Name	The name used by your application when accessing the JIRA server that acts as user manager. Note that you will also need to define your application to that JIRA server, via the '**Other Applications**' option in the 'Users, Groups & Roles' section of the 'Administration' menu.
Application Password	The password used by your application when accessing the JIRA server that acts as user manager.

JIRA Permissions

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

Advanced JIRA Settings

Setting	Description
Enable Nested Groups	Enable or disable support for nested groups. Before enabling nested groups, please check to see if nested groups are enabled on the JIRA server that is acting as user manager. When nested groups are enabled, you can define a group as a member of another group. If you are using groups to manage permissions, you can create nested groups to allow inheritance of permissions from one group to its sub-groups.

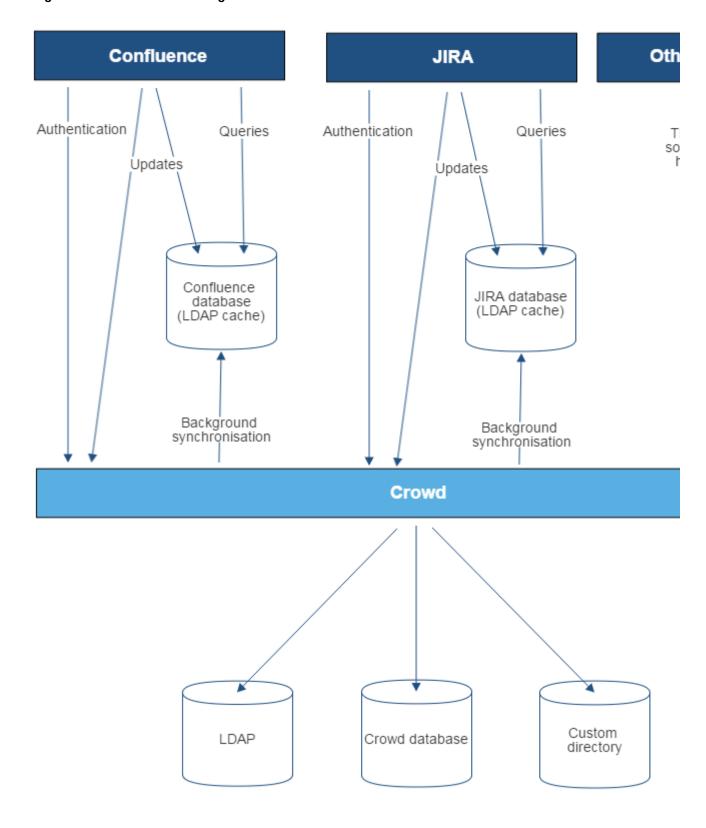| Synchronisation Interval (minutes)	Synchronisation is the process by which the application updates its internal store of user data to agree with the data on the directory server. The application will send a request to your directory server every x minutes, where 'x' is the number specified here. The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

Diagrams of Some Possible Configurations

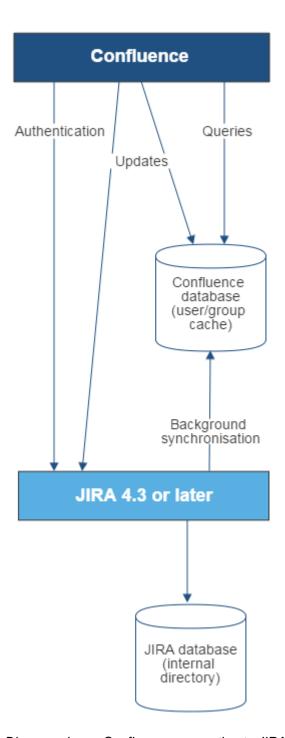*Diagram above: Confluence, JIRA and other applications connecting to Crowd for user management.*

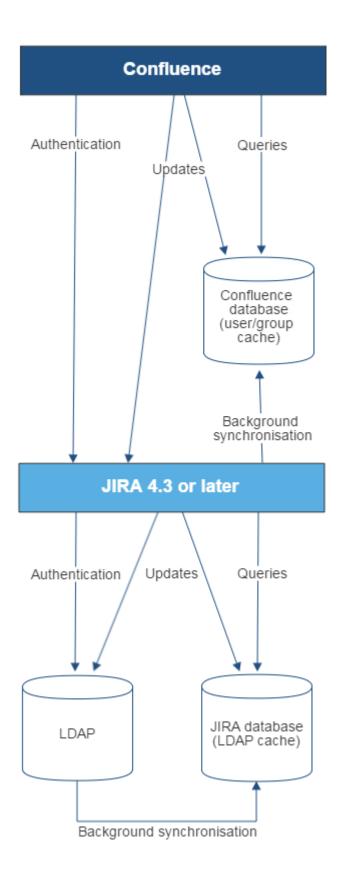*Diagram above: Confluence connecting to JIRA for user management.*

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

Troubleshooting

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

Error	Message	Cause

error.jirabaseurl.connection.refused	Connection refused. Check if an instance of JIRA 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.

On this page:
• Option 1 – Manually Recreate Users and Groups in Confluence
• Option 2 – Transfer Crowd/JIRA application Users and Groups to the Confluence Database

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

Option 1 – Manually Recreate Users and Groups in Confluence

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

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

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

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

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

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

Step 1. Create Backups

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

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

Step 2. Replace Confluence User Management

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

Move Groups to Confluence

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

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

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

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

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

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

Move Users to Confluence

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

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

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

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

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

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

Delete the External Application directory

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

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

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

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

- List the directories and their order using

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

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

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

c. Check that the internal directory is enabled.

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

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

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

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

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

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

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

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

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

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

Reset passwords

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

Managing Multiple Directories

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

On this page:
- Overview
- Configuring the Directory Order
- Effect of Directory Order
  - Login
  - Permissions
  - Updating Users and groups

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

Overview

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

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

Configuring the Directory Order

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

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

Notes:

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

Effect of Directory Order

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

Login

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

Permissions

Aggregating membership (default)

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

Example:

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

For administrators upgrading to Confluence 5.7 or later:

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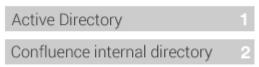
∨ Here's an example scenario...

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

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

Confluence User Directories

| Active Directory | 1 |
| Confluence internal directory | 2 |

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

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

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

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

To Confluence his group membership looks like this:

Confluence group membership
confluence-users
developers
sydney

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

Non-aggregating membership

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

⌄ Turning on non-aggregating membership...

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

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

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

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

Example:

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

Updating Users and groups

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

Example 1:

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

Example 2:

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

RELATED TOPICS

Configuring User Directories

- Configuring the Internal Directory
- Connecting to an LDAP Directory
- Connecting to an Internal Directory with LDAP Authentication
- Connecting to Crowd or JIRA for User Management
- Managing Multiple Directories
- Managing Nested Groups
- Synchronising Data from External Directories

- Diagrams of Possible Configurations for User Management
- User Management Limitations and Recommendations
- Requesting Support for External User Management
- Disabling the Built-In User Management

Managing Nested Groups

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

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

Enabling Nested Groups

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

Notes:

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

On this page:

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

Related pages:

- Configuring User Directories

Effect of Nested Groups

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

Login

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

Permissions

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

Viewing Lists of Group Members

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

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

Adding and Updating Group Memberships

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

If you try to remove a user from a flattened list, the following will happen:

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

Examples

Example 1: User is Member of Sub-Group

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

- `staff`
- `marketing`

Memberships:

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

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

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

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

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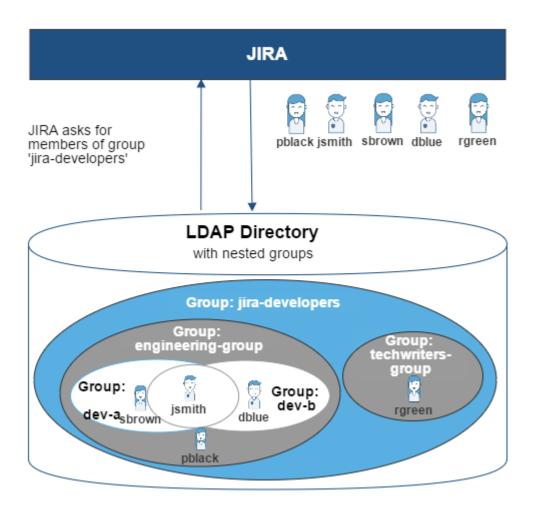- Add the '**techwriters-group**' as a sub-group of '**jira-developers**'.

Group memberships are now:

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

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

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

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

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

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

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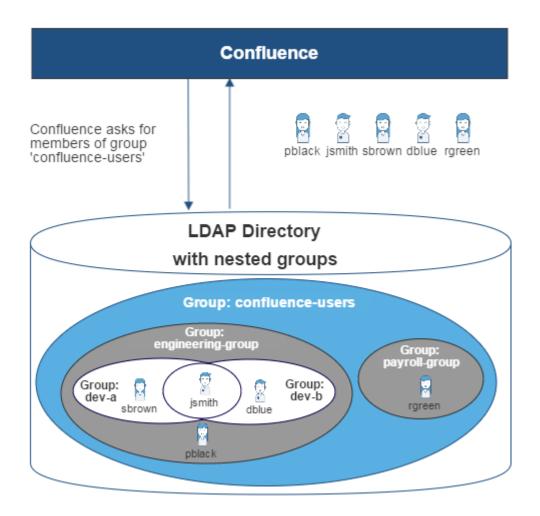- Add the '**payroll-group**' as a sub-group of '**confluence-users**'.

Group memberships are now:

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

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

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

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

Notes

- **Possible impact on performance.** Enabling nested groups may result in slower user searches.

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

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

Synchronising Data from External Directories

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

On this page:

- Affected Directory Types
- How it Works
- Finding the Time Taken to Synchronise
- Manually Synchronising the Cache
- Configuring the Synchronisation Interval

Related pages:

- Configuring User Directories

Affected Directory Types

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

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

Finding the Time Taken to Synchronise

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

Manually Synchronising the Cache

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

Screen snippet: User directories, showing information about synchronisation

Configuring the Synchronisation Interval

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

You can set the '**Synchronisation Interval**' on the directory configuration screen. The synchronisation

interval is the period of time to wait between requests for updates from the directory server.

The length you choose for your synchronisation interval depends on:

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

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

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

Diagrams of Possible Configurations for User Management

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

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

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

On this page:

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

Related pages:

- Configuring User Directories
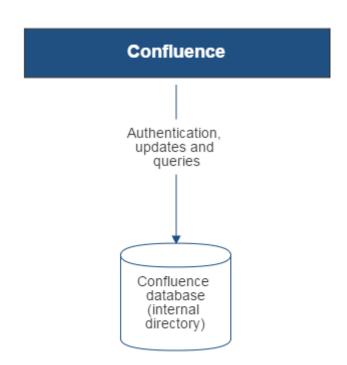
Confluence Internal Directory

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

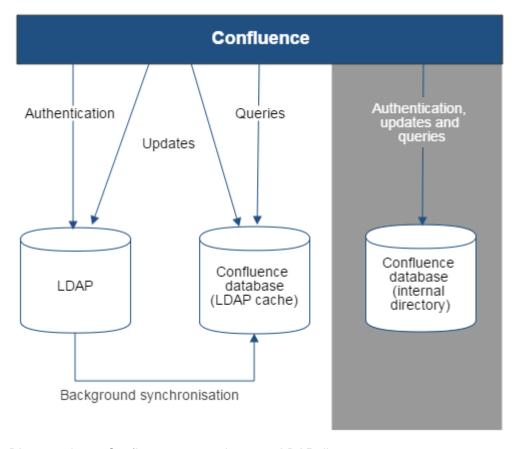**Confluence with Read/Write Connection to LDAP**

Diagram above: Confluence connecting to an LDAP directory.

Confluence with Read-Only Connection to LDAP, with Local Groups
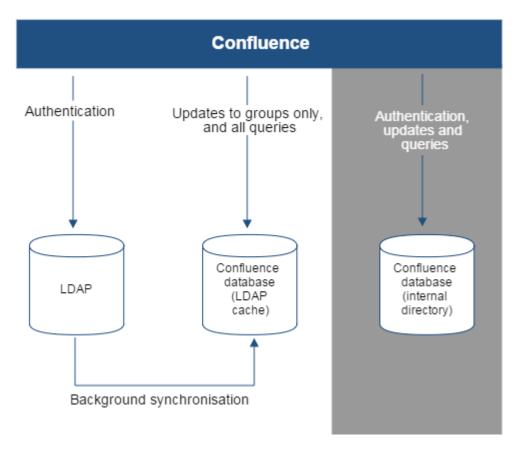
Diagram above: Confluence connecting to an LDAP directory with permissions set to read only and local groups.
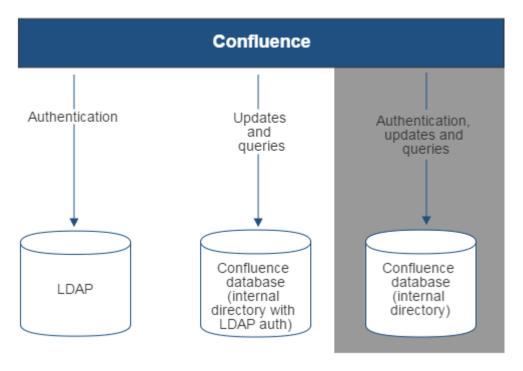
Confluence Internal Directory with LDAP Authentication

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

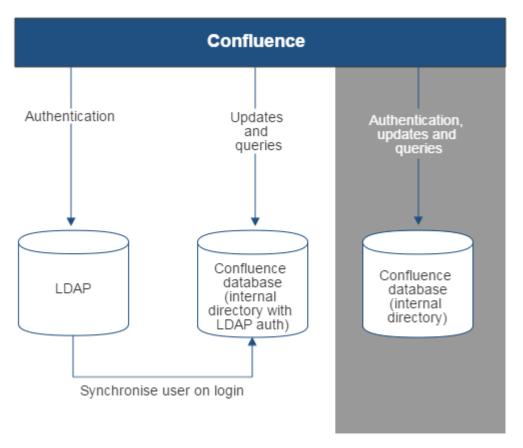**Confluence with LDAP Authentication, Copy Users on First Login**

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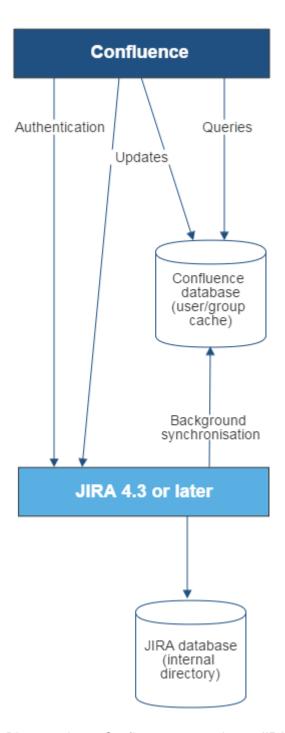
Confluence Connecting to JIRA

Diagram above: Confluence connecting to JIRA for user management.

Confluence Connecting to JIRA and JIRA Connecting to LDAP

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

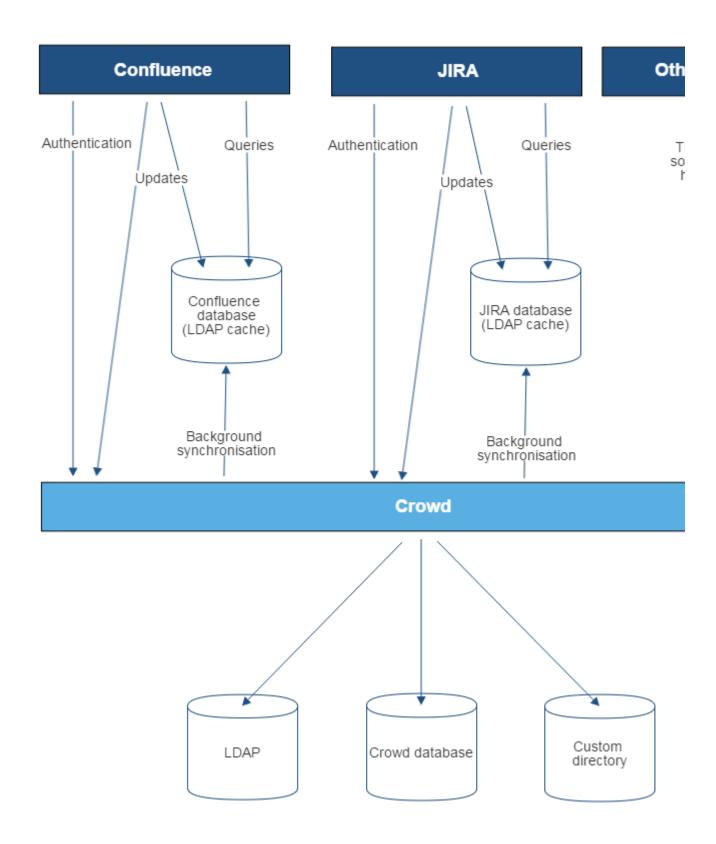**Confluence and JIRA Connecting to Crowd**

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

User Management Document 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.

Recommendations for Connecting to LDAP

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

Optimal Number of Users and Groups in your LDAP Directory

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

This recommendation affects connections to LDAP directories:

- Microsoft Active Directory
- All other LDAP directory servers

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

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

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

Your environment	Recommendation

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

Our Test Results

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

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

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

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

Redundant LDAP is Not Supported

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

Specific Notes for Connecting to Active Directory

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

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

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

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

Recommendations for Connecting to JIRA for User Management

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

Single Sign-On Across Multiple Applications is Not Supported

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

Custom Application Connectors are Not Supported

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

Custom Directories are Not Supported

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

Load on your JIRA instance

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

JIRA Cloud applications not supported

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

Recommendations

Your environment	Recommendation
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. • 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.	Your environment meets the optimal requirements for using a JIRA application for user management.

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** under Confluence Administration

2. Click '**User Directories**' in the left-hand panel.
3. **Edit** the relevant directory.
4. Click '**Test Settings**'.
5. The results of the test will appear at the top of the screen.

Please refer to our 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.pr operties`.
  Change this section:

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

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

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

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

Disabling the Built-In User Management

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

You might disable Confluence's internal user management:

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

To disable management of users and groups within Confluence:

⚙ ▾

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

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

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

Managing Add-ons or Plugins

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

There are two main types of add-ons:

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

For information about developing your own add-ons for Confluence, see the Confluence Developer documentation.

About the Universal Plugin Manager

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

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

The UPM allows you to:

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

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

See the Universal Plugin Manager documentation for more information on using the UPM.

Writing User Macros

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

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

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

On this page:

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

Related Pages:

- User Macro Module (Developer documentation)

Create a User Macro

To add a new user macro:

1. Go to

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

Macro details field	Description
Macro name	This is the name of the macro, as it appears in the code.
Visibility	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.
Description	This is the description that will appear in the macro browser. The macro browser's search will pick up matches in both the title and description.

Categories	Select one or more macro browser categories for your macro to appear in.
Icon URL	Enter an absolute URL (for example `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.
Documentation URL	If you have documentation for your macro, enter the URL here.
Macro Body Processing	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 `$body` variable. Options for processing the macro body include: • **No macro body** Select this option if your macro does not have a body. • **Escaped** Confluence will add escape characters to the HTML markup in the macro body. Use this if you want to show actual HTML markup in the rendered page. For example, if the body is `<b>Hello World</b>` it will render as `<b>Hello World</b>`. • **Unrendered** HTML in the body will be processed within the template before being output. Ensure that HTML is ultimately output by the template. • **Rendered** Confluence will recognise HTML in the macro body, and render it appropriately. For example, if the body is `<b>Hello World</b>` it will render as **Hello World.**
Template	This is where you write the code that determines what the macro should do. • Use HTML and Confluence-specific XML elements in the macro template. Details of Confluence's storage format are in Confluence Storage Format. • You can use the Velocity templating language. Here is more information on the Velocity project. • If your macro has a body, your template can refer to the macro body text by specifying '`$body`'. • Each parameter variable you use must have a matching metadata definition. Use `@param` to define metadata for your macro parameters. • When using the information passed using parameters, refer to your parameters as $paramXXX where 'XXX' is the parameter name that you specifed in the `@param` metadata definition. • Use `@noparams` if your macro does not accept parameters. See User Macro Template Syntax for more information and examples.

> **Do you need a plugin instead?**
> If you want to distribute your user macro as a plugin, please refer to the developer's guide to the User Macro plugin module. If you want to create more complex, programmatic macros in Confluence, you may need to write a Macro plugin.

Edit a user macro

To edit a user macro:

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

4. Click **Save**

Delete a user macro

To delete a user macro:

1. Go to

   ⚙ ▾

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

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

Best practices

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

Add a descriptive header to your macro template

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

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

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

Expose your parameters in the macro browser

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

Supply default values for macro parameters

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

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

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

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

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

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

Example user macros

> ⌄ Hello World

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

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

Template	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. |
| | ```
## @noparams
<div class="aui-message closeable">
Hello World!
$body
</div>
``` |

Using the 'Hello World' macro on a page

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

The result is:

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

~ NoPrint

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

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

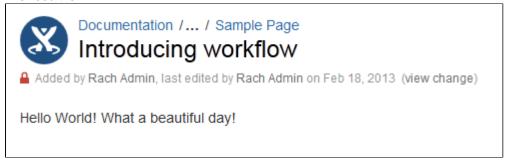| Description	Hides text from printed output.
Categories	Confluence Content
Icon URL	You can leave this field blank
Documentation URL	You can leave this field blank
Macro body processing	Rendered

Template	Enter the code below in the template field.
	
	 `## @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 recognise the NoPrint macro

See Advanced PDF Export Customisations.

⌄ Colour and Size

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

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

Template	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.

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

Alternatively, you can also use explicitly-named parameters in your macro. These macro parameters will appear as variables with the name $param<x> where <x> is the name of your parameter.

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

∨ Formatted Panel

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

(Title)

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

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

Template	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:parameter ac:name="titleBGColor">#ccc</ac:parameter> <ac:parameter ac:name="borderStyle">solid</ac:parameter> <ac:parameter ac:name="borderColor">#6699CC</ac:parameter> <ac:parameter ac:name="borderWidth">2</ac:parameter> <ac:parameter ac:name="titleColor">#000000</ac:parameter> <ac:parameter ac:name="title">$!paramTitle</ac:parameter> <ac:rich-text-body>$body</ac:rich-text-body> </ac:structured-macro>```

Explanation of the code in the macro template

Below is a breakdown of the user macro template code.

Item	Description

`## @param Title:title=Title|type=string|desc=Title`	@param defines the metadata for your macro parameters.
	
	`@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="panel">`	This calls the Confluence Panel ma cro. 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:name="titleBGColor">#ccc</ac:parameter>`
`<ac:parameter ac:name="borderStyle">solid</ac:parameter>`
`<ac:parameter ac:name="borderColor">#6699CC</ac:parameter>`
`<ac:parameter ac:name="borderWidth">2</ac:parameter>`
`<ac:parameter ac:name="titleColor">#000000</ac:parameter>`
	Sets the parameters for the macro: the background colour, border style, border colour, border width and title colour.
	
	To discover the names of the parameters for a Confluence macro, view the storage format as described above.

`<ac:parameter ac:name="title">$!paramTitle</ac:parameter>`	Enters the value stored in the 'Title' parameter into the title section of the macro. The ! tells the macro to leave the title blank, when there is no data in the "Title" parameter.
`<ac:rich-text-body>$body</ac:rich-text-body>`	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.
`</ac:structured-macro>`	This command marks the end of the macro.

User Macro Template Syntax

See Writing User Macros for an introduction to writing a user macro.

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

Accessing your macro's body

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

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

Example: Let's assume your macro is called `hello world`.
Enter the following code in your template:

```
Hello World: $body
```

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

```
From Matthew
```

The wiki page will display the following:

> *Hello World: From Matthew*

Using parameters in your user macro

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

How your macro parameters are used on a Confluence page

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

Defining the parameters

A parameter definition in the template contains:

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

Format:

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

Additional notes:

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

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 st
enum	Offers a list of values for selection. You can specify the values to appear in a dropdown i browser. Example of specifying the enum values: ```## @param colour:title=Colour|type=enum|enumValues=Grey,Red,Yellow,Gre``` *Note about i18n:* Confluence does not support internationalisation of the enum values.Th user sees is the one passed to the macro as the parameter value, with the capitalisation 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-content	Offers a control allowing the user to search for a page or blog post. Example:
	``` ## @param page:title=Page|type=confluence-content|required=true|desc=S a page do use ```
username	Search for user.
	``` ## @param user:title=Username|type=username|desc=Select user 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 st from. Format: dd/mm/YYYY ``` *Note about dates:* A user can enter a date in any format, you should validate the date for user macro.
int	Confluence accepts this type, but treats it in the same way as 'string'. Example with a def
	``` ## @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 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
$paramfoo, $parambar, ... $param<name>	Named parameters ("foo", "bar") passed to your macro.	String
$config	The BootstrapManager object, useful for retrieving Confluence properties.	BootstrapManager
$renderContext	The PageContext object, useful for (among other things) checking $renderContext.outputType	PageContext
$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).	ContentEntityObject

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

Controlling parameter appearance in the editor placeholder

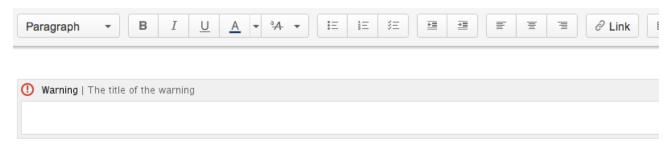You can determine which macro parameters should appear in the placeholder in the Confluence editor.

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

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

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

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

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

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

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

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

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

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

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

Customising your Confluence Site

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

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

We've documented the customisations under two broad headings:

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

Related pages:

- Integrating Confluence with Other Applications
- Tracking Customisations Made to your Confluence Installation

- Confluence Administrator's Guide

Changing the Look and Feel of Confluence

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

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

theme. No other theme will have an impact from this change. An individual space can be configured to have its own look and feel through the space administration screens.

Ways to customise the look and feel of your site:

- Add your own site logo. See Changing the Site Logo.
- Change the **colour scheme** of the user interface. See Customising Colour Schemes.
- Use **themes** for advanced layout customisation. See Working with Themes.

- Change the appearance of the **dashboard**. See Customising the Confluence Dashboard.
- Change the **site or space layouts**, which determine how the controls are laid out in the site. This does not change the actual page layouts, but it does change the way the surrounding controls appear in the page. See Customising Site and Space Layouts.

Customising the Confluence Dashboard

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

Editing the site welcome message

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

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

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

Using a page as the site landing page

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

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

Advanced customisations

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

There are two locations that you can add content to:

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

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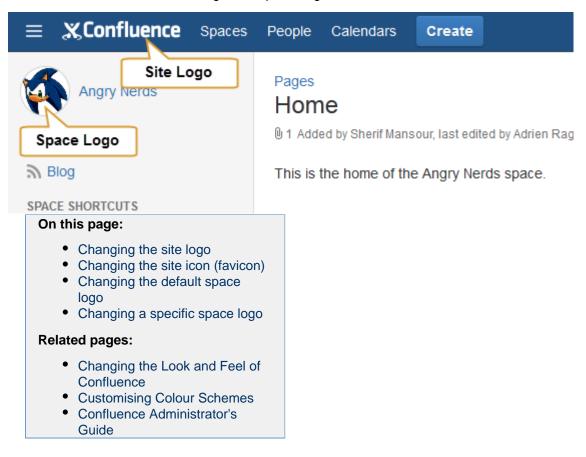
Changing the Site Logo

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

You can change:

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

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

Changing the site logo

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

To change the site logo:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Site Logo and Favicon**.
3. Choose **Browse** to upload a new logo.
4. Choose **Show Logo Only** or **Show Logo and Title** depending on whether you wish the SIte Title to display in the header.
5. Choose **Save**.

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

If you would prefer to use the default colour scheme with your custom logo go to

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

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

Changing the site icon (favicon)

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

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Site Logo and Favicon**
3. Locate your image file and choose **Upload**.

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

Changing the default space logo

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

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

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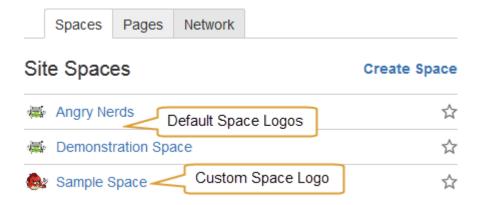
To change the default space logo:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Default Space Logo** in the left-hand panel.
3. Choose **Logo:ON**
4. Choose **Browse** to upload a new logo
5. Choose **Upload Logo**
6. Choose **Save**.

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

Changing a specific space logo

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

See Changing a Space's Logo to find out how to change the logo in a specific space.

Customising Colour Schemes

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

To change the site's colour scheme:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Colour Scheme** in the left-hand panel.
3. Choose **Edit**.
4. Enter standard HTML/CSS2 colour codes, or use the colour-picker

   to choose a new colour from the palette provided.
5. Choose **Save**. Any changes you make will immediately be reflected across the Confluence site.

> **On this page**
> - Reset your colour scheme after uploading a site logo
> - Notes
>
> **Related pages:**
>
> - Page Templates
> - Working with Themes
> - Changing the Look and Feel of Confluence
> - Confluence Administrator's Guide

Some UI elements below are for specific themes, and colour changes may not take effect for other themes.
- **Top Bar** - the top navigation bar background
- **Top Bar Text** - the text on the top navigation bar
- **Header Button Background** - buttons on the top navigation bar (e.g. Create button)
- **Header Button Text** - the text on buttons on the top navigation bar
- **Top Bar Menu Selected Background** - background colour of top navigation bar menu items when selected (e.g. spaces)
- **Top Bar Menu Selected Text** - text colour of top navigation bar menu items when selected
- **Top Bar Menu Item Text** - text on top navigation bar drop down menus (e.g. help or cog)
- **Menu Item Selected Background** - highlight colour on top navigation bar drop down menu items
- **Menu Item Selected Text** - text colour on highlighted top navigation bar drop down menu items
- **Search Field Background** - the background colour of the search field on the header
- **Search Field Text** - the colour of the text in the search field on the header
- **Page Menu Selected Background** - the background colour of the drop down page menu when selected
- **Page Menu Item Text** - the text of the menu items in the drop down page menu
- **Heading Text** - all heading tags throughout the space
- **Space Name Text** - the text of the current space name located above the page title
- **Links** - all links throughout the space
- **Borders and Dividers** - table borders and dividing lines
- **Tab Navigation Background** - the background colour of the tab navigation
- **Tab Navigation Text** - the text of the tab navigation when highlighted
- **Tab Navigation Background Highlight** - the background colour of the tab navigation when highlighted
- **Tab Navigation Text Highlight** - the text of the tab navigation elements when highlighted

Screenshot: Editing the colour scheme

Custom Colour Scheme

A custom colour scheme which can be edited.

Selected

Edit

The following colours can be customised for this colour scheme.

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

Reset   Save   Cancel

Reset your colour scheme after uploading a site logo

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

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

To reset the colour scheme:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Colour Scheme** in the left-hand panel.
3. Choose **Edit**.
4. Choose **Reset.**

Notes

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

Styling Confluence with CSS

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

Introduction

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

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

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

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

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

On this page:
- Introduction
- Considerations for Using Custom CSS
- Getting Started
- CSS Resources

Related pages:
- Basic Styling Tutorial
- Styling Fonts in Confluence

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.
Custom CSS does not apply to Admin screens

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

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

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

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

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

Getting Started

Editing the CSS

To edit a space's CSS style sheets:

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

If your space uses the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Click **Stylesheet** in the left-hand panel under the heading 'Look and Feel'.
3. Choose **Edit**.
4. Paste your custom CSS into the text field.
5. Save your changes. The new CSS will be visible on all content pages in the space.

To edit your global CSS stylesheet:

1. Choose the **cog icon**
   ⚙ ▾
   , then choose **General Configuration** under Confluence Administration
2. Choose **Stylesheet**.
3. Choose **Edit**.

4. Paste your custom CSS into the text field.
5. Choose **Save**.

Note:

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

Follow the Tutorial

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

CSS Resources

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

Basic Styling Tutorial

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

CSS Editing Quick-Start

To edit a space's CSS style sheets:

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

If your space uses the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Click **Stylesheet** in the left-hand panel under the heading 'Look and Feel'.
3. Choose **Edit**.
4. Paste your custom CSS into the text field.
5. Save your changes. The new CSS will be visible on all content pages in the space.

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 customising.
3. Note the page ID of the page where you uploaded the new graphic. (in this example, the page ID was '**658833839**'.
4. Compose your custom CSS for the header. The example below loads the new graphic (called '**header. png**') from a specific page (denoted by page ID '**658833839**') in the same space.

```
#header .aui-header  {

background-image:url('../download/attachments/658833839/header.p
ng');
        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 customisation, as it already includes all of the elements that can be changed. When you work on the space stylesheet it styles all content pages in the space. Build and test it at space-level, before considering applying the new stylesheet to your entire site. Once you are satisfied with your space design, test it thoroughly until you are confident that it has no problems. Then, you can look into advanced customisation of the Confluence CSS such as adjusting the Search page, the Dashboard and other integral pages.
Use the Right Tools

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

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

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

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

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

Notes

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

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

Styling Fonts in Confluence

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

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

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

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

Notes

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

⚙ ▾

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

Working with Themes

Themes are pre-defined style sets that you can apply to Confluence, to alter the appearance of your site. This is a way of personalising the 'look and feel' of Confluence. You can apply a theme to your entire Confluence site and to individual spaces. Choose a specific theme if you want to add new functionality or significantly alter the appearance of Confluence.

Confluence comes with a selection of themes. After a theme is installed on your Confluence site, any space administrator can apply a theme to a space.

A site administrator can install new themes as add-ons via the Confluence Administration Console.

By default when you create a new space, the space will have the Confluence default theme.

To look at the themes installed on your Confluence site:

1. Choose the **cog icon**

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

2. Choose **Themes** in the left-hand panel.
3. You will see a list of all installed themes.

Useful add-ons

Before installing an add-on (also called a plugin) into your Confluence site, please check the add-on's information page to see whether it is supported by Atlassian, by another vendor, or not at all. See our guidelines on add-on support.*Visit the Atlassian Marketplace to* search for additional themes *you can add to your site.*

Applying a Theme to a Site

You can use a theme to personalise the 'look and feel' of Confluence. Some themes simply change the basic styling, others add new functionality or significantly alter the appearance of Confluence. You can apply a theme to your entire Confluence site and to individual spaces.

Confluence comes with a selection of themes.In addition, a site administrator can install new themes as plugins via the Confluence Administration Console .Provided that the theme is available in your Confluence site, any space administrator can apply a theme to a space. By default when you create a new space, the space will have the Confluence default theme.

To apply a theme across the site:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Themes** in the left-hand panel.
3. The screen will display all available themes. Select a radio button to choose a theme.
4. Choose **Confirm**.

Screenshot: Applying a theme

Site Theme

Current Theme

The current theme controls the layout and colours of this space.

Default Theme

This is the original Confluence look and feel. Page content spans the full width of the screen.

Choose New Theme

To change the theme of this space, select one below.

○ Documentation Theme

This theme is well suited for structured content, such as documentation. It features a table of contents (page tree) on the left, making it easier to see the structure of a space and move from page to page. You can customise the left-hand panel, page header and page footer.

Find more themes...

Confirm

Creating a Theme

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

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

Related pages:

- Applying a Theme to a Site
- Apply a Theme to a Space
- Configure the Documentation Theme
- Confluence Administrator's Guide

Customising Site and Space Layouts

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

When you edit a site layout, you'll be modifying the default decorators in every space in your site, except

for those that have already been edited in a space. See Customise Space Layouts for more information on how to edit the decorators for a single space.

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

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

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

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

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

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

- **Site layouts**: These are used to define the controls that surround each page in the site. For example, the header, footer and dashboard.

- **Content layouts**: These control the appearance of content such as pages and blog posts. They do not change the way the pages themselves are displayed, but allow you to alter the way the surrounding comments or attachments are displayed.

- **Export layouts**: These control the appearance of spaces and pages when they are exported to HTML.

Editing a site decorator file

To edit a site decorator:

1. Go to

   > **General Configuration** > **Layouts** (under **Look and Feel**)
2. Click **Create Custom** next to the decorator `.vmd` file you want to modify.
3. Make your changes and click **Update**.

If something goes wrong: Hit **Reset Default** to revert to the original layouts.

Using Velocity macros

When editing Custom Decorator Templates, there are a number of macros available to define complex or variable parts of the page such as menus and breadcrumbs. You may insert these macros anywhere in your templates. More information on Working With Decorator Macros.

Advanced customisations

Overriding Velocity templates

The `velocity` directory is at the front of Confluence's Velocity template search path. As such, you can override *any* of Confluence's Velocity templates by placing an identically named file in the right place. While we don't recommend you do this unless you know exactly what you're doing, it does give you complete control over the look of every aspect of Confluence. It also means that you can edit your templates in a text-editor if you wish, rather than through the web interface.

Caching

Velocity is configured to cache templates in memory. When you edit a page from within Confluence, it knows to reload that page from disk. If you are editing the pages on disk, you will either have to turn off velocity's caching temporarily in `WEB-INF/classes/velocity.properties`, or restart the server to make your changes visible.

Location of Velocity files

You will find the Velocity files in your Confluence installation directory. The primary Velocity files are located in the `<CONFLUENCE-INSTALLATION>\confluence\decorators` directory. For example, you will find the following files in that directory: `main.vmd, space.vmd, form-aui.vmd, global.vmd`, and more.

Finding the layout via the URL

If the layout has changed so extensively as to not be visible, you can browse to the URL directly:

```
http://<confluence base
url>/admin/resetdecorator.action?decoratorName=decorators/main.vmd
```

Substitute the base URL and the appropriate `.vmd` file.

Upgrading Customised Site and Space Layouts

As Confluence evolves, so do the default site and space layouts that drive the rendering of every page. As new functionality is added or current functionally is changed, the default layouts are modified to support these changes.

> **Related pages:**
>
> - Customising Site and Space Layouts

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

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

To reapply your custom layouts, you need to:

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

Step 1. Obtain your Custom Layouts

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

If not, you should be able to find your customisations using the following method. This method extracts all site- and space-level layouts from your Confluence site as a single output. From this output, you should be able to identify your customisations.

> This method is handy to use if you have:
>
> - Many spaces with space layout customisations, or
> - Do not have an independent record of your site or space layout customisations.

Custom layouts are stored in the `DECORATOR` table within your Confluence database. You can `SELECT` for the source of the layout using SQL like this:

```
mysql> select SPACEKEY,DECORATORNAME,BODY from DECORATOR;
+----------+--------------------+------+
| SPACEKEY | DECORATORNAME      | BODY |
+----------+--------------------+------+
| NULL     | decorators/main.vmd | ... |
+----------+--------------------+------+
1 row in set (0.03 sec)
```

*This example was tested on **MySQL**, but should be applicable to all SQL databases.*

Step 2. Reapply your Customisations

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

What are 'major' and 'minor' releases? Major release upgrades are ones where the 1st digit of Confluence's version number or the 1st digit after the 1st decimal place differ after the upgrade, for example, when upgrading from Confluence 3.0 to 3.1, or 2.8 to 3.0. Minor release upgrades are ones where the 1st digit of Confluence's version number and the 1st digit after the 1st decimal place remain the same after the upgrade, for example, when upgrading Confluence 3.0 to 3.0.1.

If you have made Confluence site-wide layout customisations:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Select **Layouts** in the left-hand navigation panel. The decorators are grouped under **Site**, **Content** and **Export** layouts.
3. Ensure you have all your customisations available (preferably in a form which can be copied and pasted).
4. Click **Reset Default** next to the layout whose customisations need to be reapplied.
5. Click **Create Custom** next to the same layout and reapply your customisations (by copying and pasting them) into the appropriate locations within the new default layout.
6. Click the **Save** button.
7. Repeat this procedure from step 4 for each layout whose customisations need to be reapplied.

If you have made space-specific layout customisations:

1. Go to the space and choose **Space tools** > **Look and Feel** from the bottom of the sidebar
2. Choose **Layout**. The decorators are grouped under **Site**, **Content** and **Export** layouts.
3. Ensure you have all your customisations available (preferably in a form which can be copied and pasted).
4. Click **Reset Default** next to the layout whose customisations need to be reapplied.
5. Click **Create Custom** next to the same layout and reapply your customisations (by copying and pasting them) into the appropriate locations within the new default layout.
6. Click the **Save** button.
7. Repeat this procedure from step 5 for each layout whose customisations need to be reapplied.

If your space is using the Documentation theme:

1. Go to a page in the space.
2. Choose **Browse** > **Space Admin** from the header

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
3. Choose **Layout** from the left hand panel.
4. Follow the steps above.

Step 3. Test your Modifications Carefully

Changes may interact unpredictably with future versions of Confluence. When upgrading, you should always

test your custom modifications thoroughly before deploying them on a live site. It's beyond the scope of Atlassian Support to test and deploy these changes.

Turning Off Caching

Velocity is configured to cache templates in memory. When you edit a page from within Confluence, it knows to reload that page from disk. If you are editing the pages on disk, you will either have to turn off Velocity's caching temporarily in `WEB-INF/classes/velocity.properties`, or restart the server to make your changes visible.

The `velocity.properties` file is available in the `confluence-x.x.x.jar` file, where `x.x.x` is the Confluence version number. The JAR file is located in the `WEB-INF/lib` directory. If you wish to make modification to the files in the JAR, we recommend the following steps:

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

Working With Decorator Macros

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

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

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

Macro	Usage
`#breadcrumbs()`	Draws the "You are here" breadcrumbs list, like the one found above the page name in the default template.
`#includePage(pageTitle)`	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.
`#searchbox()`	Inserts a search box into the page, like the one to the far right of the breadcrumbs in the default template.
`#globalnavbar(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:
`#globalnavbar("table")`	Displays the navigation bar in its default mode: drawn as a table of links with coloured backgrounds and mouse-over effects.
`#globalnavbar("text")`	Displays the navigation bar as series of text links separated by | characters.
`#usernavbar()`	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.
`#helpicon()`	Draws the help icon, and link to the Confluence help page.
`#printableicon()`	On pages where a printable version is available, draws the printable page icon, linking to the printable version of the page. Otherwise, draws nothing

`#pagetitle(class)`	When you are viewing a page in a Confluence space, draws the name of the space that page is in. Otherwise, writes the word "CONFLUENCE".The "class" argument is the CSS class that the title should be drawn in. Unless you have customised your Confluence installation's CSS file, you should call this with "spacenametitle" as the class: `#pagetitle("spacenametitle")`
`#poweredby()`	Writes out the "Powered by Confluence" and Confluence version-number boilerplate found at the bottom of the default template.
`#bottomshadow()`	Draws the fading shadow-effect found at the bottom of the content area in the default template.
`#dashboardlink()`	Inserts a link to the dashboard page.

Custom Decorator Templates

About Decorators

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

What this means for Confluence is that you can customise the look and feel of parts of your Confluence site through editing decorators, for example:

- The "Main" decorator defines the generic header and footer
- The "Page" decorator defines how a page is displayed
- The "Printable" decorator defines the look and feel of the printable versions of pages.

You can view and edit these decorators from within Confluence. Changes to the decorators will affect all spaces in that Confluence installation.

The decorator that is used to draw Confluence's administrative pages cannot be edited from within Confluence. This means that if you make a mistake that renders the rest of the site unuseable, the administrative pages should still be available for you to fix the template.

Browsing the Default Decorators

At any time, you can browse the default decorators that come packaged with Confluence by following the "View Default" links on the "Site Layouts" page. The template browser also allows you to view the "#parsed" templates that are included within the template when it is compiled. While you can't edit these included templates, you will probably have to copy some or all of them into your custom template as you do your customisation.

Editing Custom Decorators

To edit Confluence decorators you will need a good knowledge of HTML, and some understanding of the Velocity templating language.

To edit a decorator:

1. Go to **Confluence Admin** > **Layouts**.
2. Choose **Create Custom** beside the decorator you wish to edit.
3. Save your changes.

If you make a mistake or want to undo your changes, choose **Reset Default** beside the edited decorator.

Alternatively, the custom templates are stored in the DECORATOR table in the database. If you have somehow managed to render Confluence completely unuseable through editing your templates, delete the relevant entries from the DECORATOR table.

Macros

Some parts of the page are drawn using Velocity macros, including the navigation bar. The macros you should know about when editing decorators are described in Working With Decorator Macros.

For Advanced Users

The `velocity` directory is at the front of Confluence's velocity template search path. As such, you can override *any* of Confluence's velocity templates by placing an identically named file in the right place.

While we don't recommend you do this, it does give you complete control over the look of every aspect of Confluence. It also means that you can edit your templates in a text-editor if you wish, rather than through your browser.

There are, however, two important caveats:

1. Velocity is configured to cache templates in memory. When you edit a page from within Confluence, it knows to reload that page from disk. If you are editing the pages on disk, you will either have to turn off velocity's caching temporarily in `WEB-INF/classes/velocity.properties`, or restart the server to make your changes visible.
2. Changes may interact unpredictably with future versions of Confluence. When upgrading, you should always test your custom modifications thoroughly before deploying them on a live site.

Customising a Specific Page

If you'd like to change the appearance of a specific page, you can modify the corresponding Velocity template. Here's how to find out which one:

1. Access the page. Note the name of the action. For example, the "Contact Administrators" page is `<baseUrl>/administrators.action`.
2. Browse to <confluence-install>/confluence/WEB-INF/lib/confluence-x.y.jar. Copy the file.
3. Unzip or unjar the file using a standard unzipper or the java jar utility.
4. Open xwork.xml. Search the file for the name of the action corresponding to the page you'd like to modify. You'll see an entry like:

```
<action name="administrators"
class="com.atlassian.confluence.user.actions.AdministratorsAction">
        <interceptor-ref name="defaultStack"/>
        <result name="success" type="velocity">/administrators.vm</result>
    </action>
```

5. The file to look for is the vm or vmd file. In the above example, it's administrators.vmd. Because there is no context path (just a / before the name of the file), 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.

Customising the Login Page

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

Notes:

- Customisations to the Confluence login page will need to be reapplied when you upgrade Confluence. Consider this before making drastic changes to the layout, and be sure to keep a list of what you have changed for your upgrade process later.
- Please test your changes on a **test** Confluence site first.

Only administrators with access to the server where Confluence is running can modify the Confluence login page.

Related pages:

- Changing the Site Logo
- Velocity Template Overview
- Customising Site and Space Layouts
- Changing the Look and Feel of Confluence
- Modify Confluence Interface Text

To change the login page:

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

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

Modify Confluence Interface Text

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

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

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

Replace "x.x.x" with your Confluence version, for example for 4.3.2, it
will be named "confluence-4.3.2.jar".
Within this File, the relevant file to edit is
:\com\atlassian\confluence\core\ConfluenceActionSupport.properties.
```

Refer to Editing jar files for reference.

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

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

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

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

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

Steps For Modification

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

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

4. Restart Confluence

Modify Keyboard Shortcuts

Confluence provides a set of keyboard shortcuts. You could customise the shortcuts by making modifications inside the `ConfluenceActionSupport.properties` file.

- To disable a particular shortcut, you can simply just comment out a respective line of code. One may like to disable the shortcut to one of the navigation links: *View, Edit, Attachments, Info* . For instance, to disable shortcut to *Attachments*one would comment out the following line:

```
#navlink.attachments.accesskey=a
```

- To modify an access key, one could simply just change the letter, bearing in mind the fact that the letter must be unique.

Customising the eMail Templates

> Customising the Confluence email templates is **not supported**. If you do decide to edit the templates we strongly recommend you use a test instance of Confluence.
>
> Any customisations you make to the Confluence email notification templates will need to be reapplied after upgrading Confluence.

Email notification templates are contained within the `confluence-email-notifications` plugin, which is a bundled plugin (add-on) that is installed automatically when you install Confluence.

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

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

To change the email notification templates:

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

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

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

Changing the Default Behaviour and Content in Confluence

Confluence comes with some handy default settings that determine what people see when they first enter the Confluence site, and the default content that is put into new spaces and other areas of Confluence.

Confluence administrators can change the settings to customise the behaviour and the default content of their Confluence site:

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

Related pages:

- Changing the Look and Feel of Confluence
- Customising your Confluence Site
- Confluence Administrator's Guide

Administering Site Templates

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

Administrators can import templates, to make them available to other people using Confluence. See Importing Templates.

Confluence also provides 'system templates' which contain default content for the site welcome message (see Editing the Site Welcome Message) and default space content (see Customising Default Space Content).

Administrators can also disable templates and blueprints, to stop them appearing in the Create and Create Space dialogs anywhere in their Confluence site.

To **disable** a template or blueprint across the entire Confluence site:

- Choose the **cog icon**

  ⚙ ▾

  , then choose **General Configuration** under Confluence Administration
- Choose **Global Templates and Blueprints**.
- Choose **Disable** next to the template, page blueprint or space blueprint you wish to disable.

Administrators can re-enable these templates and blueprints at any time.

Importing Templates

A template is a predefined page that can be used as a prototype when creating new pages. Templates are useful for giving pages a common style or format.

You can create your own templates within Confluence. See Create a Template.

In addition, you can download pre-defined templates from the Atlassian Marketplace in the form of a template bundle. Each template bundle contains one or more templates, created by Atlassian or third parties. Here is a summary of the steps required:

- Download the template bundle from the Atlassian Marketplace.
- Install the template bundle into your Confluence site.
- Make the templates available by importing them into the site or into an individual space.

You need 'System Administrator' permission to install template bundles into your Confluence site. You need 'Confluence Administrator' permission to manage the existing template bundles on your Confluence site. See Global Permissions Overview.

On this page:

- Step 1. Check the template bundles installed on your Confluence site
- Step 2. (Optional) Download and install additional template bundles from the Atlassian Marketplace
- Step 3. Import the templates to make them available to users
- Notes

Related pages:

- Pages and Blogs
- Page Templates

Step 1. Check the template bundles installed on your Confluence site

To see the template bundles that are currently available for import on your Confluence site:

1. Log in to Confluence as a System Administrator or Confluence Administrator.
2. Choose the **cog icon**

   ⚙️ ▾

   , then choose **General Configuration** under Confluence Administration
3. Choose **Import Templates** in the left-hand panel. You will see a list of the template bundles installed on your Confluence site, and the templates included in each bundle.

Step 2. (Optional) Download and install additional template bundles from the Atlassian Marketplace

Follow the steps below if you want to add more template bundles to your site.

Before installing an add-on (also called a plugin) into your Confluence site, please check the add-on's information page to see whether it is supported by Atlassian, by another vendor, or not at all. See our guidelines on add-on support.

To upload more templates:

1. Go to the Atlassian Marketplace and download the template bundle that you need. It will be in the form of a JAR file. Save the JAR file somewhere in your file system.
2. Log in to Confluence as a System Administrator.
3. Choose the **cog icon**

   ⚙️ ▾

   , then choose **General Configuration** under Confluence Administration
4. Choose **Manage Add-ons** in the left-hand panel.
5. Choose **Upload Add-on**.
6. Browse to find the template bundle that you downloaded, and upload it to Confluence. The template bundle will appear in the list under 'User-installed Add-ons'.

Step 3. Import the templates to make them available to users

You now have one or more template bundles on your site. The templates are not available until you have 'imported' them.

To import a template:

1. Log in to Confluence as a System Administrator or Confluence Administrator.
2. Choose the **cog icon**

   ⚙️ ▾

   , then choose **General Configuration** under Confluence Administration
3. Choose **Import Templates** in the left-hand panel. You will see the template bundles installed on your Confluence site and the templates included in each bundle.
   Note: You can see a preview of the template by choosing the template name.
4. Select the templates to be imported by ticking the check boxes next to the relevant template names.
5. Choose the import destination for the templates in the **Import To** dropdown menu. If you want the templates to be available to only a specific space, choose the name of the space, otherwise choose **Global Templates** to make the templates available to all spaces.
6. Choose **Import**.

Screenshot: Importing a template

Import Templates

The following template package plugins were found. To import templates, tick the checkboxes for the desired templates, select where to import the templates to, and click the Import button.

Human Resources Templates (8)

Check All Uncheck All

- ☐ HOW-TO Guide
- ☐ Induction Tasks
- ☐ Job Description
- ☐ Meeting Minutes
- ☐ Recruitment Dashboard
- ☐ Space Home Page
- ☐ Time Sheet
- ☑ Wiki Induction

Import To: Global Templates ▾

> **Global Templates**
> **Spaces**
>   Demonstration Space
>   Documentation

Import

Notes

- **Building your own template bundles.** You can build a template bundle as an add-on (also called a 'plugin') and then upload it to your Confluence site. You can then import the templates from your custom template bundle, as described above. You will need some programming knowledge to develop a template bundle. See Creating A Template Bundle.

- **Duplicate template names.** If a template with the same name already exists on import, a duplicate template of the same name will be created. You will need to check the templates and rename them manually.

- **Removing the template.** Removing the add-on that contains a template will not remove the template from your Confluence site if you have already imported it. You will need to remove the template manually via the administration console or space administration screen.

Changing the Site Title

The site title appears in your browser's title bar. By default, it is set to 'Confluence'.

To change the title of your Confluence site:

1. Choose the **cog icon**
   ⚙ ▾
   , then choose **General Configuration** under Confluence Administration
2. Choose '**General Configuration**' in the left-hand panel.
3. Choose '**Edit**' at the top of the '**Site Configuration**' screen.
4. Enter a new title for your site in the input field next to '**Site Title**'.
5. Choose '**Save**'.

Related pages:

- Changing the Site Logo
- Editing the Site Welcome Message
- Customising your Confluence Site
- Confluence Administrator's Guide

Choosing a Default Language

Administrators can define a default language to be applied to all spaces in your Confluence site. Note that individual users can select a language preference for their session.

Related pages:

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

Setting the defaultlanguage

To change the default language for the Confluence site:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Select '**Languages**' in the 'Configuration' section of the left-hand panel
3. Choose **Edit** and select the language you want to use as the default language for your Confluence site

Screenshot: Default languages in Confluence

> # Global Default Language
>
> Installed Languages   Deutsch (Deutschland)
>
> **English (United Kingdom)**
> español (España)
> français (France)
> 日本語 (日本)
>
> 🔒 Find more language packs...
>
> Edit

Other settings that affect the language

Individual users can choose the language that Confluence will use to display screen text and messages. Note that the list of supported languages depends on the language packs installed on your Confluence site.

The language used for your session will depend on the settings below, in the following order of priority from highest to lowest:

- The language preference defined in your user profile. Note that you need to be logged in for this setting to take effect.
- The language that you choose by clicking an option at the bottom of the Confluence login screen. Confluence stores this value in a cookie. When the cookie expires, the setting will expire too.
- The language set in your browser. The browser sends a header with a prioritised list of languages. Confluence will use the first supported language in that list. Your Confluence administrator can disable this option by setting a system property.
- The default language for your site, as defined by your Confluence site administrator.

Showing User Interface Key Names for Translation

This feature is useful if you are working on creating translations of the Confluence user interface. After opening the Confluence dashboard, you can add this text to the end of your Confluence URL:

```
?i18ntranslate=on
```

Then press Enter.

This will cause each element of the user interface to display its special **key name**. This makes it easier to find the context for each key within the user interface. You can then search for the key on http://translations.atlassian.com where you can enter an appropriate translation for your custom language pack.

The key names are displayed with a 'lightning bolt' graphic. For example:

Dashboard⚡title.dashboard   👤 Invite Users⚡easyuser.add.users.button   ⊕ Create Space⚡dashboard.button.add.space

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

```
?i18ntranslate=off
```

Configuring the Administrator Contact Page

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

See the explanation of Confluence Groups for Administrators.

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

On this page:

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

Related pages:

- Configuring Captcha for Spam Prevention

Customising the Administrator Contact Message

You can customise the message that is presented to the user on the '**Contact Site Administrators**' page. **To edit the administrator contact message:**

1. Choose the **cog icon**
   ⚙▾
   , then choose **General Configuration** under Confluence Administration
2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit** at the top of the 'Site Configuration' section.
4. Enter your text in the **Custom Contact Administrators Message** box. You can enter any text or Confluence wiki markup .
5. Choose **Save**.

The Default Administrator Contact Message

By default, the 'contact administrators message' looks much like the highlighted area in the screenshot below, starting with 'Please enter information...'.

Screenshot: The default 'Contact Site Administrators' message

Contact Site Administrators

Please enter information about your request for the site administrators. If you are reporting an error please be sure you include information on what you were doing and the time the problem occurred.

To	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**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit** at the top of the 'Site Configuration' section.
4. Select **on** or **off** for the 'Contact Administrators Form'.
5. Choose **Save**.

Configuring Spam Prevention

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

Configuring the Site Home Page

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

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

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

Related pages:

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

To use a page as your site home page:

1. Go to

   ⚙ ▾

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

> **Note about permissions**
> Before changing the site homepage you should check that the default 'confluence-users' or 'users' groups have permissions to view the space the page was created in, and that the page itself is not restricted to particular people or groups.
>
> If your site is public, you'll also need to make sure anonymous users have permissions to view the space, otherwise anonymous users will be directed to the dashboard instead.

Accessing the dashboard with a site homepage set

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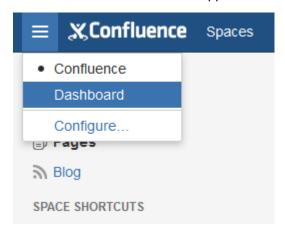
To add the Confluence Dashboard to the Application Navigator:

1. Go to

   > **General Configuration** > **Application Navigator**.
2. Enter the name for your link, for example, 'Dashboard'.
3. Enter the URL for your site dashboard, for example, `https://yoursite.com/wiki/dashboard.action`.
4. Choose **Add**.

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

Customising Default Space Content

Confluence Administrators can edit the template that is used to create the home page for new sites. This default content appears on the home page when a new space is created. There is a different template for site spaces, personal spaces and space blueprints.

The default content in the template only appears for new spaces (those that are created after you have defined the content). Changes to the template do not affect existing home pages.

Edit the default home page for a blank space

To edit the default (blank) space content template:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Global Templates and Blueprints** in the left-hand panel.

3. Choose **Edit** next to 'Default Space Content' or 'Default Personal Space Content' depending on whether you want to customise the content for new site space or personal space home pages.
4. Enter the content that you want to appear on the home page for new blank spaces. you can add variables, macros and other content in the 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
> - Confluence Administrator's Guide

The following variables are available to be added to the default space content templates.

- **$spaceKey** - inserts the space key into the site space homepage
- **$spaceName** - inserts the space name into the site space homepage
- **$userFullName** - inserts the user (owner of the personal space) into the personal space homepage
- **$userEmail** - inserts the email address of the user (owner of the personal space) into the personal space homepage.

Default space templates differ from ordinary page templates in that they do not present the user with a form to complete, so variables should be limited to those listed in the **Variables** menu.

Some macros, such as the Table of Contents macro, may not display correctly when you preview the template as they are designed to work on a page. The macros will display correctly on the home page when you create a new space. For more information on editing a template, including adding macros see - Adding Content to a Template.

Reset the original default content

To reset the original default content:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **Global Templates and Blueprints** in the left-hand panel.
3. Choose **Reset to default** next to the template you wish to reset.

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

Editing the Site Welcome Message

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

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

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

> **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

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. Scroll down to the System templates and choose **Edit** next to **Default Welcome Message**.

3. Add your content and choose **Save**.

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

Screenshot: Default site welcome message

Hints for using the template editor

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

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

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

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

Allowing other people to edit the site welcome message

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

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

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

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

Integrating Confluence with Other Applications

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

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

Related Topics

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

Linking to Another Application

Application Links (sometimes called "AppLinks") is a bundled plugin that allows you to link Atlassian applications to each other. Linking two applications allows you to share information and access one application's functions and resources from within the other.

Atlassian only recommends using OAuth authentication for application links, because of the greater security inherent with that protocol. We no longer recommend the Trusted Applications and Basic authentication types.

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

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **Application Links** in the left-hand panel. The Application Links configuration page appears and lists any links you already have set up.
3. Enter the URL of the application you want to link to, then click **Create new link**.
   - If you check **The servers have the same set of users...** then user impersonation with 2-Legged OAuth authentication will be configured for this link. You can change this later if necessary.
   - If you are *not* an admin on both servers you won't be able to set up a 2-way (reciprocal) application link. If you want to go ahead and create a 1-way link anyway, clear the **I am an administrator on both instances** checkbox.
4. Use the wizard to finish configuring the link. If the application you are linking to does not have the Application Links plugin, you must supply additional information to set up a link with OAuth authentication.

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

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

- Edit the settings of the application link (for example, to change the authentication type of the link) using the **Edit** link.
- Specify the default instance if you have multiple links to the same type of application (for example, to multiple JIRA servers) using the **Make Primary** link. See Making a primary link for links to the same application type for more information.

> **Having trouble integrating your Atlassian products with application links?**
> We've developed a guide to troubleshooting application links, to help you out. Take a look at it if you need a hand getting around any errors or roadblocks with setting up application links.

Configuring Workbox Notifications

You can view and manage in-app notifications and tasks in your Confluence workbox. In addition, you can receive notifications from JIRA applications nd other Confluence servers in your Confluence

workbox. To make this possible, your Confluence server must be linked to the other server(s) via appli cation 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?
- Enabling Confluence workbox and in-app notifications
- Configuring the polling intervals
- Including notifications from JIRA
- Stopping JIRA applications from sending notifications to Confluence
- Including notifications from another Confluence server
- Sending Confluence notifications to another Confluence server
- Disabling workbox and in-app notifications in Confluence

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.

Enabling Confluence workbox and in-app notifications

Confluence workbox and in-app notifications are disabled by default.

To enable workbox and in-app notifications:
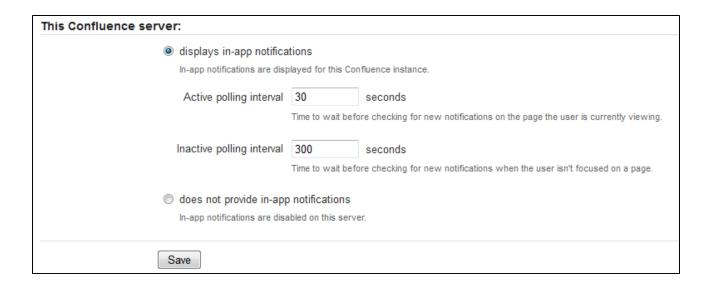
1. Choose the **cog icon**

    , then choose **General Configuration** under Confluence Administration
2. Choose **In-app Notifications** in the left-hand panel
3. Choose **displays in-app notifications** (or **displays in-app notifications from other servers**)
    The workbox icon

    will appear in the Confluence top menu bar and will be visible to all users.

Screenshot: Simple configuration with Confluence workbox and in-app notifications enabled for this server only

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 Desk .

To include notifications from a JIRA application:

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

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **In-app Notifications** in the left-hand panel of the Confluence administration console.
3. Choose **displays in-app notifications from other servers**.

   - Your JIRA application will appear in the list of linked applications below this option.
   - People will see JIRA notifications in their workbox, as described in Workbox Notifications.
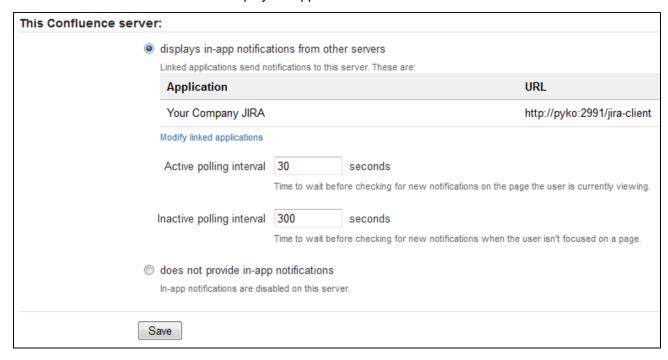
Notes:

- JIRA sends its notifications to the Confluence server that is configured as the **primary** application link.
- Your JIRA server must be running **JIRA 5.2 or later**.
- The following plugins must be present and enabled in JIRA. The plugins are shipped with JIRA 5.2

and later:
- 'Workbox – Common Plugin'
- 'Workbox – JIRA Provider Plugin'
- You do not need to configure JIRA. The plugins are enabled by default in JIRA, and JIRA will automatically send notifications to Confluence.
- Confluence can display notifications from more than one server.

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

Stopping JIRA applications from sending notifications to Confluence

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

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

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

- 'Workbox – Common Plugin'
- 'Workbox – JIRA Provider Plugin'

Including notifications from another Confluence server

Confluence workbox can include notifications from another Confluence server.

Let's assume that you have two Confluence servers, *ConfluenceChatty* and *ConfluenceQuiet*. Let's also assume that you want *ConfluenceChatty* to display a workbox, and to include notifications from *ConfluenceQuiet*.

To include notifications from other Confluence servers:

1. Connect *ConfluenceChatty* and *ConfluenceQuiet* via application links. In *ConfluenceChatty*:
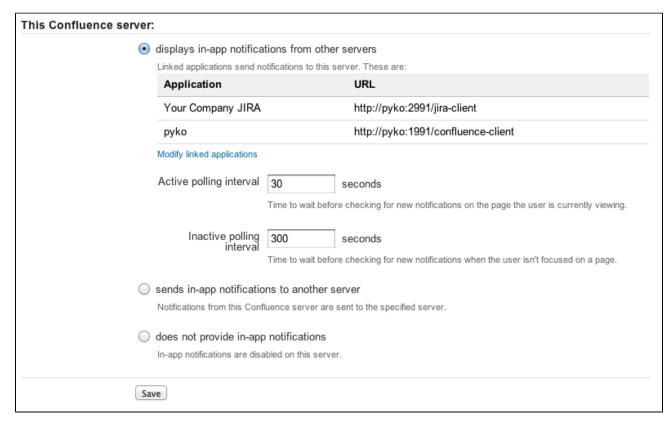   - Choose the **cog icon**

     , then choose **General Configuration** under Confluence Administration
   - Choose **Application Links** in the left-hand panel.
   - Set up the link as described in Linking to Another Application.
2. Configure the notification settings in *ConfluenceChatty*:
   - Choose **In-app Notifications** in the left-hand panel of the Confluence administration console.

- Choose **displays in-app notifications from other servers**.
3. Configure the notification settings in *ConfluenceQuiet*:
    - Choose **In-app Notifications** in the left-hand panel of the Confluence administration console.
    - Choose **sends in-app notifications to another server**.
    - Select the Confluence server that will display the workbox – in our example, this is *Confluence Chatty*. (The entry for *ConfluenceChatty* will appear here only if you have already configured *ConfluenceChatty* to display in-app notifications.)

Notes:

- Your Confluence servers must be running **Confluence 4.3.3 or later**.
- Confluence can display notifications from more than one server.
- Confluence can send notifications to only one server.
- Only one of the linked Confluence servers can display the in-app notifications.

Screenshot: This Confluence server displays in-app notifications from itself, from JIRA, and from another Confluence server

Sending Confluence notifications to another Confluence server

You can configure Confluence to send all notifications to a different Confluence server. In this case, the current Confluence server will not display the workbox.

To send notifications to another Confluence server: Follow the instructions in our example for *ConfluenceQuiet* above.

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

This Confluence server:

- displays in-app notifications from other servers
  Linked applications send notifications to this server.

- sends in-app notifications to another server
  Notifications from this Confluence server are sent to the specified server.

	Application	URL
◉	pyko	http://pyko:1991/confluence-client

- does not provide in-app notifications
  In-app notifications are disabled on this server.

Save

Disabling workbox and in-app notifications in Confluence

If you choose **does not provide in-app notifications**:

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

Integrating JIRA and Confluence

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

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

On this page:

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

Installing JIRA and Confluence together

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

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

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

Use JIRA and Confluence together

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

Delegate user management to JIRA

If you already have a JIRA application you can choose to delegate user management to JIRA, and manage all your users in one place. You can control which JIRA groups also have permissions to use Confluence.

Your license tiers for each application do not need to be the same.

See Configuring JIRA Integration in the Setup Wizard to delegate user management to JIRA when installing Confluence for the first time.

See Connecting to Crowd or JIRA for User Management to delegate user management to JIRA for an existing Confluence site.

Connect JIRA and Confluence with an application link

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

If you delegated user management to JIRA as part of Confluence's setup process, an application link to JIRA (using Trusted Applications authentication) will be all set up and ready to go.

You can choose to change the authentication to OAuth for additional security. See Configuring authentication for an application link to find out about changing the authentication to OAuth for additional security.

> **Having trouble integrating your Atlassian products with application links?**
> We've developed a guide to troubleshooting application links, to help you out. Take a look at it if you need a hand getting around any errors or roadblocks with setting up application links.

Registering External Gadgets

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

There's two ways to register external gadgets:

- **Subscribe to all of the external application's gadgets:** You can add all the gadgets from your JIRA application, Bamboo, FishEye or Crucible site – or from another Confluence site – to your Confluence gadget directory. People can then pick and choose the gadgets to add to their Confluence pages.
- **Register the external gadgets one by one:** If you cannot subscribe to an application's gadgets, you will need to add the gadgets one by one. This is necessary for applications and websites that do not support gadget subscription, and for applications where you cannot establish a trusted relationship via Application Links.

Both methods are described below. First, consider whether you need to set up a trust relationship between Confluence and the other application.

> **On this page:**
>
> - Setting up a trust relationship with the other application
> - Subscribing to all of the application's gadgets
> - Registering individual gadgets
> - Removing access to external gadgets
>
> **Related pages:**
>
> - Configuring the Whitelist
> - The big list of Atlassian gadgets
>
> - Linking to Another Application

Setting up a trust relationship with the other application

In addition to registering the external gadgets, we recommend that you set up an OAuth or Trusted Application relationship between the application that serves the gadget (the service provider) and Confluence (the consumer). The trust relationship is required for gadgets that access restricted data from the external web application.

See how to configure OAuth or Trusted Applications Authentication, using Application Links.

If the external web application provides anonymous access to all the data you need in the gadgets, then you

do not need a trust relationship.

For example, if your gadgets will retrieve data from JIRA and your JIRA server includes projects and issues that are restricted to logged-in users, then you will need a trust relationship between Confluence and JIRA. If you do not set up the trust relationship, then the gadgets will show only the information that JIRA makes visible to anonymous users.

Subscribing to all of the application's gadgets

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

To subscribe to another site's gadgets:

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

Screenshot: Subscribing to a gadget feed

External Gadgets

Only add gadgets that you trust! Gadgets can allow unwanted or malicious code onto your web page.

| Gadget Specifications | Gadget Feeds | Gadget Whitelist |

Add a new Gadget Feed

Gadget Feed URL [   ]

[Add]

Added Gadget Feeds

Gadget Feed URL	Actions
https://pug.jira.com/rest/gadgets/1.0/g/feed	[Delete]
https://pug.jira.com/rest/gadgets/1.0/g/feed/	[Delete]

Registering individual gadgets

If you cannot subscribe to an application's gadgets, you will need to register the gadgets one by one. This is necessary for applications and websites that do not support gadget subscription, and for applications where you cannot establish a trusted relationship via Application Links.

First you will need to get the gadget URL and copy it to your clipboard.

Getting a gadget's URL from an Atlassian application

If your application is another Atlassian application:

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

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

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

```
http://my-app.my-server.com:port/rest/gadgets/1.0/g/my-plugin.key:my-gadget/my-p
ath/my-gadget.xml
```

For example:

```
http://mycompany.com/jira/rest/gadgets/1.0/g/com.atlassian.streams.streams-jira-
plugin:activitystream-gadget/gadgets/activitystream-gadget.xml
```

To find a gadget's URL in JIRA:

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

To find a gadget's URL in Confluence:

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

Getting a gadget's URL from another application

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

Registering the gadget for use in Confluence

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

To register the gadget in Confluence:

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

Screenshot: Registering external gadgets one by one

External Gadgets

Only add gadgets that you trust! Gadgets can allow unwanted or malicious code onto your web page.

| Gadget Specifications | Gadget Feeds | Gadget Whitelist |

You can add gadgets from Atlassian applications such as Confluence, JIRA and others. You can also add gadgets from other websites such as iGoogle. Many public gadgets will work on a Confluence page. Some gadgets may rely on specific OpenSocial features that will not work properly in Confluence.

If you are adding gadgets from another Atlassian Application you need to either setup the other application to trust Confluence using Trusted Applications (make sure you add the '/rest' path to the allowed URL paths) or add Confluence as an OAuth Consumer to the other application.

A gadget's URL looks something like this: http://example.com/my-gadget-location/my-gadget.xml

Add a new Gadget

Gadget Specification URL

[   ]

[Add]

Added Gadgets

Gadget Specification URL	**Actions**
https://pug.jira.com/rest/gadgets/1.0/g/com.atlassian.streams.streams-jira-plugin:activitystream-gadget /gadgets/activitystream-gadget.xml	[Delete]
https://pug.jira.com/rest/gadgets/1.0/g/com.atlassian.jira.gadgets:created-vs-resolved-issues-chart-gadget /gadgets/createdvsresolved-gadget.xml	[Delete]

Removing access to external gadgets

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

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

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

Configuring the Office Connector

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

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

On this page:

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

Related pages:

- Office Connector Prerequisites
- Office Connector Limitations and Known Issues

Enabling and disabling the Office Connector

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

To enable or disable the Office Connector modules:

1. Go to

> **Add-ons**
2. Choose **System** from the filter drop down and then search for **Office Connector**
3. Expand the Office Connector add-on listing. From here you can:
    * Choose **Configure** to specify preferences for the Office Connector (this opens the configuration screen described below)
    * Click **Disable** to disable all modules of the add-on
    * Expand the **modules** list to enable or disable selected Office Connector modules

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

Configuring the Office Connector Options

Users with System Administrator permissions can configure the behaviour of the Office Connector.

To set the configuration options for the Office Connector:

1. Go to

    > **General Configuration** > **Office Connector**
    Screenshot: Configuring the Office Connector options

Configure Office Connector plugin

User Interface

Edit in word button location: ○ Page action icon
                                ○ View page tab

Importing From Word

Warnings: ☐ Show a warning before allowing a user to perform an import
Advanced Formatting Options: ☐ Use the footnote macro for Word footnotes
Maximum imported image size: | 1200 | height (pixels)
                               | 900 | width (pixels)

System Resource Usage

Temporary storage for viewfile macro: ○ Confluence home directory: C:\Program Files\Atlassian\Application Data\Confluence\viewfile
                                        ○ No directory specified in the properties file. Please set a correct value in resources/directories.properties.

Maximum file space for cache(MB): | 500 |
Number of Conversion Queues: | 6 | Manage Queues

Security

Authentication: ☐ Allow authentication tokens in the URL path

                  Submit

2. Set the configuration options as described in the table below

Option	Default Value	Description
Edit in word button location	Page action icon	Where the button for editing the content in Word is located. You can configure action icon or from the view page tab. **Note:** This setting has no effect in modern Confluence themes, including the C documentation theme.
Warnings: Show a warning before allowing a user to perform an import	Disabled	If this option is enabled, the user will receive a warning when importing a Word user when they are about to overwrite existing content.
Advanced Formatting Options: Use the footnote macro for Word footnotes	Disabled	**Note:** This feature requires a third party add-on that is not supported for Conflu If this option is enabled, a Confluence page created from an imported Word do from Adaptavist to render any footnotes contained in the document. Note that y add-on onto your Confluence site. For more information about this add-on and add-on.
Authentication: Allow authentication tokens in the URL path	Disabled	If this option is enabled, the Office Connector will use authentication tokens in t This needs to be enabled to edit Office 2013 documents.

Temporary storage for viewfile macro	The Confluence Home directory.	The {viewfile} macro will cache data temporarily. This option allows you to set t settings are:
		
		• **Confluence home directory** – The temporary file will be stored in your Co
		• **A directory specified in the directories.properties file** – You can
		  Office Connector's directories.properties file:
		   1. Locate the OfficeConnector-x.xx.jar file (where x.xx is the versi
		     Home directory and copy it to a temporary location
		   2. Unzip the JAR file and find the resources/directories.properti
		     like this:
		
		<pre>#Complete the following line to set a custom ca
		#If resetting to blank, don't delete anything b
		the '='
		com.benryan.confluence.word.edit.cacheDir=</pre>
		   3. Edit the last line, adding the path to your required temporary location di
		     example:
		    • On Windows:
		<pre>com.benryan.confluence.word.edit.cacheDir=c:'</pre>
		    • On Linux:
		<pre>com.benryan.confluence.word.edit.cacheDir=/hc</pre>
		   4. Save the file, recreate the JAR and put it back in your Confluence Hom
		     JAR.

Maximum file space for cache (MB)	500	This is the maximum size of the cache used by the {viewfile} macro. (See abov
Number of Conversion Queues	6	This is the maximum number of threads used to convert PowerPoint, Excel file this setting to manage Confluence performance, by limiting the number of threa does not consume too many resources. Click **Manage Queues** to view attachments that are still pending conversion.

Managing your Confluence License

Your license entitles you to run Confluence and be eligible for support and upgrades for a specified period. It also defines the number of users who are entitled to use Confluence.

To quickly check the status of your license you can go to

⚙ ▾

> **General Configuration** > **Support Tools**.

You'll need need Confluence Administrator or System Administrator permissions to view and edit your license.

Viewing your license details

To view your Confluence license:

1. Go to

   ⚙ ▾

   > **General Configuration**.
2. Choose **License Details** in the left-hand panel.

The License Details page tells you:

- The type of license (for example: Commercial, Academic, Community, or Evaluation).
- Number of users you are licensed for, and how many are currently in use.
- Your license expiry date, for support and upgrade eligibility.
- Your server ID which is generated when you install Confluence for the first time and remains the same for the life of the installation (including after upgrades or changes to your license).
- Your support entitlement number (SEN).

Updating your license

If you change your license (for example to a license with more users), or migrate from Confluence Cloud and you will need to update your license.

To update your Confluence license:

1. Go to

   ⚙ ▾

   > **General Configuration** > **License Details**
2. Enter your new license in the **License** field.
3. Choose **Save**.

Understanding the user count for your license

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

The License Details page will indicate the number of users currently signed up (your registered user count). It:

Screenshot: License details

License Details

This page shows your current licensing information.
You can use the form below to update the license Confluence is running with.

Organisation	**Atlassian Docs**
Date Purchased	**Dec 08, 2013**
License Type	**Confluence: Developer**
	25 (5 signed up currently) [Refresh]
Support Period	Your commercial Confluence support and updates are available until **Dec 08, 2014**.
Support Entitlement Number	**SEN-L2180**
Server ID	**BC1S-1234-8T5O-ABCD** (Atlassian sales or support may ask you to provide this ID)

[Save]

- includes only users who have the 'can use' global permissions for the Confluence site.
- does not include anonymous users, who may access your Confluence site if you have allowed anonymous access.
- does not include deactivated users.

Exceeding your licensed user count

If you exceed the number of users included in your license, your Confluence instance will become read-only, that means no users will be able to create or edit content until you reduce the number of users.

Reducing your user count

You can reduce your user count by removing or deactivating users who do not require access to Confluence. See Delete or Disable Users.

If you have connected Confluence to an LDAP directory, you may want configure Confluence to only synchronise a subset of users from LDAP rather than all users. See How to change the number of users synchronized from LDAP to Confluence in the Knowledge Base. This can be a complicated process and we recommend that you only use this method if necessary.

Downgrading your license

If you decide to downgrade your Confluence license to pay for fewer users you need to ensure that the number of users currently signed up (as shown on the License Details page) is lower 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 Finding Your Support Entitlement Number for more general information about how Atlassian Support uses this number.

Managing Confluence Data

This page is an overview of recommended techniques for managing the data on your Confluence site. This is of interest to Confluence administrators – people with System Administrator or Confluence Administrator permissions.

- Database Configuration
- Site Backup and Restore
- Attachment Storage Configuration
- Confluence Data Model
- Finding Unused Spaces
- Data Import and Export
- Import a Text File

Related pages:

- Managing Add-ons or Plugins
- Integrating Confluence with Other Applications
- Getting Started as Confluence Administrator
- Confluence Administrator's Guide

Database Configuration

This document provides information on connecting

Confluence to an external database.

Choosing an external database

Note: Take time to choose your database wisely. The XML backup built into Confluence is not suited for migration or backup of large data sets. If you need to migrate later, you'll need to use a third party database migration tool.

Below is more information on selecting and migrating to an external database:

- Migrating to a Different Database
- Supported Databases
- Database Troubleshooting

About the embedded H2 database

Your Confluence installation includes an embedded H2 database, to enable you to try Confluence without setting up an external database. The embedded H2 database is **only** supported while you are evaluating Confluence. You must migrate to a supported external database before using Confluence as a production system.

To find out if you are still using the embedded database, go to

⚙ ▾

> **General Configuration** > **Support Tools**.

Database setup

Here are the setup instructions for the supported databases:

- Database Setup for Oracle
- Database Setup For MySQL
- Database Setup for PostgreSQL
- Database Setup for SQL Server

Database troubleshooting

For solving database-related problems:

- Troubleshooting External Database Connections
- How to Interpret DB2 Error Codes
- Database Troubleshooting

Obtain technical support from Troubleshooting Problems and Requesting Technical Support.

Notes

Issue CONF-12599 requests a more robust strategy for migrating large Confluence sites.

Database JDBC Drivers

This page provides the download links for the JDBC drivers for all databases currently supported for Confluence. You will need to make the driver available to your application server, as described in the appropriate setup guide .

Note: We bundle some JDBC drivers with Confluence, as shown below. If you are using a direct JDBC connection, you do not need to download or install the drivers that are bundled. If you are connecting via a datasource, or if you are using a database whose driver is not bundled, you will need to download and install the drivers manually.

> **Related pages:**
> - Database Configuration
> - Supported Platforms

Database	JDBC driver bundled with Confluence?	JDBC drivers	Notes	More information
PostgreSQL		**9.2-1004 JDBC 4 driver download**	The JDBC 4 driver will work under the 1.7 JVM. If you want to use other drivers, you can download it from the PostgreSQL website. However, we recommend that you use the bundled JDBC 4 driver.	Database Setup for PostgreSQL
Microsoft SQL Server		**jTDS 1.2.2 driver download**	The above version is the version bundled with Confluence. All our testing is done on that version. We do not know of any issues with later versions, so you are free to use them if you have tested them and find there are no issues in your environment. However, later versions are technically not supported. That means that if you do run into any problems, Atlassian Support may require you to move back to the above fully-tested version for troubleshooting.	Database setup for Microsoft SQL Server
MySQL		**Connector\J 5.1.30 driver download**	Note: In Confluence 5.1 and later, the MySQL drivers are no longer included in the Confluence distribution. For more information please refer to the Confluence 5.1 Upgrade Notes.	Database setup for MySQL
Oracle		**JDBC driver downloads**	• For Oracle 11.1, use the 11.1.0.7.0 driver (ojdbc6.jar). • For Oracle 11.2, use the 11.2.0.4.0 driver (ojdbc6.jar). • For Oracle 12c use the 12.1.0.1 driver (ojdbc7.jar) We recommend using the thin drivers only. See the Oracle JDBC driver FAQ.	Database setup for Oracle

Database Setup for Oracle

This page provides instructions for configuring Confluence to use an Oracle database. The setup process involves configuration of your Oracle server and your Confluence site.

Step 1. Check the prerequisites

Check the following before you start:

- Make sure your version of Oracle is supported. See Supported Platforms. If your version of Oracle is not supported, please

upgrade to a supported version before installing Confluence.

- If you have been evaluating Confluence and wish to transfer your data to a new database, consult the following guide first: Migrating to Another Database.
- If you are migrating from another database, consult the following guide first: Migrating to Another Database.
- **Note:** This database can only be set up by an Oracle database administrator (DBA). Oracle is difficult to set up. If you are not a DBA, and you do not have access to an experienced Oracle DBA, we recommend that you choose an alternative database. For a list of supported databases, see Supported Platforms. If you are evaluating Confluence, we recommend that you start with an alternative database and only consider migrating to Oracle after approval from your DBA. If you request Atlassian's technical support for difficulties with Oracle setup, we will assume you have the high level of skill required for an Oracle setup.

On this page:

- Step 1. Check the prerequisites
- Step 2. Install your Oracle server
- Step 3. Set up your Oracle user with schema-creation privileges
- Step 4. Install Confluence
- Step 5. Determine your JDBC URL
- Step 6. Download and install the Oracle database driver
- Step 7. Set up your database connection in the Confluence Setup Wizard
- Troubleshooting

Related pages:

- Database Configuration
- Known Issues for Oracle
- Confluence Installation and Upgrade Guide

Step 2. Install your Oracle server

If you do not already have an operational Oracle database server, download the installation package from the Oracle download page and follow the instructions in the Oracle documentation.

Then follow the steps below, to deploy Confluence to a schema in your Oracle server.

Note: Your database should be configured to use the same character encoding as Confluence. The recommended encoding is AL32UTF8 (the Oracle equivalent of Unicode UTF-8). See Configuring Database Character Encoding.

Step 3. Set up your Oracle user with schema-creation privileges

In this step you will create a Confluence user in Oracle and grant the appropriate roles to the user, so that the user can set up a connection, can create objects in its own schema, and can configure the schema.

To create the user and assign its privileges:

1. Access the command line interface to Oracle via the 'sqlplus' command.

    ```
    sqlplus user/password <as sysdba|as sysoper>
    ```

    You must add the 'as sysdba' or 'as sysoper' option if you are logging in with the user 'sys'. This determines which sys role you are using.
    Once logged in, you can type arbitrary SQL commands.
2. Create a Confluence user (<user>) in Oracle, and grant the appropriate roles only to the user:
    - connect role is required to set up a connection.
    - resource role is required to allow the user to create objects in its own schema.
    - Create table, sequence and trigger are required to configure the schema.

```
create user <user> identified by <password> default tablespace
<tablespace_name> quota unlimited on <tablespace_name>;
grant connect to <user>;
grant resource to <user>;
grant create table to <user>;
grant create sequence to <user>;
grant create trigger to <user>;
```

Notes:
- Do not grant the user the `select any table` permission. That permission can cause problems with other schemas. See the bug report CONF-3613.
- When you create a user, specify the `tablespace` for the table objects as shown above.

3. Add a local `all_objects` view to the user's schema, to prevent a conflict that can occur when a table exists in another schema with the same name as one of the Confluence tables. This is a workaround for the bug CONF-3613:

```
create view <user>.all_objects as
select *
from sys.all_objects
where owner = upper('<user>');
```

Step 4. Install Confluence

Install Confluence if you have not done so already. See the Confluence Installation Guide. **Stop immediately after the installation, before opening the Confluence Setup Wizard in your browser**, and follow the steps below.

If you have already got part-way through the Confluence Setup Wizard, stop at the database setup step and follow the steps below. You will be able to restart the setup wizard at the same step later.

Step 5. Determine your JDBC URL

The JDBC thin driver for Oracle use three different styles of URL:

- New style:

  New Style

  ```
  jdbc:oracle:thin:@//[HOST][:PORT]/SERVICE
  ```

- Old style:

  Old Style

  ```
  jdbc:oracle:thin:@[HOST][:PORT]:SID
  ```

- 'tnsnames' style:

tnsnames

```
jdbc:oracle:thin:@(DESCRIPTION=
                     (SDU=32768)
                     (enable=broken)
                     (LOAD_BALANCE=yes)
                     (FAILOVER=yes)
                     (ADDRESS=
                       (PROTOCOL=TCP)
                       (HOST=dbserver1.example.com)
                       (PORT=1525))
                     (ADDRESS=
                       (PROTOCOL=TCP)
                       (HOST=dbserver2.example.com)
                       (PORT=1525))
                     (CONNECT_DATA=
                       (SERVICE_NAME=CONFDB)))
```

Notes:

- The `tnsnames` style is required for connecting to an Oracle RAC cluster. For easy reading, we have split the example above over multiple lines, but you should compact it into a single line. These values may need more analysis than documented here, so you should seek the assistance of an experienced DBA.
- If you use the new style URL, then SERVICE can be either an SID or Service Name.
- If you use the old style URL, then SERVICE can only be the SID.

To determine the host, port, service name, and/or SID, execute the following command as the user running Oracle. (By default, the user is "oracle"):

```
lsnrctl status
```

For reference, here is a sample output:

```
SNRCTL for Linux: Version 11.2.0.2.0 - Beta on 29-JUN-2012 15:20:59
Copyright (c) 1991, 2010, Oracle.  All rights reserved.
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=EXTPROC_FOR_XE)))
STATUS of the LISTENER
-----------------------
Alias                     LISTENER
Version                   TNSLSNR for Linux: Version 11.2.0.2.0 - Beta
Start Date                06-JUN-2012 08:36:34
Uptime                    23 days 6 hr. 44 min. 25 sec
Trace Level               off
Security                  ON: Local OS Authentication
SNMP                      OFF
Default Service           XE
Listener Parameter File
/u01/app/oracle/product/11.2.0/xe/network/admin/listener.ora
Listener Log File
/u01/app/oracle/diag/tnslsnr/<HOSTNAME>/listener/alert/log.xml
Listening Endpoints Summary...
   (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC_FOR_XE)))
   (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=<HOSTNAME>)(PORT=1521)))

(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=<HOSTNAME>)(PORT=8080))(Presentation=HT
TP)(Session=RAW))
Services Summary...
Service "PLSExtProc" has 1 instance(s).
  Instance "PLSExtProc", status UNKNOWN, has 1 handler(s) for this service...
Service "XE" has 1 instance(s).
  Instance "XE", status READY, has 1 handler(s) for this service...
Service "XEXDB" has 1 instance(s).
  Instance "XE", status READY, has 1 handler(s) for this service...
The command completed successfully
```

Notes:

- The host and port are determined by the line containing `PROTOCOL=tcp`, **without** `Presentation=HTTP`.
- Under `Services Summary`, each service which has an instance with READY status is a connectable service. The name following `Service` is a service name for connecting to the database name following `Instance` on the next line.
- The SID is the name of the database instance, as defined by the `$ORACLE_SID` variable when you have sourced the Oracle environment to your shell.

For example, assuming that you are running Confluence on the same server as the Oracle database, with the above `lsnrctl` status output, you would use one of the following URLs:

```
jdbc:oracle:thin:@//localhost:1521/XE
jdbc:oracle:thin:@localhost:1521:XE
```

The URL can be used in either a direct JDBC connection or using a Tomcat datasource.

For further information on Oracle JDBC URLs, see the Oracle JDBC FAQ.

Step 6. Download and install the Oracle database driver

Decide whether you will set up a **direct JDBC connection or a datasource connection** to Oracle, to suit your environment. If unsure, choose direct JDBC.

To set up a direct JDBC connection:

If you plan to set up a direct JDBC connection to Oracle, you will need to copy the Oracle JDBC driver to your Confluence installation.

1. Download the latest compatible database driver. Links to the appropriate database drivers are available on this page: Database JDBC Drivers.
2. Copy the driver JAR file to the `<Confluence installation>/confluence/WEB-INF/lib` folder in your new Confluence installation.

To set up a datasource connection:

If you plan to set up a datasource connection to Oracle, follow the steps described in Configuring an Oracle Datasource in Apache Tomcat.

Step 7. Set up your database connection in the Confluence Setup Wizard

Start Confluence, and go to the Confluence Setup Wizard in your browser. Follow these steps to set up the new configuration:

1. Follow the initial steps in the Confluence Setup Guide.
2. When prompted to choose an evaluation or production installation, choose **production installation**.
3. When prompted to choose an embedded or external database, select **Oracle xx** from the dropdown list, where 'xx' is your Oracle version, and choose **External Database**.
4. Choose either the direct JDBC or the datasource connection, to suit the choice you made earlier when setting up the Oracle database driver.
   - For the JDBC connection: Enter the **database URL** to match the JDBC URL you determined in the previous section. Enter the **user name** (for example, `confluenceuser`) and **password** you chose when adding the Confluence database to Oracle.
   - For a datasource connection: Set the JNDI name to `java:comp/env/jdbc/confluence`.
5. Wait a while, as Confluence will create the schema in Oracle.

Congratulations! Confluence is now using your Oracle database to store its data.

Troubleshooting

- The following page contains common issues encountered when setting up your Oracle database to work with Confluence: Known Issues for Oracle.
- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If none of the above describes your issue, please create a support ticket at http://support.atlassian.com and be sure to include your logs (found in `<CONFLUENCE-INSTALLATION>/logs` and `<CONFLUENCE-HOME>/logs`).

Configuring an Oracle Datasource in Apache Tomcat

This page tells you how to set up an Oracle datasource connection for Confluence.

Step 1. Shut down Tomcat

1. Run `bin/shutdown.sh` or `bin/shutdown.bat` to bring Tomcat down while you are making these changes.
2. Make a backup of your `<CONFLUENCE_HOME>/confluence.cfg.xml` file and your `<CONFLUENCE_INSTALLATION>/conf/server.xml` file, so that you can easily revert if you have a problem.

Step 2. Install the Oracle database driver

1. Download the Oracle JDBC driver. Links are available on this page: Database JDBC Drivers.
2. Copy the driver JAR file into the `lib` folder of

On this page:
- Step 1. Shut down Tomcat
- Step 2. Install the Oracle database driver
- Step 3. Configure Tomcat
- Step 4. Configure the Confluence web application
- Step 5. Restart Tomcat

Related pages:

- Database Setup For Oracle
- Confluence Installation and Upgrade Guide

your Tomcat installation: `<TOMCAT-INSTALL
ATION>/lib`.

Step 3. Configure Tomcat

1. If you are using the Confluence distribution, edit the `conf/server.xml` file in your Tomcat installation. If you are running your own Tomcat instance, edit the XML file where you declared the Confluence Context descriptor.
2. Find the `Context` element in the `Host` element：

```
<Host name="localhost" debug="0" appBase="webapps" unpackWARs="true"
autoDeploy="false">
    <Context path="" docBase="../confluence" debug="0" reloadable="true">
        <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
        <Manager pathname="" />
    </Context>
</Host>
```

3. Insert the DataSource `Resource` element into the `Context` element, directly after the opening `<Cont
ext.../>` line, before `Manager,` as shown here:

```
<Host name="localhost" debug="0" appBase="webapps" unpackWARs="true"
autoDeploy="false">
    <Context path="" docBase="../confluence" debug="0" reloadable="true">
        <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
    <!-- If you're using Confluence 5.7 or below; change maxTotal to
maxActive -->
        <Resource
        name="jdbc/confluence"
        auth="Container"
        type="javax.sql.DataSource"
        driverClassName="oracle.jdbc.OracleDriver"
        url="jdbc:oracle:thin:@hostname:port:sid"
        username="<username>"
        password="<password>"
        connectionProperties="SetBigStringTryClob=true"
    accessToUnderlyingConnectionAllowed="true"
        maxTotal="25"
        maxIdle="10"
        maxWaitMillis="10000"
        />

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

4. Change the `username` and `password` to match your Oracle login.
5. Change the `url` to match the URL for your Oracle database. See how to find your Oracle URL. For example:

```
jdbc:oracle:thin:@example.atlassian.com:1521:confluencedb
```

6. If required, choose different `maxTotal` and `maxIdle` values. These define the number of database connections that will be allowed at one time, and the number that will be kept open even when there is no database activity.

Step 4. Configure the Confluence web application

Configure Confluence to use this datasource:

1. Edit this file in your Confluence installation: `<CONFLUENCE_INSTALLATION>/confluence/WEB-IN F/web.xml`.
2. Insert the following element just before `</web-app>` near the end of the file:

```
<resource-ref>
  <description>Connection Pool</description>
  <res-ref-name>jdbc/confluence</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
```

Step 5. Restart Tomcat

Run `bin/startup.sh` or `bin/startup.bat` to start Tomcat with the new settings.

Database Setup for SQL Server

This page provides instructions for configuring Confluence to use the Microsoft SQL Server database.

Step 1. Check the prerequisites

Check the following before you start:

- Check that your version of SQL Server is supported. See Supported Platforms. If your version is not supported, please upgrade to a supported version of SQL Server before installing Confluence.
- If you have been evaluating Confluence and wish to transfer your data to a new database, consult the following guide first: Migrating to Another Database.
- If you are migrating from another database, consult the following guide first: Migrating to Another Database.

On this page:

- Step 1. Check the prerequisites
- Step 2. Install SQL Server
- Step 3. Set up your SQL Server database and user
- Step 4. Install Confluence and the SQL Server database driver
- Step 5. Set up your database connection in the Confluence Setup Wizard
- Troubleshooting

Related pages:

- Database Configuration
- Known issues for SQL Server
- Confluence Installation and Upgrade Guide

Step 2. Install SQL Server

If you do not already have an operational SQL Server database, download the installation package from the Microsoft SQL Server download page and follow the instructions on MSDN.

Step 3. Set up your SQL Server database and user

In this step you will create a database within SQL Server to hold your Confluence data, and a database user with authority to access that database. The database user should be in the `db_owner` role.

1. Identify which character encoding to use. To do this, check the encoding currently used by your applic ation server and Confluence. All three must use compatible encoding. For example, the default SQL Server encoding of UCS-2 is compatible with UTF-8.
2. Using your SQL administrator permissions, create a new database in SQL Server.
3. If you set your application server and Confluence to use an encoding incompatible with UCS-2, specify that character encoding for the database.
4. Set the default collation for the database to be 'SQL_Latin1_General_CP1_CS_AS' (**case sensitive**). You can do this by issuing the following SQL query:

```
ALTER DATABASE <database_name> COLLATE SQL_Latin1_General_CP1_CS_AS
```

Note: if you receive an error stating 'The database could not be exclusively locked to perform the operation', you may need to prevent other connections by setting the mode to single user for the transaction:

```
ALTER DATABASE <database_name> SET SINGLE_USER WITH ROLLBACK IMMEDIATE;
<your ALTER DATABASE query>
ALTER DATABASE <database_name> SET MULTI_USER;
```

5. Configure the database to use the isolation level, 'Read Committed with Row Versioning'. You can do this by issuing the following SQL query:

   Determine if READ_COMMITTED_SNAPSHOT is enabled

   ```
   SELECT is_read_committed_snapshot_on FROM
   sys.databases WHERE name= 'YourDatabase'
   ```

   Return value:
   1 = READ_COMMITTED_SNAPSHOT option is ON. Read operations under the read-committed isolation level are based on snapshot scans and do not acquire locks.

   0 = READ_COMMITTED_SNAPSHOT option is OFF (default). Read operations under the read-committed isolation level use share locks.

   ```
   ALTER DATABASE <database_name>
       SET READ_COMMITTED_SNAPSHOT ON
       WITH ROLLBACK IMMEDIATE;
   ```

6. Using your SQL administrator permissions, create a new SQL user account for Confluence (for example, `confluenceuser`). Give this user full create, read and write permissions for the database tables. Note that Confluence must be able to create its own schema.

Step 4. Install Confluence and the SQL Server database driver

Decide whether you will set up a **direct JDBC connection or a datasource connection** to SQL Server, to suit your environment. If unsure, choose direct JDBC.

Install Confluence if you have not done so already. See the Confluence Installation Guide.

- If you plan to set up a **direct JDBC connection** to SQL Server, you can run the Confluence installation and move directly on to the Confluence Setup Wizard, as described below. The SQL Server JDBC driver is bundled with Confluence, as documented on this page: Database JDBC Drivers
.
- If you plan to set up a **datasource connection** to SQL Server:
  - Stop immediately after the Confluence installation, before opening the Confluence Setup Wizard in your browser. If you have already got part-way through the Confluence Setup Wizard, stop at the database setup step. You will be able to restart the setup wizard at the same step later.
  - Follow the steps described in Configuring a SQL Server Datasource in Apache Tomcat.

Step 5. Set up your database connection in the Confluence Setup Wizard

Start Confluence, and go to the Confluence Setup Wizard in your browser. Follow these steps to set up the new configuration:

1. Follow the initial steps in the Confluence Setup Guide.
2. When prompted to choose an evaluation or production installation, choose **production installation**.

3. When prompted to choose an embedded or external database, select **Microsoft SQL Server** from the dropdown list and choose **External Database**.
4. Choose either the direct JDBC or the datasource connection, to suit the choice you made earlier when setting up the SQL Server database driver.
   - For the JDBC connection:

     - When prompted for a **Driver Class Name,** enter the following:

       ```
       net.sourceforge.jtds.jdbc.Driver
       ```

     - When prompted for the **Database URL**, use this format:

       ```
       jdbc:jtds:sqlserver://<server>:<port>/<database>
       ```

       ℹ️ If MS SQL is clustered, use this format:

       ```
       jdbc:jtds:sqlserver://<server>:<port>/<database>;instance=<insta
       nce>
       ```

     - Enter the username (for example, `confluenceuser`) and password you chose earlier.
   - For a datasource connection: Set the **Datasource Name** to `java:comp/env/jdbc/conflue
     nce`

Congratulations! Confluence is now using your SQL Server database to store its data.

Troubleshooting

- If you get the following error message , verify that you have given the `confluenceuser` user all the required database permissions when connecting from `localhost`.

  ```
  Could not successfully test your database: : Server connection
  failure during transaction. Due to underlying exception:
  'java.sql.SQLException: Access denied for user
  'confluenceuser'@'localhost' (using password: YES)'
  ```

- The following page contains common issues encountered when setting up your SQL Server database to work with Confluence: Known Issues for SQL Server.
- Additional ports may be required to be opened. See this support document from Microsoft about the ports required for SQL Server.
- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If none of the above describes your issue, please create a support ticket at http://support.atlassian.co m and be sure to include your logs (found in `<CONFLUENCE-INSTALLATION>/logs` and `<CONFLUE
  NCE-HOME>/logs`).

Configuring a SQL Server Datasource in Apache Tomcat

This page tells you how to set up a SQL Server datasource connection for Confluence.

Step 1. Shut down Tomcat

1. Run `bin/shutdown.sh` or `bin/shutdown.
   bat` to bring Tomcat down while you are
   making these changes.
2. Make a backup of your `<CONFLUENCE_HOME`

`>/confluence.cfg.xml` file and your `<CON
FLUENCE_INSTALLATION>/conf/server
.xml` file, so that you can easily revert if you
have a problem.

Step 2. Install the SQL Server database driver

1. Download the SQL Server JDBC driver:
   - Links are available on this page: Datab
     ase JDBC Drivers.
   - Unpack the archive file you have
     downloaded, and find the JAR file
     called something like this: `jtds-x.x.
     x.jar`, where `x.x.x` is a version
     number.
2. Alternatively, you can get the driver from your
   Confluence installation: `/confluence/WEB-
   INF/lib/jtds-x.x.x.jar`.
3. Put the JAR file into the `lib` folder of your
   Tomcat installation: `<TOMCAT-INSTALLATIO
   N>/lib`.

Step 3. Configure Tomcat

1. Edit the `conf/server.xml` file in your Tomcat installation.
2. Find the following lines:

```
<Context path="" docBase="../confluence" debug="0" reloadable="true">
        <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
```

3. Insert the DataSource `Resource` element inside the `Context` element, directly after the opening `<Co
   ntext.../>` line, before `Manager`:

```
<!-- If you're using Confluence 5.7 or below; change maxTotal to maxActive
-->
<Resource name="jdbc/confluence" auth="Container"
type="javax.sql.DataSource"
    username="yourDatabaseUser"
    password="yourDatabasePassword"
    driverClassName="net.sourceforge.jtds.jdbc.Driver"
    url="jdbc:jtds:sqlserver://localhost:1433/yourDatabaseName"
    maxTotal="25"
    maxIdle="10"
    validationQuery="select 1" />
```

   - Replace the `username` and `password` parameters with the correct values for your database
   - In the `url` parameter, replace the word '`yourDatabaseName`' with the name of the database
     your Confluence data will be stored in.

Notes:

- If switching from a direct JDBC connection to datasource, you can find the above details in your `<CON
  FLUENCE_HOME>/confluence.cfg.xml` file.
- Here are the configuration properties for Tomcat's standard data source resource factory (`org.apach
  e.tomcat.dbcp.dbcp.BasicDataSourceFactory`):
   - driverClassName — Fully qualified Java class name of the JDBC driver to be used.
   - maxTotal — The maximum number of active instances that can be allocated from this pool at
     the same time.
   - maxIdle — The maximum number of connections that can sit idle in this pool at the same time.
   - maxWaitMillis — The maximum number of milliseconds that the pool will wait (when there are

no available connections) for a connection to be returned before throwing an exception.
- password — Database password to be passed to our JDBC driver.
- url — Connection URL to be passed to our JDBC driver. (For backwards compatibility, the property driverName is also recognized.)
- user — Database username to be passed to our JDBC driver.
- validationQuery — SQL query that can be used by the pool to validate connections before they are returned to the application. If specified, this query MUST be an SQL SELECT statement that returns at least one row.

- Why is the `validationQuery` element needed? When a database server reboots, or there is a network failure, all the connections in the connection pool are broken and this normally requires a Application Server reboot. However, the Commons DBCP (Database Connection Pool) which is used by the Tomcat application server can validate connections before issuing them by running a simple SQL query, and if a broken connection is detected, a new one is created to replace it. To do this, you will need to set the "validationQuery" option on the database connection pool.

Step 4. Configure the Confluence web application

1. Edit this file in your Confluence installation: `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`.
2. Insert the following element just before `</web-app>` near the end of the file:

```
<resource-ref>
    <description>Connection Pool</description>
    <res-ref-name>jdbc/confluence</res-ref-name>
    <res-type>javax.sql.DataSource</res-type>
    <res-auth>Container</res-auth>
</resource-ref>
```

If you are changing an existing Confluence installation over to using a Tomcat datasource:

1. Edit the `<CONFLUENCE_HOME>/confluence.cfg.xml` file.
2. Delete any line that contains a property that begins with `hibernate.`
3. Insert the following at the start of the `<properties>` section.

```
<property name="hibernate.setup">true</property>
<property
name="hibernate.dialect">net.sf.hibernate.dialect.SQLServerIntlDialect</property>
<property
name="hibernate.connection.datasource">java:comp/env/jdbc/confluence</property>
```

Step 5. Restart Tomcat

Run `bin/startup.sh` or `bin/startup.bat` to start Tomcat with the new settings.

Database Setup For MySQL

This page provides instructions for configuring Confluence to use the MySQL database.

Step 1. Check the prerequisites

Check the following before you start:

- Check that your version of MySQL is supported. See Supported Platforms.
  - ⌄ Note about MariaDB and Percona Server...
    We do not currently support MariaDB or Percona Server. Both are known to have issues with Confluence. See

CONF-36471			CONF-29060
NEW	and		OPEN

for more information.

- If you have been evaluating Confluence using the embedded database and wish to transfer your data to a new database, consult the following guide first: Migrating to Another Database.
- If you are migrating from another database, consult the following guide first: Migrating to Another Database.

Step 2. Install MySQL Server

If you do not already have an operational MySQL database server, install 'MySQL Community Edition'. Download the installation package from the MySQL download page and follow the instructions in the My SQL documentation.

On this page:
- Step 1. Check the prerequisites
- Step 2. Install MySQL Server
- Step 3. Configure MySQL Server
- Step 4. Set up your MySQL database and user
- Step 5. Install Confluence
- Step 6. Download and install the MySQL database driver
- Step 7. Check settings for internationalisation
- Step 8. Set up your database connection in the Confluence Setup Wizard
- Troubleshooting

Related pages:

- Configuring Database Character Encoding
- Database Configuration
- Known issues for MySQL
- Confluence Installation and Upgrade Guide

Step 3. Configure MySQL Server

In this step, you will configure your MySQL database server.

Note: If you intend to connect Confluence to an existing MySQL database server, we strongly recommend that you reconfigure this database server by running through the configuration steps in the MySQL installation wizard as described below .

To configure MySQL Server:

1. Run the MySQL installation wizard:
   a. If you are connecting Confluence to your existing MySQL server, choose **Reconfigure Instance**.
   b. Choose **Advanced Configuration**.
   c. Choose the **type of MySQL Server** that best suits your hardware requirements. This will affect the MySQL Server's usage of memory, disk and CPU resources. Refer to the MySQL documentation for further information.
   d. Choose **Transactional Database Only** to ensure that your MySQL database will use **InnoDB** a s its default storage engine.
   *It is **highly recommended** that you only use the InnoDB storage engine with Confluence. Avoid using the MyISAM storage engine as this can lead to data corruption.*
   e. Set the **InnoDB Tablespace** settings to your requirements. (The default settings are acceptable.)
   f. Set the approximate **number of concurrent connections** permitted to suit your Confluence usage requirements. You can use one of the presets or enter a number manually. Refer to the MySQL documentation for further information.
   g. For the **networking options**, ensure the **Enable TCP/IP Networking** and **Enable Strict Mode** options are selected (default). Refer to the MySQL documentation on setting the networking an d server SQL modes for further information.
   h. For the MySQL server's **default character set**, choose **Best Support For Multilingualism** (in other words, UTF-8). This will ensure Confluence's support for internationalisation. For more information, see Configuring Database Character Encoding.
   i. For the Windows configuration option, choose whether or not to install the MySQL Server as a Windows service. If your hardware is going to be used as a dedicated MySQL Server, you may wish to choose the options to **Install As Windows Service** (and **Launch the MySQL Server automatically**). Refer to the MySQL documentation for further information.
   Note: If you choose not to install the MySQL Server as a Windows Service, you will need to ensure that the database service has been started before running Confluence.

j. Select **Modify Security Settings** to enter and set your MySQL Server (root) access password.
2. Edit the `my.cnf` file (`my.ini` on Windows operating systems) in your MySQL server. Locate the `[my sqld]`section in the file, and add or modify the following parameters:
   (Refer to MySQL Option Files for detailed instructions on editing `my.cnf` and `my.ini`.)
   Locate the `[mysqld]`section in the file, and add or modify the following parameters:
   - Specify the default character set to be UTF-8:

```
[mysqld]
...
character-set-server=utf8
collation-server=utf8_bin
...
```

   - Set the default storage engine to InnoDB:

```
[mysqld]
...
default-storage-engine=INNODB
...
```

   - Specify the value of `max_allowed_packet` to be at least 256M:

```
[mysqld]
...
max_allowed_packet=256M
...
```

   - Specify the value of `innodb_log_file_size` to be at least 256M for MySQL 5.5 and below:

```
[mysqld]
...
innodb_log_file_size=256M
...
```

   NB: This should be set to at least 2G for MySQL 5.6 and above.
   - Ensure the sql_mode parameter does not specify NO_AUTO_VALUE_ON_ZERO

```
// remove this if it exists
sql_mode = NO_AUTO_VALUE_ON_ZERO
```

3. Restart your MySQL server for the changes to take effect:
   - On Windows, use the Windows Services manager to restart the service.
   - On Linux:
     - Run one of the following commands, depending on your setup: `'/etc/init.d/mysqld stop'` or `'/etc/init.d/mysql stop'` or `'service mysqld stop'`.
     - Then run the same command again, replacing `'stop'` with `'start'`.
   - On Mac OS X, run `'sudo /Library/StartupItems/MySQLCOM/MySQLCOM restart'`.

Step 4. Set up your MySQL database and user

In this step you will create a database within MySQL to hold your Confluence data, and a database user with authority to access that database.

To create the database and user privileges:

1. Run the '`mysql`' command as a MySQL super user. The default user is 'root' with a blank password.
2. Create an empty Confluence database schema by running this command:

```
CREATE DATABASE confluence CHARACTER SET utf8 COLLATE utf8_bin;
```

3. Create the Confluence database user by running this command. Replace 'confluenceuser' and 'confluencepass' with a username and password of your choice. If Confluence is not running on the same server as your MySQL database server, replace 'localhost' with the hostname or IP address of the Confluence server:

```
GRANT ALL PRIVILEGES ON confluence.* TO 'confluenceuser'@'localhost'
IDENTIFIED BY 'confluencepass';
```

Step 5. Install Confluence

Install Confluence if you have not done so already. See the Confluence Installation Guide. **Stop immediately after the installation, before opening the Confluence Setup Wizard in your browser**, and follow the steps below.

If you have already got part-way through the Confluence Setup Wizard, stop at the database setup step and follow the steps below. You will be able to restart the setup wizard at the same step later.

Step 6. Download and install the MySQL database driver

If you are **upgrading Confluence to a later version, and you are already using the recommended MySQL driver** (JDBC Connector/J 5.1), you can skip the instructions in this section. The Confluence upgrade task will automatically copy over your existing driver to the upgraded installation.

If you are installing Confluence, or you are upgrading Confluence and not using the recommended MySQL driver (JDBC Connector/J 5.1), follow the steps below.

Choose whether you will set up a **direct JDBC connection or a datasource connection** to MySQL, to suit your environment. If unsure, choose direct JDBC.

To set up a direct JDBC connection:

If you plan to set up a direct JDBC connection to MySQL, you will need to copy the MySQL JDBC driver to your Confluence installation.

1. Get the MySQL driver:
   - If you are **installing Confluence**, download the recommended MySQL driver . Links to the appropriate database drivers are available on this page: Database JDBC Drivers.
   You can download either the `.tar.gz` or the `.zip` archive. Extract the driver JAR file (for example, `mysql-connector-java-x.x.x-bin.jar`, where `x.x.x` is a version number) from the archive.
   - If you are **upgrading Confluence to a later version, and you are not using the recommended MySQL driver** (JDBC Connector/J 5.1), copy the driver JAR file from your existing Confluence installation before you upgrade. The driver will be in the `<Confluence installation> /confluence/WEB-INF/lib` folder.
2. Copy the driver JAR file to the `<Confluence installation> /confluence/WEB-INF/lib` folder in your new or upgraded Confluence installation.

To set up a datasource connection:

If you plan to set up a datasource connection to MySQL, follow the steps described in Configuring a MySQL Datasource in Apache Tomcat.

Step 7. Check settings for internationalisation

If you are using a existing database, use the **status** command to verify database character encoding information. The results should be UTF-8. See Configuring Database Character Encoding.

Step 8. Set up your database connection in the Confluence Setup Wizard

Start Confluence, and go to the Confluence Setup Wizard in your browser. Follow these steps to set up the new configuration:

1. Follow the initial steps in the Confluence Setup Guide.
2. When prompted to choose an evaluation or production installation, choose **production installation**.
3. When prompted to choose an embedded or external database, select **MySQL** from the dropdown list and choose **External Database**.
   (**Note:** you'll see a warning that a driver is required. You downloaded or copied the driver in step 6 so you can ignore this warning)
   Choose either the direct JDBC or the datasource connection, to suit the choice you made earlier when setting up the MySQL database driver.
      - For the JDBC connection: Enter the username (for example, `confluenceuser`) and password you chose earlier.
      - For a datasource connection: Set the JNDI name to `java:comp/env/jdbc/confluence`

Congratulations! Confluence is now using your MySQL database to store its data.

Troubleshooting

- If you get the following error message , verify that you have given the `confluenceuser` user all the required database permissions when connecting from `localhost`.

  ```
  Could not successfully test your database: : Server connection
  failure during transaction. Due to underlying exception:
  'java.sql.SQLException: Access denied for user
  'confluenceuser'@'localhost' (using password: YES)'
  ```

- The following page contains common issues encountered when setting up your MySQL database to work with Confluence: Known issues for MySQL
- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If none of the above describes your issue, please create a support ticket at http://support.atlassian.com and be sure to include your logs (found in `<CONFLUENCE-INSTALLATION>/logs` and `<CONFLUENCE-HOME>/logs`).
- There's currently a known issue for MySQL 5.5
  CONF-37829 - Improve select query that used in Confluence for copying a page
  RESOLVED

Configuring a MySQL Datasource in Apache Tomcat

This page tells you how to set up a MySQL datasource connection for Confluence.

Step 1. Shut down Tomcat

1. Run `bin/shutdown.sh` or `bin/shutdown.bat` to bring Tomcat down while you are making these changes.
2. Make a backup of your `<CONFLUENCE_HOME>/confluence.cfg.xml` file and your `<CONFLUENCE_INSTALLATION>/conf/server.xml` file, so that you can easily revert if you have a problem.

Step 2. Install the MySQL database driver

1. Download the MySQL JDBC driver. Links are available on this page: Database JDBC Drivers.
2. Unpack the archive file you have downloaded,

On this page:

- Step 1. Shut down Tomcat
- Step 2. Install the MySQL database driver
- Step 3. Configure Tomcat
- Step 4. Configure the Confluence web application
- Step 5. Restart Tomcat

Related pages:

- Database Setup For MySQL
- Confluence Installation and Upgrade Guide

and find the JAR file called something like
this: `mysql-connector-java-x.x.x-bin`
`.jar`, where `x.x.x` is a version number.
3. Copy the JAR file into the `lib` folder of your
Confluence installation: `<CONFLUENCE_INST`
`ALLATION>/lib`.

Step 3. Configure Tomcat

1. Edit the `conf/server.xml` file in your Tomcat installation.
2. Find the following lines:

```
<Context path="" docBase="../confluence" debug="0" reloadable="true">
     <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
```

3. Insert the DataSource `Resource` element within the `Context` element, directly after the opening `<Co`
`ntext.../>` line, before `Manager`:

```
<!-- If you're using Confluence 5.7 or below; change maxTotal to maxActive
-->
<Resource name="jdbc/confluence" auth="Container"
type="javax.sql.DataSource"
         username="yourusername"
         password="yourpassword"
         driverClassName="com.mysql.jdbc.Driver"

url="jdbc:mysql://localhost:3306/confluence?useUnicode=true&amp;characterEn
coding=utf8"
         maxTotal="25"
         maxIdle="10"
         defaultTransactionIsolation="READ_COMMITTED"
         validationQuery="Select 1" />
```

- Replace the `username` and `password` parameters with the correct values for your database.
- In the url parameter, replace the word '`confluence`' with the name of the database your
  Confluence data will be stored in.

Notes

- If switching from a direct JDBC connection to a datasource connection, you can find the above details
  in your `<CONFLUENCE_HOME>/confluence.cfg.xml` file.

- The configuration properties for Tomcat's standard datasource resource factory (`org.apache.tomca`
  `t.dbcp.dbcp.BasicDataSourceFactory`) are as follows:
    - `driverClassName` – Fully qualified Java class name of the JDBC driver to be used.
    - `maxTotal` – The maximum number of database connections in the pool at the same time.
    - `maxIdle` – The maximum number of connections that can sit idle in this pool at the same time.
    - `maxWaitMillis` – The maximum number of milliseconds that the pool will wait (when there
      are no available connections) for a connection to be returned before throwing an exception.
    - `password` – Database password to be passed to your JDBC driver.
    - `url` – Connection URL to be passed to your JDBC driver. (For backwards compatibility, the
      property `driverName` is also recognised.)
    - `user` – Database username to be passed to your JDBC driver.
    - `validationQuery` – SQL query that can be used by the pool to validate connections before
      they are returned to the application. If specified, this query must be an SQL SELECT statement
      that returns at least one row.
- Why is the `validationQuery` element needed? When a database server reboots, or there is a
  network failure, all the connections in the connection pool are broken and this normally requires an
  application server reboot. However, the Commons DBCP (Database Connection Pool) which is used

by the Tomcat application server can validate connections before issuing them by running a simple SQL query, and if a broken connection is detected, a new one is created to replace it. To do this, you will need to set the `validationQuery` option on the database connection pool.

Step 4. Configure the Confluence web application

1. Edit this file in your Confluence installation: `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`.
2. Insert the following element just before `</web-app>` near the end of the file:

```
<resource-ref>
    <description>Connection Pool</description>
    <res-ref-name>jdbc/confluence</res-ref-name>
    <res-type>javax.sql.DataSource</res-type>
    <res-auth>Container</res-auth>
</resource-ref>
```

If you are changing an existing Confluence installation over to using a Tomcat datasource:

1. Edit the `<CONFLUENCE_HOME>/confluence.cfg.xml` file.
2. Delete any line that contains a property that begins with `hibernate.`
3. Insert the following at the start of the `<properties>` section.

```
<property name="hibernate.setup"><![CDATA[true]]></property>
<property
name="hibernate.dialect"><![CDATA[com.atlassian.hibernate.dialect.MySQLDialect]]></property>
<property
name="hibernate.connection.datasource"><![CDATA[java:comp/env/jdbc/confluence]]></property>
```

Step 5. Restart Tomcat

Run `bin/startup.sh` or `bin/startup.bat` to start Tomcat with the new settings.

Database Setup for PostgreSQL

This page provides instructions for configuring Confluence to use a PostgreSQL database.

Step 1. Check the prerequisites

Check the following before you start:

- Check that your version of PostgreSQL is supported. See Supported Platforms. If your version is not supported, please upgrade to a supported version of PostgreSQL before installing Confluence.
- If you have been evaluating Confluence and wish to transfer your data to a new database, consult the following guide first: Migrating to Another Database.
- If you are migrating from another database, consult the following guide first: Migrating to Another Database.

Step 2. Install PostgreSQL

If you do not already have an operational PostgreSQL database, install it now.

1. Download and install PostgreSQL. Please note the following information when installing PostgreSQL:
   - The **password** that you are prompted to provide during the installation process is for the 'postgres' account, which is the database root-level account, sometimes called the super user ('postgres'). Remember this username and password. You will need it each time you log in to the database.
   - The **default port** for PostgreSQL is 5432. If you decide to change the default port, please ensure that your new port number does not conflict with any services running on that port. You will also need to remember to update all further mentions of the database port.
   - Choose the **locale** that best fits your geographic location.
   - Do **not** launch **Stack Builder** at the completion of the installer.

Step 3. Set up your PostgreSQL database and user

Next you need to create a database within PostgreSQL to hold your Confluence data, and a database user with authority to access that database.

1. Create a database user (for example `confluenceuser`)
   - Your new user must be able to **create database objects** and **create roles**.
2. Create a database (for example `confluence`)
   - **Owner** is your new database user (for example `confluenceuser`)
   - **Character encoding** should be `utf8` encoding.

You can use pgAdmin as an alternative to the command line to complete this step. If you used the graphical installer when installing PostgreSQL, pgAdmin will be already installed on your computer.

Step 4. Install Confluence and the PostgreSQL database driver

Decide whether you will set up a **direct JDBC connection or a datasource connection** to PostgreSQL, to suit your environment. If unsure, choose direct JDBC.

Install Confluence if you have not done so already. See the Confluence Installation Guide.

- If you plan to set up a **direct JDBC connection** to PostgreSQL, you can run the Confluence installation and move directly on to the Confluence Setup Wizard, as described below. The PostgreSQL JDBC driver is bundled with Confluence, as documented on this page: Database JDBC Drivers.
- If you plan to set up a **datasource connection** to PostgreSQL:
  - Stop immediately after the Confluence installation, before opening the Confluence Setup Wizard in your browser. If you have already got part-way through the Confluence Setup Wizard, stop at the database setup step. You will be able to restart the setup wizard at the same step later.
  - Follow the steps described in Configuring a PostgreSQL Datasource in Apache Tomcat.

Step 5. Set up your database connection in the Confluence Setup Wizard

Start Confluence, go to the Confluence Setup Wizard in your browser, and follow these steps:

1. When prompted to choose an evaluation or production installation, choose **production installation**.
2. When prompted to choose an embedded or external database, select **PostgreSQL** from the dropdown

list and choose **External Database**.
3. Choose either the **direct JDBC or the datasource connection**, to suit the choice you made earlier.
   - **For the JDBC connection**:
     - When prompted for a **Driver Class Name,** enter:

       ```
       org.postgresql.Driver
       ```

     - When prompted for the **Database URL**, use this format:

       ```
       jdbc:postgresql://<server>:<port>/<database>
       ```
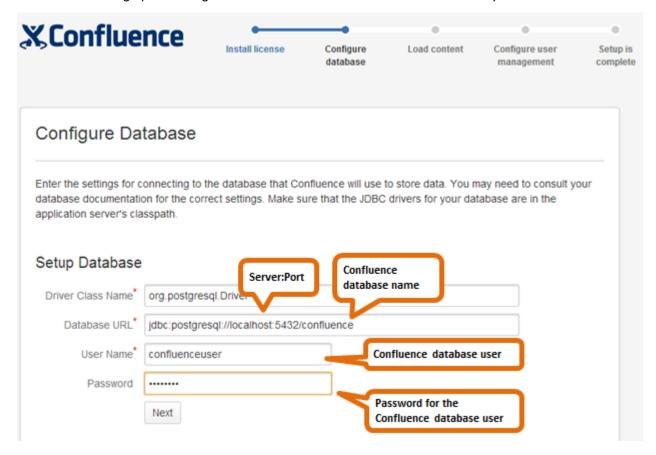
       For example: `jdbc:postgresql://localhost:5432/confluence`

       Note: If you need to connect to an SSL database, add the `ssl=true` parameter in the database URL. For example: `jdbc:postgresql://localhost:5432/confluence?ssl=true`
     - Enter the username (for example `confluenceuser`) and password you chose earlier when setting up your Confluence database.
   - **For a datasource connection**: Set the **Datasource Name** to the following: `java:comp/env/jdbc/confluence`

That's it - Confluence is now using your PostgreSQL database to store its data.

Screenshot: Setting up the PostgreSQL JDBC connection in the Confluence Setup Wizard

Notes

- If the server that is hosting the PostgreSQL database is not the same server as Confluence, then please ensure that the Confluence server can contact the database server. Please also refer to the PostgreSQL documentation on how to set up pg_hba.conf. If the `pg_hba.conf` file is not set properly, remote communication to the PostgresSQL server will fail.
- Running SQL queries: For ongoing maintenance of your server, you can continue to use PGAdmin III as

your SQL browser.

Troubleshooting

- If you get the following error message, verify that you have given the `confluenceuser` user all the required database permissions when connecting from `localhost`.

  ```
  Could not successfully test your database: : Server connection
  failure during transaction. Due to underlying exception:
  'java.sql.SQLException: Access denied for user
  'confluenceuser'@'localhost' (using password: YES)'
  ```

- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If you are unable to connect to the database from Confluence and they are on different machines, most likely you have a firewall in between the two machines or your `pg_hba.conf` file is misconfigured. Verify that your firewall is set to allow connections through 5432 or double check your hba configuration.
- The following page contains common issues encountered when setting up your PostgreSQL database to work with Confluence: Known issues for PostgreSQL.
- If none of the above describes your issue, please create a support ticket at http://support.atlassian.com a nd be sure to include your logs (found in `<CONFLUENCE-INSTALLATION>/logs` and `<CONFLUENCE-HOME>/logs`).

Configuring a PostgreSQL Datasource in Apache Tomcat

This page tells you how to set up a PostgreSQL datasource connection for Confluence.

Step 1. Shut down Tomcat

1. Run `bin/shutdown.sh` or `bin/shutdown.bat` to bring Tomcat down while you are making these changes.
2. Make a backup of your `<CONFLUENCE_HOME>/confluence.cfg.xml` file and your `<CONFLUENCE_INSTALLATION>/conf/server.xml` file, so that you can easily revert if you have a problem.

> **On this page:**
>
> - Step 1. Shut down Tomcat
> - Step 2. Install the PostgreSQL Server database driver
> - Step 3. Configure Tomcat
> - Step 4. Configure the Confluence web application
> - Step 5. Restart Tomcat
>
> **Related pages:**
>
> - Database Setup for PostgreSQL

Step 2. Install the PostgreSQL Server database driver

1. Download the PostgreSQL Server JDBC driver JAR file.
   - Links are available on this page: Database JDBC Drivers.
   - Alternatively, you can get the driver from your Confluence installation: `/confluence/WEB-INF/lib/postgresql-x.x-x.jdbcx.jar`, where 'x' represents a version number.
2. Copy the JAR file into the `lib` folder of your Tomcat installation: `<TOMCAT-INSTALLATION>/lib`.

Step 3. Configure Tomcat

1. Edit the `conf/server.xml` file in your Tomcat installation.
2. Find the following lines:

```
<Context path="" docBase="../confluence" debug="0" reloadable="true">
        <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
```

3. Insert the DataSource `Resource` element inside the `Context` element, directly after the opening `<Context.../>` line, before `Manager`:

```
<!-- If you're using Confluence 5.7 or below; change maxTotal to maxActive
-->
<Resource name="jdbc/confluence" auth="Container"
type="javax.sql.DataSource"
            username="postgres"
            password="postgres"
            driverClassName="org.postgresql.Driver"
            url="jdbc:postgresql://localhost:5432/yourDatabaseName"
            maxTotal="25"
            maxIdle="10"
            validationQuery="select 1" />
```

- Replace the `username` and `password` parameters with the correct values for your database
- In the `url` parameter, replace the word '`yourDatabaseName`' with the name of the database your Confluence data will be stored in.

Notes:

- If switching from a direct JDBC connection to datasource, you can find the above details in your `<CONFLUENCE_HOME>/confluence.cfg.xml` file.
- Here are the configuration properties for Tomcat's standard data source resource factory (`org.apache.tomcat.dbcp.dbcp.BasicDataSourceFactory`):
    - driverClassName — Fully qualified Java class name of the JDBC driver to be used.
    - maxTotal — The maximum number of database connections in the pool at the same time.
    - maxIdle — The maximum number of connections that can sit idle in this pool at the same time.
    - maxWaitMillis — The maximum number of milliseconds that the pool will wait (when there are no available connections) for a connection to be returned before throwing an exception.
    - password — Database password to be passed to our JDBC driver.
    - url — Connection URL to be passed to our JDBC driver. (For backwards compatibility, the property driverName is also recognized.)
    - user — Database username to be passed to our JDBC driver.
    - validationQuery — SQL query that can be used by the pool to validate connections before they are returned to the application. If specified, this query MUST be an SQL SELECT statement that returns at least one row.
- Why is the `validationQuery` element needed? When a database server reboots, or there is a network failure, all the connections in the connection pool are broken and this normally requires a Application Server reboot. However, the Commons DBCP (Database Connection Pool) which is used by the Tomcat application server can validate connections before issuing them by running a simple SQL query, and if a broken connection is detected, a new one is created to replace it. To do this, you will need to set the "validationQuery" option on the database connection pool.

Step 4. Configure the Confluence web application

1. Edit this file in your Confluence installation: `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`.
2. Insert the following element just before `</web-app>` near the end of the file:

```
<resource-ref>
    <description>Connection Pool</description>
    <res-ref-name>jdbc/confluence</res-ref-name>
    <res-type>javax.sql.Datasource</res-type>
    <res-auth>Container</res-auth>
</resource-ref>
```

If you are changing an existing Confluence installation over to using a Tomcat datasource:

1. Edit the `<CONFLUENCE_HOME>/confluence.cfg.xml` file.
2. Delete any line that contains a property that begins with `hibernate.`
3. Insert the following at the start of the `<properties>` section:

```
<property name="hibernate.setup"><![CDATA[true]]></property>
<property
name="hibernate.dialect"><![CDATA[net.sf.hibernate.dialect.PostgreSQLDialec
t]]></property>
<property
name="hibernate.connection.datasource"><![CDATA[java:comp/env/jdbc/confluen
ce]]></property>
```

Step 5. Restart Tomcat

Run `bin/startup.sh` or `bin/startup.bat` to start Tomcat with the new settings.

Embedded H2 Database

To enable you to try Confluence without setting up an external database, your Confluence installation includes an embedded H2 database.

The embedded H2 database is used by default when you choose the **Trial** installation path.

Related pages:

- Confluence Home and other important directories
- Database Configuration

The embedded database files are stored in your Confluence home directory `<confluence-home>/database`.

The embedded H2 database is **only** supported while you are evaluating Confluence. You must migrate to a supported external database before using Confluence as a production system.

To find out if you are still using the embedded database, go to

⚙ ▾

> **General Configuration** > **Support Tools**.

Connect to the embedded H2 database using DB Visualizer

DBVisualizer is just one database administration tool. You can use any administration tool that supports embedded H2 databases. The steps will be similar.

1. Shut down Confluence.
2. Back up your `<confluence-home>/database` directory.
3. Launch DBVisualizer.
4. Choose **Create new database connection** and follow the prompts to set up the connection. The information you'll need is:
   - **Database driver:** H2 embedded
   - **Database Userid:** sa
   - **Database password:** leave this field blank
   - **Database filename:**`<confluence-home>/database/h2db`
     ⚠ leave off the `.h2.db` file extension.
5. Connect to the database.

Refer to the DBVisualizer documentation for help using DBVisualizer.

Connect to the embedded H2 database using the H2 console

Alternatively you can connect using the browser based H2 console. The easiest way to access the console is to double click the H2 database jar file at `<installation-directory>\confluence\WEB-INF\lib\h2-x.x.x.jar`.

Migrating to Another Database

This document describes how to migrate your Confluence data from your existing database to another database. The instructions are designed primarily for migrating from an evaluation to a production database. Large data sets will require third party database migration tools.

This page covers the following scenarios:

- Moving from the embedded, trial database to a supported external database.
- Moving from one external database to another, for example from Oracle to PostgreSQL (provided your dataset is not large)
- Upgrading to a new version of the same external database. Note: you don't need to migrate your data if you're upgrading the database in place.

Note: If you are moving your database from one server to another you can change the JDBC URL in `<confluence.home>/confluence.cfg.xml` (if you are using a direct JDBC connection) or in the definition of your datasource (if you are connecting via a datasource).

On this page:

Related pages:

Limitations of database migration

Note: The XML export built into Confluence is not suited for the backup or migration of large data sets. There are a number of third party tools that may be able to assist you with the data migration. If you would like help in selecting the right tool, or help with the migration itself, we can put you in touch with one of the Atlassian Experts.

Database migration

There are two ways you can perform the migration, both described on this page:

1. **Method one** is the standard procedure.
2. Use **method two** if the total size of attachments in your installation exceeds 500MB.

Method one – standard procedure

Step 1: Take note of your add-ons

Take note of the add-ons (plugins) currently installed and enabled in Confluence, so that you can reinstate them later. Make a note of the following for each add-on:

- Add-on name
- Version
- Enabled or disabled status. This is useful if you have enabled or disabled modules yourself, making your configuration differ from the default.

Step 2: Back up your data

1. Create an XML backup of your existing data, via the **Confluence administration console**. See Manually Backing Up the Site. Make a note of the location where you put the XML file. You will need it later to import your Confluence data into your new database.
2. Shut down Confluence.
3. Make a copy of the Confluence Home directory. This is a precautionary measure, to ensure you can recover your data if it is mistakenly overwritten.
4. If you are using an external database, make a separate backup using the utilities that were installed with that database. This also is a precautionary measure.

Step 3: Set up the new database

Choose the database setup instructions for your new database, and follow those instructions to do the following:

- Install the database server.
- Perform any required configuration of the database server, as instructed.
- Add the Confluence database and user. Make a note of the username and password that you define in this step. You will need them later, when running the Confluence Setup Wizard.

Step 4. Install Confluence (same version number) in a new location

Now you will install Confluence again, with a different home directory path and installation path.

Note: You must use the same version of Confluence as the existing installation. (If you want to upgrade Confluence, you must do it as a separate step.) For example, if your current site is running Confluence 5.1.2, your new installation must also be Confluence 5.1.2.

When running the Confluence installer:

- Choose **Custom Install**. (Do **not** choose to upgrade your existing installation.)
- Choose a **new destination directory**. This is the installation directory for your new Confluence. It must not be the same as the existing Confluence installation.
- Choose a **new home directory**. This is the data directory for your new Confluence. It must not be the same as the existing Confluence installation.

Step 5. Download and install the database driver if necessary

Note that Confluence bundles some database drivers, but you'll need to install the driver yourself if it is not bundled. Follow the database setup instructions for your new database, to download and install the database driver if necessary.

Step 6. Run the Confluence setup wizard and copy your data to your new database

When running the Confluence setup wizard:

- Enter your license key, as usual.
- Choose **Production Installation** as the installation type.
- In the database configuration step, choose your new database type from the dropdown menu, then choose **External Database**.
- Choose the connection type: **Direct JDBC** or **Datasource**. If you are not sure which, choose 'Direct JDBC'. This is the most common connection type.
- When prompted for the database **user and password**, supply the credentials you defined earlier when adding the Confluence database to your database server.
- On the load content step, choose **Restore From Backup**. This is where you will import the data from your XML backup. There are two options for accessing the XML file:
    - Browse to the location of your XML backup on your network, and choose **Upload and Restore**.
    - Alternatively, put the XML file in the Confluence home directory of the new site (`<CONFLUENCE -HOME-DIRECTORY>\restore`) then choose **Restore**.

Note: If you choose not to restore during the Confluence setup wizard, you can do the import later. Go to the Confluence administration console and choose to restore an XML backup. See Site Backup and Restore.

Step 7. Re-install your add-ons

Re-install any add-ons (plugins) that are not bundled with Confluence.

- Use the same version of the add-on as on your old Confluence site.
- The data created by the add-ons will already exist in your new Confluence site, because it is included in the XML backup.

Step 8. Check settings for new machine

If you are moving Confluence to a different machine, you need to check the following settings:

- Configure your new base URL. See Configuring the Server Base URL.
- Check your application links. See Linking to Another Application.
- Update any gadget subscriptions from external sites pointing to this Confluence site. For example, if your JIRA site subscribes to Confluence gadgets, you will need to update your JIRA site.
- Review any other resources that other systems are consuming from Confluence.

Method two – for installations with a large volume of attachments

Before you start

These instructions only apply to attachments stored in the file system. If you store attachments in the database see Attachment Storage Configuration to find out how to migrate between different attachment storage methods.

Step 1: Take note of your add-ons

Take note of the add-ons (plugins) currently installed and enabled in Confluence, so that you can reinstate them later. Make a note of the following for each add-on:

- Add-on name
- Version
- Enabled or disabled status. This is useful if you have enabled or disabled modules yourself, making your configuration differ from the default.

Step 2: Back up your data

1. Create an XML backup of your existing data, via the **Confluence administration console**. See Manu ally Backing Up the Site. Make a note of the location where you put the XML file. You will need it later to import your Confluence data into your new database.
2. Shut down Confluence.
3. Make a copy of the attachments directory (`<CONFLUENCE-HOME-DIRECTORY>\attachments`) in your Confluence Home directory. You will need it later to copy your Confluence attachments data into your new Confluence installation.
4. If you are using an external database, make a separate backup using the utilities that were installed with that database. This also is a precautionary measure.

Step 3: Set up the new database

Choose the database setup instructions for your new database, and follow those instructions to do the following:

- Install the database server.
- Perform any required configuration of the database server, as instructed.
- Add the Confluence database and user. Make a note of the username and password that you define in this step. You will need them later, when running the Confluence Setup Wizard.

Step 4. Install Confluence (same version number) in a new location

Now you will install Confluence again, with a different home directory path and installation path.

Note: You must use the same version of Confluence as the existing installation. (If you want to upgrade Confluence, you must do it as a separate step.) For example, if your current site is running Confluence 5.1.2, your new installation must also be Confluence 5.1.2.

When running the Confluence installer:

- Choose **Custom Install**. (Do **not** choose to upgrade your existing installation.)
- Choose a **new destination directory**. This is the installation directory for your new Confluence. It must not be the same as the existing Confluence installation.
- Choose a **new home directory**. This is the data directory for your new Confluence. It must not be the same as the existing Confluence installation.

Step 5. Download and install the database driver if necessary

Note that Confluence bundles some database drivers, but you'll need to install the driver yourself if it is not bundled. Follow the database setup instructions for your new database, to download and install the database driver if necessary.

Step 6. Run the Confluence setup wizard and copy your data to your new database

When running the Confluence setup wizard:

- Enter your license key, as usual.
- Choose **Production Installation** as the installation type.
- In the database configuration step, choose your new database type from the dropdown menu, then choose **External Database**.
- Choose the connection type: **Direct JDBC** or **Datasource**. If you are not sure which, choose 'Direct JDBC'. This is the most common connection type.
- When prompted for the database **user and password**, supply the credentials you defined earlier when adding the Confluence database to your database server.
- On the load content step, choose **Restore From Backup**. This is where you will import the data from your XML backup. There are two options for accessing the XML file:
    - Browse to the location of your XML backup on your network, and choose **Upload and Restore**.
    - Alternatively, put the XML file in the Confluence home directory of the new site (`<CONFLUENCE -HOME-DIRECTORY>\restore`) then choose **Restore**.

Note: If you choose not to restore during the Confluence setup wizard, you can do the import later. Go to the Confluence administration console and choose to restore an XML backup. See Site Backup and Restore.

Step 7: Copy your attachments across

Copy the contents of the attachments directory (`<CONFLUENCE-HOME-DIRECTORY>\attachments`) from your old Confluence Home directory to your new Confluence Home directory.

Step 8. Re-install your add-ons

Re-install any add-ons (plugins) that are not bundled with Confluence.

- Use the same version of the add-on as on your old Confluence site.
- The data created by the add-ons will already exist in your new Confluence site, because it is included in the XML backup.

Step 9. Check settings for new machine

If you are moving Confluence to a different machine, you need to check the following settings:
- Configure your new base URL. See Configuring the Server Base URL.
- Check your application links. See Linking to Another Application.
- Update any gadget subscriptions from external sites pointing to this Confluence site. For example, if your JIRA site subscribes to Confluence gadgets, you will need to update your JIRA site.
- Review any other resources that other systems are consuming from Confluence.

A note about case sensitivity in your database

'Collation' refers to a set of rules that determine how data is sorted and compared. Case sensitivity is one aspect of collation. Other aspects include sensitivity to kana (Japanese script) and to width (single versus double byte characters).

Case sensitive or case insensitive collation – how should you create your Confluence database? What about when you are migrating your existing Confluence instance from one database to another?

Setting up a New Confluence Instance

For new Confluence instances, we recommend using case sensitive collation for your Confluence database. This is the default collation type used by many database systems.

Note: Even if the database is configured for case sensitive collation, Confluence reduces all usernames to lower case characters before storing them in the database. For example, this means that 'joebloggs', 'joeBloggs' and 'JoeBloggs' will be treated as the same username.

Migrating an Existing Confluence Instance to a Different Database

The default Confluence configuration uses case sensitive database collation. This is typical of databases created under default conditions. If you are migrating from this type of configuration to a new database, we recommend that the new database uses case sensitive collation. If you use case insensitive collation, you may encounter data integrity problems after migration (for example, via an XML import) if data stored within your original Confluence site required case sensitive distinctions.

Troubleshooting

See our troubleshooting guide if you're unable to restore your XML backup.

Configuring Database Character Encoding

The database used with Confluence should be configured to use the same character encoding as Confluence. **The recommended encoding is Unicode UTF-8 (the equivalent for Oracle databases is AL32UTF8).**

There are two places where character encoding may need to be configured:

- when creating the database
- when connecting to the database (JDBC connection URL or properties).

The configuration details for each type of database are different. Some examples are below.

On this page:

- JDBC connection settings
- Creating a UTF-8 database
- Updating existing database to UTF-8

Related pages:

- Configuring Character Encoding
- Known issues for MySQL

JDBC connection settings

MySQL

Append "useUnicode=true to your JDBC URL:

```
jdbc:mysql://hostname:port/database?useUnicode=true&characterEncoding
=utf8
```

If you are modifying `confluence.cfg.xml` directly rather than via the Confluence Installation GUI, you'll need to escape out the & in the URL string as this is a reserved XML token and will break the syntax when the XML is parsed. An effective URL could be similar to:

```
<property
name="hibernate.connection.url">jdbc:mysql://hostname:port/data
base?useUnicode=true&amp;characterEncoding=utf8</property>
```

Creating a UTF-8 database

MySQL

1. Create a UTF-8 database with binary UTF-8 collation.
   ⓘ Binary UTF-8 provides case-sensitive collation.

   ```
   CREATE DATABASE confluence CHARACTER SET utf8 COLLATE utf8_bin;
   ```

2. You will also need to set the `Server Characterset` to `utf8`. This can be done by adding the following in `my.ini` for Windows or `my.cnf` for other OS. It has to be declared in the Server section, which is the section after `[mysqld]`:

   ```
   [mysqld]
   default-character-set=utf8
   ```

   If the above option does not work, try using *character_set_server=utf8* in lieu of *default-character-set= utf8*
3. Use the **status** command to verify database character encoding information.

   Screenshot: Using the Status Command to Verify Database Character Encoding

4. In some cases, the individual tables collation and character encoding may differ from the one that the database as a whole has been configured to use. Please use the command below to ensure all tables within your Confluence database are correctly configured to use UTF-8 character encoding and binary UTF-8 collation:

```
use confluence;
show table status;
```

Check for the value listed under the **Collation** column, to ensure it has been set to `utf8_bin` (that is, case-sensitive) collation for all tables.
If not, then this can be changed by the following command, executed for each table in the Confluence database:

```
ALTER TABLE tablename CONVERT TO CHARACTER SET utf8 COLLATE utf8_bin;
```

Please substitute the <tablename> above, with each table within the confluence database.

Relevant MySQL manual for more detailed explanation:

- Specifying Character Sets and Collations documentation.
- Connection Character Sets and Collations.
- SHOW TABLE STATUS Syntax.
- ALTER TABLE Syntax.

PostgreSQL

```
CREATE DATABASE confluence WITH ENCODING 'UNICODE';
```

Or from the command-line:

```
$ createdb -E UNICODE confluence
```

For more information see the PostgreSQL documentation.

For PostgreSQL running under Windows

Please note that international characters sets are only fully supported and functional when using PostgreSQL 8.1 and above under Microsoft Windows.

For PostgreSQL running under Linux

> Please make sure you check the following to ensure proper handling of international characters in your database

When PostgreSQL creates an initial database cluster, it sets certain important configuration options based on the host enviroment. The command responsible for creating the PostgreSQL environment `initdb` will check environment variables such as `LC_CTYPE` and `LC_COLLATE` (or the more general `LC_ALL`) for settings to use as database defaults related to international string handling. As such it is important to make sure that your PostgreSQL environment is configured correctly **before you install Confluence**.

To do this, connect to your PostgreSQL instance using `pgsql` and issue the following command:

```
SHOW LC_CTYPE;
```

If LC_CTYPE is set to either "C" or "POSIX" then certain string functions such as converting to and from upper and lower case will not work correctly with international characters. Correct settings for this value take the form `<LOCALE>.<ENCODING>` (`en_AU.UTF8` for example).

If your LC_CTYPE is incorrect please check the PostgreSQL documentation for information on configuring database localisation. It is not easy to change these settings with a database that already contains data.

Updating existing database to UTF-8

MySQL database with existing data

> **For an existing database**
> If you're using a existing database, confirm the Character Encoding by executing the query:
> *SHOW VARIABLES LIKE 'character%';* and *SHOW VARIABLES LIKE 'collation%';*.
> The results should be UTF-8.

> Before proceeding with the following changes, please **backup your database**.

This example shows how to change your database from latin1 to utf8, where your database is named "confluence".

1. Dump the database (except the plugindata table) to a text file using the mysqldump tool from the command-line :
```
mysqldump -p --default-character-set=latin1 -u <username>
--skip-set-charset --ignore-table='`confluence`.plugindata'
--ignore-table='`confluence`.attachmentdata' confluence >
confluence_database.sql
```
2. Dump the plugindata and attachmentdata tables to a text file using mysqldump separately. This is done separately as the recode step below can corrupt the binary data in these tables :
```
mysqldump -p --default-character-set=latin1 -u <username>
--skip-set-charset confluence attachmentdata plugindata >
confluence_blobtables.sql
```

3. copy `confluence_database.sql` to `confluence_utf8.sql`
4. Open `confluence_utf8.sql` in a text editor and change all character sets from 'latin1' to 'utf8'
5. Encode all the latin1 characters as UTF-8:
   `recode latin1..utf8 confluence_utf8.sql` (the recode utility is described at http://directory.fsf.org/recode.html; it can actually be downloaded from http://recode.progiciels-bpi.ca/, and is available for Ubuntu via `apt-get`)

In MySQL:

1. `DROP DATABASE confluence;`
2. `CREATE DATABASE confluence CHARACTER SET utf8 COLLATE utf8_bin;`

Reimport the UTF-8 text file, and also the plugindata and attachmentdata dumps:

1. `mysql -u <username> -p --default-character-set=utf8 --max_allowed_packet=64M confluence < /home/confluence/confluence_utf8.sql`
2. `mysql -u <username> -p --default-character-set=latin1 --max_allowed_packet=64M confluence < /home/confluence/confluence_blobtables.sql`

To support large imports, the parameter '--max_allowed_packet=64M' used above sets the maximum size of an SQL statement to be very large. In some circumstances, you may need to increase it further, especially if attachments are stored in the database.

Finally, since the plugindata and attachmentdata tables were not actually converted before, you need to instruct MySQL to convert all the text fields manually:

1. `ALTER TABLE plugindata CONVERT TO CHARACTER SET utf8;`
2. `ALTER TABLE attachmentdata CONVERT TO CHARACTER SET utf8;`

Testing database encoding

See Troubleshooting Character Encodings for a number of tests you can run to ensure your database encoding is correct.

Configuring database query timeout

If database queries are taking too long to perform, and your application is becoming unresponsive, you can configure a timeout for database queries. There is no default timeout in Confluence.To configure a database query timeout, do the following on your test server:

1. Shut down Confluence.

2. Extract `databaseSubsystemContext.xml` from the `confluence-x.x.x.jar` that is in `confluence/WEB-INF/lib/`, and put a copy in `confluence/WEB-INF/classes/`.

3. Edit `confluence/WEB-INF/classes/databaseSubsystemContext.xml` to add the defaultTimeout property to the "transactionManager" bean:

```
<bean id="tenantedTransactionManager"
class="org.springframework.orm.hibernate.HibernateTransactionManager"
plugin:available="true">
      <property name="sessionFactory" ref="sessionFactory"/>
      <property name="defaultTimeout" value="120"/>
   </bean>
```

The timeout is measured in seconds and will forcibly abort queries that take longer than this. In some cases, these errors are not handled gracefully by Confluence and will result in the user seeing the Confluence error page.

4. Start Confluence.

Once the timeout is working properly in your test environment, migration the configuration change to Confluence.

⚠️ You will need to reapply these changes when upgrading Confluence, as the original `databaseSubsystemC`

`ontext.xml` file changes from version to version.

Troubleshooting External Database Connections

A common administration issue when configuring Confluence is identifying database connectivity problems. This page tells you about a helper utility, in the form of a JSP page, that can help you to isolate database connectivity issues. It checks whether you can connect to a database with your application server. If your application server cannot connect to the database, Confluence will not be able to connect to the database either.

Introduction to the Atlassian Database Check Utility

You can use this utility to:

- Check that your application server can successfully query your database, either via immediate JDBC connectivity or a datasource in the context of your application server.
- Pinpoint problems in your configuration which may occur if the above is failing.

This is what the utility does:

- Check that a JDBC driver can be loaded into memory and view what is already loaded.
- Connect to a JDBC URL and do a 'select 1' from the database.
- Find a DataSource in the JNDI environment and do the above.
- View the System classpath (to ensure that the JDBC JAR file is there).

Using the Utility

If you have already set up Confluence completely

1. Download the attached testdatabase.jsp to your `<confluence-install>\confluence` directory.
2. Restart Confluence
3. Go to http://MY-CONF-SERVER:MY-CONF-PORT/testdatabase.jsp.
4. Check that your database driver is loaded into memory. If not, check the system classpath for the JDBC driver file, and that the driver is in the `<confluence-install>\lib` directory (for Confluence version 2.10 onwards) or `<confluence-install>\common\lib` (for earlier versions). Here are some instructions.
5. Enter the DB settings Confluence is using and test the database. If an error appears, check that the db service is running, the location matches, and that any users specified actually exist with the right login and permissions. You may be able to find a workaround by Googling the error.

If you cannot set up Confluence because of an error in 'Configuring Database'

1. Record the DB settings you are using for your direct JDBC or datasource connection in the 'Configure Database' step of your setup.
2. Download the attached testdatabase.jsp to your `<confluence-install>\confluence` directory.
3. Rename your `<confluence-install>\confluence\WEB-INF\web.xml` file to `backup web.xml`. This disables redirection.
4. Restart Confluence.
5. Go to http://MY-CONF-SERVER:MY-CONF-PORT/testdatabase.jsp.
6. Check that your database driver is loaded into memory. If not, check the system classpath for the JDBC driver file, and that the driver is in the `<confluence-install>\common\lib` directory as described in these instructions.
7. Enter the DB settings you recorded and test the database. If an error appears, check that the db service is running, the location matches, and that any users specified actually exist with the right login and permissions. You may be able to find a workaround by Googling the error.
8. After correcting the error, rename `<confluence-install>\confluence\WEB-INF\backup web.xml` back to `web.xml`.

Notes

If you use this utility, please let us know ways in which we could improve it or leave helpful hints for others here.
ℹ️ For a comprehensive set of database instructions that might be helpful for troubleshooting, please refer to the following links:

- PostgreSQL
- MySQL

Requesting Technical Support

If you are still stuck after attempting the suggestions above, lodge a free technical support request with information on your database setup.

Surviving Database Connection Closures

When a database server reboots or a network failure has occurred, all connections in the database connection pool are broken. To overcome this issue, Confluence would normally need to be restarted.

However, database connections in the database connection pool can be validated by running a simple SQL query. If a broken database connection is detected in the pool, a new one is created to replace it.

To do this, you can specify an optional validation query for your database connection. Depending on whether you are using a direct JDBC URL, or a data source, this is configured differently.

Determining the validation query SQL for your database type

Different database types have slightly different SQL syntax requirements for their validation query. The validation query should be as simple as possible, as this is run every time a connection is retrieved from the pool.

The following validation queries are recommended for the following types of databases:

Database Type	Validation Query
MySQL	`select 1`
Microsoft SQL Server	`select 1`
Oracle	`select 1 from dual`
PostgreSQL	`select 1`

Enabling validation query using direct JDBC

To ensure Confluence validates database connections in the database connection pool:

1. Shut down Confluence
2. Edit the `confluence.cfg.xml` file at the root of your Confluence Home and other important directories
3. Add the property "hibernate.c3p0.validate" and set it to "true", and add the property "hibernate.c3p0.preferredTestQuery" and set it to the value of the query you determined above for your database type. See this excerpt of the file with the two added properties for details:

   ```
   confluence.cfg.xml (excerpt)

   ...
   <property name="hibernate.c3p0.acquire_increment">1</property>
   <property name="hibernate.c3p0.idle_test_period">100</property>
   <property name="hibernate.c3p0.max_size">30</property>
   <property name="hibernate.c3p0.max_statements">0</property>
   <property name="hibernate.c3p0.min_size">0</property>
   <property name="hibernate.c3p0.timeout">30</property>
   <property name="hibernate.c3p0.validate">true</property>
   <property name="hibernate.c3p0.preferredTestQuery">select 1</property>
   ...
   ```

4. Save `confluence.cfg.xml`
5. Restart Confluence

Ensuring validation query using a data source

To ensure Confluence validates database connections in the database connection pool:

1. Shut down Confluence (or the Tomcat installation running Confluence).
2. Edit the `conf/server.xml` file in your Confluence Install Directory, or in the Tomcat installation's CATALINA_HOME directory.
3. Find the Resource element for your data source, and add the "validationQuery" field, with the value of the query you determined above for your database type. See this excerpt of the file with this added for details:

server.xml (excerpt)

```
...
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
          username="postgres"
          password="postgres"
          driverClassName="org.postgresql.Driver"
          url="jdbc:postgresql://localhost:5432/yourDatabaseName"
          maxTotal="25"
          maxIdle="10"
          validationQuery="select 1" />
...
```

4. Save `conf/server.xml`
5. Restart Confluence (or the Tomcat installation running Confluence).

Results and Considerations

You should now be able to recover from a complete loss of all connections in the database connection pool without the need to restart Confluence or the application server running Confluence.

⚠️ **Performance Considerations:**

- Setting this option has a performance impact. The overall decrease in performance should be minimal, as the query itself is quick to run. In addition, the query will only execute when you make a connection. Thus, if the connection is kept for the duration of a request, the query will only occur once per request.
- If you are running a large Confluence installation, you may wish to assess the performance impact of this change before implementing it.

Site Backup and Restore

> Atlassian recommends establishing a backup strategy using a native database tool for production installations of Confluence.

By default, Confluence backs up all data and attachments once a day to an XML backup file. These files are called XML site backups, and are stored in the `backups` directory of Confluence home. You can also create XML site backups manually. This mechanism is intended for small to medium-sized deployments of Confluence. It is not intended for use with large deployments with lots of pages and attachments (see below).

- Restore your site from an XML site backup
- Manually create an XML site backup
- Configuring Backups
- User Submitted Backup & Restore Scripts

XML site backups are fine for most small to medium-sized instances of Confluence, containing a few thousand pages and attachments. However, large instances of Confluence may find that backups become slow to create and use large amounts of disk space.

Note: Plugins are not included in the XML backup. After importing your backup into a new Confluence site, you will need to re-install all plugins (add-ons) that are not bundled with Confluence. (The plugindata table is not backed up in a manual backup.)

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

Backups for large installations

XML site backups are unsuitable for installations of Confluence that contain thousands of pages, as XML backups take progressively longer to complete as the amount of text increases. Another issue with XML site backups is that Confluence instances with gigabytes of attachments will consume disk space rapidly. This is because each site backup contains all content needed for a site restore. For example, if a 1 GB instance of Confluence is backed up daily, it will create 30 GB of backups per month if left unattended. When administering a large instance, you can reduce disk space by setting XML site backups to exclude attachments, then manually scheduling a backup of your attachments from the Confluence home directory or database. The backup manager can save space by saving changed files instead of all content.

Creation Delay	Disk Usage	Recommended Backup Method
Acceptable	Acceptable	XML site backup with attachments
Acceptable	Unacceptable	XML site backup minus attachments, plus manual backup of attachments
Unacceptable	Unacceptable	Manual backup of database and attachments

Creation Delay is the time it takes to create an XML site backup *minus attachments*.
Disk Usage can be estimated by multiplying the frequency of your XML site backups by their current size.

Manual backups

Confluence's Attachment Storage Configuration can be set to store attachments in the Confluence home directory, or in the database.

Database backup
Use your Database Administration Tool to create a backup of your Confluence database. If your database is storing your attachments, importing this later will restore all content. For instances with big attachments, please note that currently Confluence migrate attachments in a single transaction: CONF-9888.

Attachment backup
If stored on the filesystem, attachments are placed under the `attachments` directory of your Confluence home directory. Copy this directory to create a backup of all attachments.

To restore from these backups, please refer to Restoring Data from other Backups.

Production Backup Strategy

Confluence automatic daily XML backup is suitable if you:

- are evaluating Confluence
- do not have database administration familiarity, and your Confluence installation is small

Once your Confluence installation reaches more than a few thousand pages, the XML backup facility can be inefficient compared to your database's own backup tools. The built in backup functionality requires a lot of memory to run and is less reliable when restoring data.

Establishing a production system backup solution

Atlassian recommends establishing an alternative 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
  ⚠ To avoid any data inconsistency and corruption, it is recommended to shut down Confluence before creating a database backup or dump.
- Create a file system backup of your Confluence home directory
- Create a file system backup of your shared home directory (Clustered instances only)

Once this is in place, disable the daily backups through the scheduled jobs feature via **'Administration Console > Administration > Scheduled Jobs'**.

We want to stress that creating these two backups is *better* than having a Confluence XML backup. It is more robust and far more reliable for large production instances. You will be able to restore your whole site, including all data, attachments and configuration information intact with these two backups. See Restoring Data from other Backups.

Which files need to be backed up?

Backing up the whole home directory is the safest option, however most files and directories are populated on startup and can be ignored. At minimum, these files/directories *must* be backed up:

- **<conf-home>/confluence.cfg.xml**
- **<conf-home>/attachments**

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

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

The location of the home directory is configured on installation and is specified in the `confluence.init.properties` file. For installation created with the automatic installer the default locations are:

- Windows `C:\Program Files\Atlassian\Application Data\Confluence`
- Linux `/var/atlassian/application-data/confluence`

For Clustered instances only: Backing up the whole shared home directory is the safest option, however some files and directories are populated at runtime and can be ignored:

- <conf-home>/thumbnails
- <conf-home>/viewfile.

How do I back up?

The commands to back up your database will vary depending on your database vendor, for example the command for PostgreSQL is `pg_dump dbname > outfile`.

You should refer to the documentation for your particular database to find out more.

How do I restore?

Our guide on Migrating Confluence Between Servers has instructions on restoring a backup using this technique.

Other processes

XML backups are described and used for other processes in Confluence, like upgrading and moving servers. Using the backup strategy described above will work for those processes too.

- Our upgrade guide does not require the use of an XML backup (although the earlier Confluence upgrade procedure, and the JIRA upgrade guide, do use XML backups).
- Our migrate server procedure– used to set up a test server – can use a SQL dump as well.
- The database migration procedure uses the XML backup for small data sets. Large data sets will require third party database migration tools.

Note: The XML export built into Confluence is not suited for the backup or migration of large data sets. There are a number of third party tools that may be able to assist you with the data migration. If you would like help in selecting the right tool, or help with the migration itself, we can put you in touch with one of the Atlassian Experts.

Configuring Backups

Confluence backs up your data regularly into a zipped XML file. By default, this backup is performed at 2.00 a.m. each day and the backup files are stored in the `backups` folder under the Confluence Home directory. The default naming convention for the backup files is '`backup-yyyy_MM_dd`'. Confluence can write backups to both local and mapped network drives.

On this page:

- Configuring Confluence Backups
- Enabling Backup Path Configuration
- Notes

From the **Backup Administration** section of Confluence's administration console, you can:

- Include or exclude attachments in backups.
- Configure a different path to store backup files. (By default, this option is not available. See below for information about enabling the configuration option.)
- Change the naming format used for the files.

✅ You can also change the schedule of this backup using Confluence's scheduled jobs feature.

ℹ️ You need to have System Administrator permissions in order to configure these options.

Configuring Confluence Backups

To configure Confluence backups:

1. Choose the **cog icon**
   ⚙️ ▾
   , then choose **General Configuration** under Confluence Administration
2. Click '**Backup Administration**' in the '**Configuration**' section.
3. Click the '**Edit**' button on the '**Backup Administration**' screen.
4. Now you can do the following:
   - To use a different naming prefix format — Enter the new format in the '**Backup File Prefix**' input field.
   - To use a different date format — Enter the date format in the '**Backup File Date Pattern**' input field using the syntax described in this document from Sun.
   - To exclude attachments from backups — Deselect '**Backup Attachments**'. By default, this feature is 'On'.
   - To specify an alternate path to store backup files (if enabled) — Select '**Custom**' and then enter the path. The directory must be on either a local drive or a mounted network drive.
     ℹ️ Notes:
     - By default, this option is not available. See below for information about enabling the configuration option.
     - Please ensure the mapped drive is on a physical server, not a Virtual Machine image.
5. '**Save**' your changes.

✅ You can disable Confluence backups through the scheduled jobs feature.

Backup Administration

Perform a backup of your site daily to a chosen directory on your filesystem.

Backup Settings

Backup File Prefix	backup-
Backup File Date Pattern	yyyy_MM_dd
Backup Path	

Custom backup paths are not enabled. More about custom backup paths

☑ Backup Attachments

Submit   Cancel

Screenshot above: Editing the Backup Configuration

Enabling Backup Path Configuration

By default, it is not possible to specify a backup path via the Confluence Administration Console. This feature is disabled by default for security reasons. Administrators can restore this functionality by updating the relevant configuration property as described below. However, we recommend that you turn the feature **off** in production environments. For production environments, please review our Production Backup Strategy.

To enable the configuration option:

1. Edit the `confluence.cfg.xml` file found in the Confluence Home and other important directories.
2. Set the value of property `admin.ui.allow.daily.backup.custom.location` to `'true'` (without the quotation marks).

   confluence.cfg.xml

   ```
   <property
   name="admin.ui.allow.daily.backup.custom.location">true</property>
   ```

3. Restart Confluence.

If the value of the above configuration property is 'true', it will be possible to specify a backup path via the Confluence Administration Console. If the value of this property is 'false' or the property is not present in the configuration file, the backup path is not configurable.

Notes

Time is derived from the Confluence server

The time zone is taken from the server on which Confluence is running.

To check the time according to the server, do the following:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Click '**System Information**' in the left-hand panel and look at the '**System Time**'.

Backup strategy for large Confluence sites

Consider using the production backup strategy if your Confluence site is large or you are encountering

problems with your automated backup.

User Submitted Backup & Restore Scripts

These scripts are user-submitted and should be used with caution as they are not covered by Atlassian technical support. If you have questions on how to use or modify these scripts, please post them to Atlassian Answers.

Delete Old Backups - Wscript Script On Windows

This script examines backup filename and deletes them if necessary, it may need to be edited.

```
'If you want 3 day old files to be deleted then insert 3 next to Date - "your
number here"
'This script will search out and delete files with this string in them
".2005-12-04-" This of course depends on the number you enter.
'You can always do a wscript.echo strYesterday or strFileName to see what the
script thinks you are searching for.

dtmYesterday = Date - 3

strYear = Year(dtmYesterday)

strMonth = Month(dtmYesterday)
If Len(strMonth) = 1 Then
    strMonth = "0" & strMonth
End If

strDay = Day(dtmYesterday)
If Len(strDay) = 1 Then
    strDay = "0" & strDay
End If

strYesterday = strYear & "-" & strMonth & "-" & strDay

strFileName = "C:\test*." & strYesterday &"-*"

Set objFSO = CreateObject("Scripting.FileSystemObject")
objFSO.DeleteFile(strFileName)
```

Delete Old Backups - Basic Bash Script For Linux

Old XML backups can be deleted automatically by inserting a nightly or weekly automation script or cron similar to the following:

```
ls -t <path to your backup dir>/* | tail -n +6 | xargs -i rm {}
```

Or, using the older form of the `tail` command if your system does not support the standard form:

```
ls -t <path to your backup dir>/* | tail +6 | xargs -i rm {}
```

Delete Old Backups - Advanced Bash Script For Linux

Old XML backups can be deleted automatically by inserting a nightly or weekly automation script or cron similar to the following. Set the BACKUP_DIR and DAYS_TO_RETAIN variables to appropriate values for your site. Between runs, more files than DAYS_TO_RETAIN builds up.

```
#!/bin/sh

# Script to remove the older Confluence backup files.
# Currently we retain at least the last two weeks worth
# of backup files in order to restore if needed.

BACKUP_DIR="/data/web/confluence/backups"
DAYS_TO_RETAIN=14

find $BACKUP_DIR -maxdepth 1 -type f -ctime +$DAYS_TO_RETAIN -delete
```

Manual Database & Home Backup - Bash Script For Linux

This backs up a mySQL database and the Confluence home directory.

```
#!/bin/bash
CNFL=/var/confluence
CNFL_BACKUP=/backup/cnflBackup/`date +%Y%m%d-%H%M%S`

rm -rf $CNFL/temp/*
mkdir $CNFL_BACKUP
mysqldump -uroot -p<password> confluence|gzip >
$CNFL_BACKUP/confluence.mysql.data.gz
tar -cjvf $CNFL_BACKUP/data.bzip $CNFL > $CNFL_BACKUP/homedir.status
```

Backup by Date - Postgres

```
export d=`date +%u`
mkdir -p /home/backup/postgres/$d

sudo -u postgres pg_dumpall | bzip2 > /home/backup/postgres/$d/sql.bz2
```

Related Topics

- Site Backup and Restore
- Backup FAQ

Manually Backing Up the Site

Confluence is configured to back up its data automatically, as a zipped XML file. You can also manually perform this backup from the **Administration Console**.

You need to have System Administrator permissions in order to perform this function.

Note:

- Atlassian recommends that you follow the Production backup strategy if your Confluence site is large or you are encountering problems with your automated backup.
- Plugins are not included in the XML backup. After importing your backup into a new Confluence site, you will need to re-install all plugins (add-ons) that are not bundled with Confluence. (The plugindata table is not backed up in a manual backup.)

Related pages:

- Restoring a Site
- Configuring Backups
- Production Backup Strategy

Creating the site backup

To manually back up your site:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Backup & Restore** in the left-hand panel.
3. Choose **Archive to backups folder** to store a copy of the backup in the same folder as Confluence's backups.
   If you do not archive the backup it will be made available for you to download, and then deleted from the server after 24 hours.
4. Choose **Backup attachments** to include attachments in your backup.
5. Choose **Backup**.
   The process will take a few minutes.

Retrieving the Backup File

Confluence stores the backup as a zipped XML file in the `'backups'` directory under the Confluence Home directory on your Confluence server. To find your Confluence Home directory, see the documentation. You will need access to the Confluence server in order to retrieve this file.

Enabling the download of the backup file via the administration console

By default, it is not possible to retrieve the backup file via the Confluence Administration Console. This feature is disabled for security reasons.

Administrators can enable this functionality by updating the relevant configuration property as described below. If this functionality is enabled, Confluence will prompt you to download the backup file when the backup process finished. However, we recommend that you turn the feature **off** in production environments.

To enable download of the backup file from the Administration Console:

1. Edit the `confluence.cfg.xml` file found in the Confluence Home and other important directories.
2. Set the value of property `admin.ui.allow.manual.backup.download` to `'true'` (without the quotation marks).
3. Restart Confluence.

If the value of the above configuration property is 'true', it will be possible to download the backup file after manually backing up the site via the Confluence Administration Console. If the value of this property is 'false' or the property is not present in the configuration file, you will need to retrieve the backup file from the file system on the Confluence server. By default, the value is 'false'.

Notes

If you experience timeout errors, please consider bypassing Apache and creating the export directly from Tomcat. This will speed up the process and prevent timeouts. For example, your URL might be something like `http://<domain>.com`. To bypass this and access Tomcat directly, use this URL: `http://localhost:8080/confluence/admin/backup.action`.

Restoring a Site

CAUTION: Restoring a backup of an entire confluence site will:

- Wipe out all Confluence content in the database. Please ensure that your database is backed up before you start.
- Log you out after the restoration process. Please make sure you know the login details contained in the data that you are about to restore.

Related pages:

- Production Backup Strategy
- Manually Backing Up the Site

This page describes how to restore data from an XML backup file into an existing Confluence installation. If you want to restore data into a new site, follow the instructions on restoring from backup during setup.

You need System Administrator permissions in order to perform this function.

Notes before you start:

- **All content replaced.** Restoring a site from backup will replace all your content, as described in the warning above.
- **Selective space restoration not possible.** You cannot select a single space to restore from the entire site backup when the backup contains more than one space.
- **Version compatibility.** Confluence supports importing site and space backups into the same major version only. In Confluence, the major version is up to the second number, eg 5.4.3's major version is 5.4. Importing an old XML backup file to a new major version (for example, Confluence 5.5 to Confluence 5.6) is not supported.

	Import into 5.4.x	**Import into 5.5.x**	**Import into 5.6.x**
Site backup from 5.5.1			
Space backup from 5.4.2			

- **XML backups should not be used to upgrade Confluence.** Upgrade Confluence by following: Upgrading Confluence.

Restoring data from an XML backup

You can restore data from an XML backup file located somewhere on your local computer or a shared drive, or you can copy the XML file into the Confluence installation and restore it from there. The second option is recommended for large backup files. Both options are described below.

To restore data from an XML backup located outside Confluence:

1. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
2. Choose **Backup and Restore** in the left-hand panel.
3. Choose **Choose File** and browse for the backup file.
4. Uncheck **Build Index** if you want to create the index at a later stage.
5. Choose **Upload and Restore**.

To restore data from an XML backup located in your Confluence installation:

1. Copy your XML backup zip file into the `restore` directory in your Confluence home directory. For example:
   - On UNIX: `/opt/java/src/confluence/deployments/conf.atlassian.com/home/restore`
   - On Windows: `C:\Program Files\Atlassian\Application Data\Confluence x.x\restore`
2. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
3. Choose **Backup and Restore** in the left-hand panel.
4. The zip file that you copied in step 1 will appear in the list of files under the heading **Restore a backup from the Confluence Home Directory** on your Confluence Administration Console. Select the zip file.
5. Uncheck **Build Index** if you want to create the index at a later stage.
6. Choose **Restore**.

Notes

- **Production backup strategy preferred.** Atlassian recommends that you follow the Production Backup Strategy for your production Confluence site, because Confluence XML backups are not recommended for non-evaluation sites.

- **Restoring from other backups.** If your daily backup zip files cannot be restored for some reason, but you have backups of both your database and your Confluence home directory, then it is still possible to restore from these backups.

Restoring a Space

You can export a space, including pages, comments and attachments, to a zip that contains an XML file and, optionally, all the attachments in the space. To import the space to another Confluence site, restore the zip as described below.

You need to have System Administrator permissions in order to restore a space from an XML zip file.

Export and import compatibility

Confluence only accepts space exports from compatible versions. To find out which versions your current Confluence version can accept space exports from, go to

⚙️ ▾

> **General Configuration** > **Backup and Restore**.

You may not be able to import spaces into earlier versions of Confluence; the **Backup and Restore** page in the later Confluence version will give you some indication of whether it is likely to be compatible.

If you need to import a space from an incompatible version, see the workaround described below.

On this page:

- Export and import compatibility
- Restoring a space from an XML backup
- Workaround for restoring spaces from Confluence 5.3 and below

Related pages:

- Restoring a Site
- Confluence Administrator's Guide

Also, Confluence won't allow you to import a space if there's already a space with the same space key. If you already have a space with the same key, you'll need to delete the existing space before importing the new one.

Restoring a space from an XML backup

You can restore data from an XML backup file located somewhere on your local computer or a shared drive, or you can copy the XML file into the Confluence restore directory and restore it from there. We recommend the second option for large backup files.

To restore data from an XML backup located outside Confluence:

1. Choose the **cog icon**
   ⚙️ ▾
   , then choose **General Configuration** under Confluence Administration
2. Choose **Backup and Restore** in the left-hand panel
3. Choose **Choose File** and browse for the backup file
4. Uncheck **Build Index** if you want to create the index at a later stage
5. Choose **Upload and Restore**

To restore data from an XML backup located in your Confluence installation:

1. Copy your XML backup zip file into the `restore` directory in your Confluence home directory. For example:
   - On UNIX: `/opt/java/src/confluence/deployments/conf.atlassian.com/home/restore`
   - On Windows: `C:\Program Files\Atlassian\Application Data\Confluence x.x\restore`
2. Choose the **cog icon**
   ⚙️ ▾

, then choose **General Configuration** under Confluence Administration
3. Choose **Backup and Restore** in the left-hand panel
4. Select the zip file that you copied in step 1 from the list of files under the heading **Restore a backup from the Confluence Home Directory**
5. Uncheck **Build Index** if you want to create the index at a later stage
6. Choose **Restore**

Workaround for restoring spaces from Confluence 5.3 and below

If you need to import a space from a version prior to Confluence 5.3, you can use a temporary Confluence installation to upgrade the space export to the right version number:

1. Download the same version of Confluence as the version you exported the space from (you can get older versions of Confluence at the Confluence Downloads Archive)
2. Install that version of Confluence on a temporary server
3. Import the space into this temporary Confluence site
4. Upgrade Confluence on your temporary site to the same version as the site where you want to import the space (see Upgrading Confluence for instructions)
5. Export the space from your temporary Confluence site (it'll now have the right version number)
6. Import the space into your production Confluence site

Restoring a Test Instance from Production

See Migrating Confluence Between Servers for a more comprehensive explanation.

Many Confluence administrators will have a production instance running the "live" version of Confluence, as well as a test instance for testing upgrades and so on. In this situation, it's quite common that the two instances are running different versions of Confluence. This document describes how to copy the data from a production instance to a test instance, where the production version may be different to the test version.

Before proceeding with this guide, ensure you have read and understood the normal procedure for upgrading Confluence.

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

Upgrading a test Confluence instance with production data

Essentially, we are copying both the production home directory and database to the test instance. We then update the database details on the test instance to point to the test database, leaving all other instance metadata (most importantly the Confluence *build number*) the same as production.

1. Shut down your test instance.
2. Restore the production database to the test database server.
3. Create a backup of the `confluence.cfg.xml` file found in the home directory of the test instance.
4. Copy the production confluence-home directory to the test application server.
5. Open the `confluence.cfg.xml` which has been copied in a text editor. Change the database settings to match the test database server. **Ensure you do not point to your production database.** (You can compare with the backup you made in Step 3 if you need to get the database settings. Don't just copy this file – you need the build number unchanged from production to indicate the database is from an older version of Confluence.)

Before starting your test instance, you need to do the following steps to ensure no contact with production systems.

Ensuring no contact with production systems

To ensure no contact with external systems, you will need to disable both inbound and outbound mail services.

1. Disable global outbound mail by running the following database query:

```
SELECT * FROM BANDANA WHERE BANDANAKEY =
'atlassian.confluence.smtp.mail.accounts';
```

2. Disable space-level mail archiving by running the following database query:

```
SELECT * FROM BANDANA WHERE BANDANAKEY =
'atlassian.confluence.space.mailaccounts';
```

Change the 'SELECT *' to a 'DELETE' in the above queries once you are sure you want to remove the specified accounts.

Once this is done, you can start your test instance without any mails being sent or retrieved. Think carefully about other plugins which may access production systems (SQL macro, etc.). These should be disabled promptly after starting the test instance.

You can create a developer license for this server and update the License Details after starting up.

See also

Upgrading Confluence
Migrating Confluence Between Servers
Restoring to a Test Instance of Confluence from Production

Restoring Data from other Backups

Typically, Confluence data is restored from the Administration Console or from the Confluence Setup Wizard.

If you are experiencing problems restoring from 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:

```
<object class="Attachment" package="com.atlassian.confluence.pages">
<id name="id">98</id>
<property name="fileName"><![CDATA[myimportantfile.doc]]></property>
...
<property name="content" class="Page" package="com.atlassian.confluence.pages"><id
name="id">10001</id>
</property>
...
</object>
```

Entities.xml Page Object

This XML describes a page. In this example, the page id is 98 and the title is Editing Your Files. The rest of the XML can be ignored:

```
<object class="Page" package="com.atlassian.confluence.pages">
<id name="id">98</id>
<property name="title"><![CDATA[Editing Your Files]]></property>
...
</object>
```

Instructions for Recovering Attachments

Each file must be individually renamed and re-uploaded back into Confluence by following the instructions

below. Choose one of the three methods:

Choice A - Recover Attachments By Filename

Best if you know each filename you need to restore, especially if you want just a few files:

1. Unzip the backup directory and open entities.xml.
2. Search entities.xml for the filename and find the attachment object with that filename. Locate its page and attachment id.
3. Using the page and attachment id from entities.xml, go to the attachments directory and open that directory with that page id. Locate the file with the attachment id.
4. Rename the file to the original filename and test it.
5. Repeat for each file.
6. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

Choice B - Restore Files By Page

Best if you only want to restore attachments for certain pages:

1. Unzip the backup directory and open entities.xml.
2. Search entities.xml for the page title and find the page object with that title. Locate its page id.
3. Go to the attachments directory and open that directory with that page id. Each of the files in the directory is an attachment that must be renamed.
4. Search entities.xml for attachment objects with that page id. Every attachment object for the page will have an attachment id and filename.
5. Rename the file with that attachment id to the original filename and test it.
6. Repeat for each page.
7. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

Choice C - Restore All Files

Best if you have a small backup but want to restore many or all the attachments inside:

> Following process is applicable to **space** export only. Site xml backups do not require page id to be updated manually due to the nature of persistent page_id's.

1. Unzip the backup directory and open entities.xml.
2. Go to the attachments directory and open any directory. The directory name is a page id. Each of the files in the directory is an attachment that must be renamed.
3. Search entities.xml for attachment objects with that page id. When one is found, locate the attachment id and filename.
4. Rename the file with that attachment id to the original filename and test it.
5. Find the next attachment id and rename it. Repeat for each file in the directory.
6. Once all files in the current directory are renamed to their original filenames, search entities.xml for the page id, eg directory name. Find the page object with that page id and locate its page title.
7. Rename the directory to the page title and move on to the next directory. Repeat for each un-renamed directory in the attachments directory.
8. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

Troubleshooting failed XML site backups

> XML site backups are only necessary for migrating to a new database. Setting up a test server or Establishing a reliable backup strategy is better done with an SQL dump.

Related pages:

- Enabling detailed SQL logging

Seeing an error when creating or importing a backup?

Problem	Solution
Exception while creating backup	Follow instructions below
Exception while importing backup	Follow Troubleshooting XML backups that fail on restore instead

Resolve Errors With Creating An XML Backup

The errors may be caused by a slightly corrupt database. If you're seeing errors such as 'Couldn't backup database data' in your logs, this guide will help you correct the error on your own. We strongly recommend that you backup your database and your Confluence home directory beforehand, so that you can restore your site from those if required. If you are unfamiliar with SQL, we suggest you contact your database administrator for assistance.

Preferable solution

The Production Backup Strategy is a very reliable and more efficient way to do backups. If you are running into problems with XML backups - whether memory related or because of problems like the one described here - use the native backup tool as an alternate solution.

To Identify And Correct The Problem

To work out where the data corruption or problems are, increase the status information reported during backup, then edit the invalid database entry:

1. Stop Confluence.
2. If you have an external database, use a database administration tool to create a manual database backup.
3. Backup your Confluence home directory. You will be able to restore your whole site using this and the database backup.
4. Open the `my_confluence_install/confluence/WEB-INF/classes/log4j.properties`and add this to the bottom and save:

   ```
   log4j.logger.com.atlassian.confluence.importexport.impl.XMLDatabinder=DEBUG
   , confluencelog
   log4j.additivity.com.atlassian.confluence.importexport.impl.XMLDatabinder=f
   alse
   ```

5. Find your atlassian-confluence.log. Move or delete all existing Confluence logs to make it easier to find the relevant logging output.
6. Restart Confluence and login.
7. Begin a backup so that the error reoccurs.
8. You must now check your log files to find out what object could not be converted into XML format. Open `confluence-home/logs/atlassian-confluence.log`. Scroll to the bottom of the file.
9. Do a search for 'ObjectNotFoundException'. You should see an error similar to this:

```
01 2005-08-24 00:00:33,743 DEBUG
[DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
object: com.atlassian.confluence.core.ContentPermission with ID:
5 to XML.
02 2005-08-24 00:00:33,743 DEBUG
[DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: type
03 2005-08-24 00:00:33,743 DEBUG
[DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: group
04 2005-08-24 00:00:33,743 DEBUG
[DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: expiry
05 2005-08-24 00:00:33,743 DEBUG
[DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: content
06 [DOCPRIV2:ERROR] LazyInitializer - Exception initializing
proxy <net.sf.hibernate.ObjectNotFoundException: No row with the
given identifier exists: 2535,
07 of class:
com.atlassian.confluence.core.ContentEntityObject>net.sf.hiberna
te.ObjectNotFoundException:
08 No row with the given identifier exists: 2535, of class:
com.atlassian.confluence.core.ContentEntityObject
09 at
net.sf.hibernate.ObjectNotFoundException.throwIfNull(ObjectNotFo
undException.java:24)
10 at
net.sf.hibernate.impl.SessionImpl.immediateLoad(SessionImpl.java
:1946)
11 at
net.sf.hibernate.proxy.LazyInitializer.initialize(LazyInitialize
r.java:53)
12 at
net.sf.hibernate.proxy.LazyInitializer.initializeWrapExceptions(
LazyInitializer.java:60)
13 at
net.sf.hibernate.proxy.LazyInitializer.getImplementation(LazyIni
tializer.java:164)
14 at
net.sf.hibernate.proxy.CGLIBLazyInitializer.intercept(CGLIBLazyI
nitializer.java:108)
15 at
com.atlassian.confluence.core.ContentEntityObject$$EnhancerByCGL
IB$$cc2f5557.hashCode(<generated>)
16 at java.util.HashMap.hash(HashMap.java:261)
17 at java.util.HashMap.containsKey(HashMap.java:339)
18 at
com.atlassian.confluence.importexport.impl.XMLDatabinder.toGener
icXML(XMLDatabinder.java:155)
```

10. Open a DBA tool such as DbVisualizer and connect to your database instance. Scan the table names in the schema. You will have to modify a row in one of these tables.
11. To work out which table, open `catalina.out`, check the first line of the exception. This says there

was an error writing the `ContentPermission` object with id 5 into XML. This translates as *the row with primary key 5 in the CONTENTLOCK table* needs fixing. To work out what table an object maps to in the database, here's a rough guide:

- Pages, blogposts, comments --> CONTENT table
- attachments --> ATTACHMENTS table
- More information can be found in the schema documentation

12. Now you must find the primary key of the incorrect row in this table. In this case, you can check the first line and see that the row has a primary key of 5.
13. Each property is written to a column, so the last property that was being written has the incorrect value. The row being written to when the exception was thrown was `CONTENT` (line 5) with a value of `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 ot the Confluence database. Locate the row in the relevant table and correct the entry. Check other rows in the table for the default column value, which may be null, 0 or blank. Overwrite the invalid row value with the default.
15. Restart Confluence.
16. Attempt the backup again. If the backup fails and you are stuck, please lodge a support request with your latest logs.

Troubleshooting "Duplicate Key" related problems

If you are encountering an error message such as:

```
could not insert:
[bucket.user.propertyset.BucketPropertySetItem#bucket.user.propertyse
t.BucketPropertySetItem@a70067d3]; SQL []; Violation of PRIMARY KEY
constraint 'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate key
in object 'OS_PROPERTYENTRY'.; nested exception is
java.sql.SQLException: Violation of PRIMARY KEY constraint
'PKOS_PROPERTYENTRY_314D4EA8'. Cannot insert duplicate key in object
'OS_PROPERTYENTRY'.
```

this indicates that the Primary Key constraint 'PK_OS_PROPERTYENTRY_314D4EA8' has duplicate entries in table 'OS_PROPERTYENTRY'.
You can locate the constraint key referring to 'PK_OS_PROPERTYENTRY_314D4EA8' in your table 'OS_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS_PROPERTYENTRY' table is:

```
SELECT ENTITY_NAME,ENTITY_ID,ENTITY_KEY,COUNT(*) FROM
OS_PROPERTYENTRY GROUP BY ENTITY_NAME,ENTITY_ID,ENTITY_KEY HAVING
COUNT(*)>1
```

To Help Prevent This Issue From Reoccuring

1. If you are using the embedded database, be aware that it is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss, which is why an external database is recommended for production use. You should migrate to an external database.
2. If you are using an older version of Confluence than the latest, you should consider upgrading at this point.

Troubleshooting XML backups that fail on restore

XML site backups are only necessary for migrating to a new database. Upgrading Confluence, Setting up a test server or Production Backup Strategy is better done with an SQL dump.

Seeing an error when creating or importing a site or space backup?

Problem	Solution
Exception while creating backup	See Troubleshooting failed XML site backups
Exception while importing backup	See instructions below

On this page:

- Resolve Errors When Attempting To Restore An XML Backup
  - Troubleshooting "Duplicate Entry" for key "cp_" or "cps_"
  - Troubleshooting "Duplicate Key" related problems
  - Troubleshooting "net.sf.hibernate.PropertyValueException: not-null" related problems
  - To Help Prevent this Issue from Recurring

Related Topics:

- Troubleshooting failed XML site backups

Resolve Errors When Attempting To Restore An XML Backup

The errors may be caused by a slightly corrupt database. You will need to find the XML backup file entry that is violating the DB rules, modify the entry and recreate the XML backup:

1. On the instance being restored, follow the instructions to disable batched updates (for simpler debugging), log SQL queries and log SQL queries **with parameters** at Enabling Detailed SQL Logging.
2. Once all three changes have been made, restart Confluence.
3. Attempt another restore.
4. Once the restore fails, check your log files to find out what object could not be converted into XML format. For Confluence distribution users, check your Confluence install directory under the `/logs/` and check both `atlassian-confluence.log` and `catalina.out` file. The correct file will contain SQL debug output.
5. Scroll to the bottom of the file and identify the last error relating to a violation of the database constraint. For example:

```
2006-07-13 09:32:33,372 ERROR
[confluence.importexport.impl.ReverseDatabinder] endElement
net.sf.hibernate.exception.ConstraintViolationException:
  could not insert:
[com.atlassian.confluence.pages.Attachment#38]
net.sf.hibernate.exception.ConstraintViolationException: could
not insert: [com.atlassian.confluence.pages.Attachment#38]
...
Caused by: java.sql.SQLException: ORA-01400: cannot insert NULL
into ("CONFUSER"."ATTACHMENTS"."TITLE")
at
oracle.jdbc.driver.DatabaseError.throwSqlException(DatabaseError
.java:112)
at oracle.jdbc.driver.T4CTTIoer.processError(T4CTTIoer.java:331)
at oracle.jdbc.driver.T4CTTIoer.processError(T4CTTIoer.java:288)
```

This example indicates a row in your attachment table with ID = 38 that has a null title.

6. Go to the server that the backup was created on. You must have a copy of the database from which the backup was created. If you do not have this, use a DBA tool to restore a manual backup of the database.
7. Open a DBA tool and connect to the original database instance and scan the table names in the schema. You will have to modify a row in one of these tables.

8. To work out which table, open `catalina.out`, check the first line of the exception. To work out what table an object maps to in the database, here's a rough guide:
   - Pages, blogposts, comments --> CONTENT table.
   - attachments --> ATTACHMENTS table.
9. To correct the example error, go to the attachment table and find that attachment object with id 38. This will have a a null title. Give a title using the other attachments titles as a guide. You may have a different error and should modify the database accordingly.
10. Once the entry has been corrected, create the XML backup again.
11. Import the backup into the new version.
12. If the import succeeds, revert the changes made in your SQL logging to re-enable disable batched updates and turn off log SQL queries and log SQL queries with parameters.
13. Restart Confluence.

Troubleshooting "Duplicate Entry" for key "cp_" or "cps_"

If you are encountering an error message such as:

```
com.atlassian.confluence.importexport.ImportExportException: Unable
to complete import because the data does not match the constraints in
the Confluence schema. Cause:
MySQLIntegrityConstraintViolationException: Duplicate entry
'1475804-Edit' for key 'cps_unique_type'
```

This indicates that the XML export came from a version of Confluence with a corrupt permissions database, caused by some 3rd party plugin. This is an issue that was fixed when CONF-22123 was implemented in Confluence 3.5.2. The simplest workaround is to export the space again after upgrading the instance to 3.5.2 or above. If that is not an option, then either the export will need to be edited manually to remove the duplicate permission entries or the source instance will need to have the offending entries removed. The following SQL queries can be used to look for such entries:

```
SELECT * FROM CONTENT_PERM WHERE USERNAME IS NULL AND GROUPNAME IS NULL;

SELECT cp.ID, cp.CP_TYPE, cp.USERNAME, cp.GROUPNAME, cp.CPS_ID, cp.CREATOR,
cp.CREATIONDATE, cp.LASTMODIFIER, cp.LASTMODDATE
FROM CONTENT_PERM cp
WHERE cp.USERNAME IS NOT NULL AND cp.GROUPNAME IS NOT NULL;

SELECT cps1.ID, cps1.CONTENT_ID, cps1.CONT_PERM_TYPE FROM CONTENT_PERM_SET cps1,
CONTENT_PERM_SET cps2
WHERE cps1.ID <> cps2.ID AND
cps1.CONTENT_ID = cps2.CONTENT_ID AND
cps1.CONT_PERM_TYPE = cps2.CONT_PERM_TYPE
ORDER BY cps1.CONTENT_ID, cps1.CONT_PERM_TYPE, cps1.CREATIONDATE ASC;

SELECT cp.ID, cp.CP_TYPE, cps.CONTENT_ID,
(SELECT scps.ID FROM CONTENT_PERM_SET scps WHERE scps.CONTENT_ID = cps.CONTENT_ID
AND scps.CONT_PERM_TYPE = cp.CP_TYPE) AS suggested_cps_id
FROM CONTENT_PERM cp, CONTENT_PERM_SET cps
WHERE cp.CPS_ID = cps.ID AND
cp.CP_TYPE <> cps.CONT_PERM_TYPE;

SELECT DISTINCT cp1.ID, cp1.CP_TYPE, cp1.USERNAME, cp1.GROUPNAME, cp1.CPS_ID,
cp1.CREATOR, cp1.CREATIONDATE, cp1.LASTMODIFIER, cp1.LASTMODDATE
FROM CONTENT_PERM cp1, CONTENT_PERM_SET cps1, CONTENT_PERM cp2, CONTENT_PERM_SET
cps2
WHERE
cp1.CPS_ID = cps1.ID AND
cp2.CPS_ID = cps2.ID AND
cp1.ID <> cp2.ID AND
cps1.CONTENT_ID = cps2.CONTENT_ID AND
cp1.CP_TYPE = cp2.CP_TYPE AND
cp1.USERNAME = cp2.USERNAME
ORDER BY cp1.CPS_ID, cp1.CP_TYPE, cp1.USERNAME, cp1.CREATIONDATE;

SELECT DISTINCT cp1.ID, cp1.CP_TYPE, cp1.USERNAME, cp1.GROUPNAME, cp1.CPS_ID,
cp1.CREATOR, cp1.CREATIONDATE, cp1.LASTMODIFIER, cp1.LASTMODDATE
FROM CONTENT_PERM cp1, CONTENT_PERM_SET cps1, CONTENT_PERM cp2, CONTENT_PERM_SET
cps2
WHERE
cp1.CPS_ID = cps1.ID AND
cp2.CPS_ID = cps2.ID AND
cp1.ID <> cp2.ID AND
cps1.CONTENT_ID = cps2.CONTENT_ID AND
cp1.CP_TYPE = cp2.CP_TYPE AND
cp1.GROUPNAME = cp2.GROUPNAME
ORDER BY cp1.CPS_ID, cp1.CP_TYPE, cp1.GROUPNAME, cp1.CREATIONDATE;

SELECT * FROM CONTENT_PERM_SET
WHERE ID NOT IN (SELECT DISTINCT CPS_ID FROM CONTENT_PERM);
```

Remove all matching entries and perform the export again.

Troubleshooting "Duplicate Key" related problems

If you are encountering an error message such as:

```
could not insert:
[bucket.user.propertyset.BucketPropertySetItem#bucket.user.propertyse
t.BucketPropertySetItem@a70067d3]; SQL []; Violation of PRIMARY KEY
constraint 'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate key
in object 'OS_PROPERTYENTRY'.; nested exception is
java.sql.SQLException: Violation of PRIMARY KEY constraint
'PKOS_PROPERTYENTRY_314D4EA8'. Cannot insert duplicate key in object
'OS_PROPERTYENTRY'.
```

This indicates that the Primary Key constraint 'PK_OS_PROPERTYENTRY_314D4EA8' has duplicate entries in table 'OS_PROPERTYENTRY'.

You can locate the constraint key referring to 'PK_OS_PROPERTYENTRY_314D4EA8' in your table 'OS_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS_PROPERTYENTRY' table is:

```
SELECT ENTITY_NAME,ENTITY_ID,ENTITY_KEY,COUNT(*) FROM
OS_PROPERTYENTRY GROUP BY ENTITY_NAME,ENTITY_ID,ENTITY_KEY HAVING
COUNT(*)>1
```

Troubleshooting "net.sf.hibernate.PropertyValueException: not-null" related problems

If you're receiving a message like:

```
ERROR [Importing data task]
[confluence.importexport.impl.ReverseDatabinder] endElement
net.sf.hibernate.PropertyValueException: not-null property references
a null or transient value:
com.atlassian.user.impl.hibernate.DefaultHibernateUser.name
```

This means there's an unexpected null value in a table. In the above example, the error is in the name column in the USERS table. We've also seen them in the ATTACHMENTS table.

Remove the row with the null value, redo the xml export, and reimport.

To Help Prevent this Issue from Recurring

1. If you are using the embedded database, be aware that it is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss, which is why an external database is recommended for production use. You should migrate to an external database.
2. If you are using an older version of Confluence than the latest, you should consider upgrading at this point.

The problem with different settings for case sensitivity varies between databases. The case sensitivity of the database is usually set through the collation that it uses. Please vote on the existing issue

Attachment Storage Configuration

Database attachment storage was deprecated in Confluence 5.5. If you currently store attachments in the database you will be able to continue to do so, but you will be unable to switch to storing attachments in the database in Confluence 5.5 or later.

System Administrators can configure where

Confluence stores attachments. Attachments can be stored in a:

- File system - locally in the Confluence home directory, or
- Database - in Confluence's configured database (deprecated)

To configure Confluence attachment storage:

- Choose the **cog icon**

  ⚙ ▾

  , then choose **General Configuration** under Confluence Administration
- Choose **Attachment Storage**.

On this page:

- Attachment Storage Options
  - Local File System
  - Database (deprecated)
- Migrating between attachment storage systems

Related pages:

- Working with Confluence Logs

Attachment Storage Options

Local File System

By default, Confluence stores attachments in the `attachments` directory within the configured Confluence home folder.

Database (deprecated)

Confluence 5.4 and earlier gives administrators the option to store attachments in the database that Confluence is configured to use.

> While storing attachments in the database can offer some advantages (such as ease of backup, and avoiding issues with some characters in attachment filenames), please be aware that the amount of space used by the database will increase because of greater storage requirements.

Migrating between attachment storage systems

You can migrate your attachments from one storage system to another. All existing attachments will be moved over to the new attachment storage system.

> When the migration occurs, all other users will be locked out of the Confluence instance. This is to prevent modification of attachments while the migration occurs. Access will be restored as soon as the migration is complete.

When migrating attachments from your database to a filesystem, the attachments are removed from the database after migration. However, when migrating attachments from a filesystem to your database, the attachments remain on the filesystem after migration.

To improve logging during the migration, add the package `com.atlassian.confluence.pages.persistence.dao` with level DEBUG. See Configuring Logging for more information.

To migrate, follow the steps below:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Click '**Attachment Storage**' in the left-hand panel. The current configuration will be displayed.
   Screenshot: Attachment storage configuration

3. Click the '**Edit**' button to modify the configuration.
4. Select the storage system you desire.
   Screenshot: Edit attachment storage

5. Click the '**Save**' button to save the changes.
6. A screen will appear, asking you to confirm your changes. Clicking 'Migrate' will take you to a screen that displays the progress of the migration.
   Screenshot: migration warning

The following external website provides further information on migrating attachments from database to file system storage that you might find helpful - http://www.scandio.de/2013/05/confluence-attachment-migration-the-safe-way-2/.

Configuring Attachment Size

Confluence gives you the option of limiting the maximum size of a single file attachment.

To configure the maximum size allowed for an attachment:

1. Go to

   > **General Configuration**.

Related pages:

- Recognised System Properties

- Files

   2. Choose **Edit**.
   3. Enter the maximum size next to **Attachment Maximum Size**'.
      The default is 100mb.
   4. Choose **Save**.

To configure the maximum 'index-able size of attachments':

By default, large attachment is defined as greater than 1 MB. The threshold for attachments that won't get excerpts can be modified using the system property `atlassian.indexing.contentbody.maxsize`, which takes a size in bytes.

For example to specify 250 kb you would use the following JVM parameter:

```
-Datlassian.indexing.contentbody.maxsize=256000
```

Limiting the size of attachment indexing has the following effects:

- Decreases the size of the index when large attachments are present.

- Prevent excerpts of large attachments being displayed in search results.

Hierarchical File System Attachment Storage

> The way attachments are stored changed significantly in Confluence 3.0. If you are upgrading from Confluence 2.10 or earlier see Upgrading Confluence for recommended upgrade paths, and read the version of the Hierarchical File System Attachment Storage page in our Confluence 3.0 documentation which provides more detail about migrating to the new file system structure.

Confluence stores attachments, such as files and images, in a file system. Confluence's attachment storage layout is designed to:

1. Limit the number of entries at any single level in a directory structure (as some file systems have a limit on the number of files that can be stored in a directory).
2. Partition attachments per space making it possible for a system admin to selectively back up attachments from particular spaces.

Attachments in Confluence have a number of identifying attributes: *content id of the file itself*, the *space id* and *content id of the page the file is attached to*. This means the file logically belongs to a piece of content which logically belongs in a space (not all content belongs to a space). For files within a space in Confluence, the directory structure is typically 8 levels, with the name of each directory level based on the following algorithm:

level	Derived From
1 (top)	Always 'ver003' indicating the Confluence version 3 storage format
2	The least significant 3 digits of the *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 *content id* 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

Within the 8th level will be a file for each version of that attachment, named to match the version number e.g. 1

An example:

Attached Files

A	B	C	
file content id: 745644 space id: 800432 page content id: 632780	file content id: 782234 space id: 800432 page content id: 620002	file content id: 771250 space id: 810032 page content id: 603101	file c *globa* page

Directory structure

Attachments directory

ver003 — top level directory partitioni

nonspaced | **182** | **32** — 3 least significant digits of s

50 | **60** — next 3 least significant digit

800432 | **810032** — space id

242 | **30** | **2** | **101** — 3 least significant digits of th content id % 250

11 | **132** | **120** | **103** — next 3 least significant digits page's content id % 250

511242 | **632780** | **620002** | **603101** — content id of the page the fi

701002 | **745644** | **782234** | **771250** — content id of the attached fi

version number of the file

To find the directory where attachments for a particular space are stored, go to `<confluence url>/admin/findspaceattachments.jsp` and enter a space key. It will return the directory on the file system where attachments for that space are stored.

File D in the above diagram is stored in a slightly different structure. Files that are not conceptually within a space replace the level 2 - 4 directories with a single directory called 'nonspaced'. Examples of such files are the global site logo and attachments on unsaved content.

Confluence Data Model

This document provides a diagram of the Confluence schema and a conceptual overview of the data model.

Notes:

- The Hibernate mapping files are the authoritative reference for the Confluence data model. These are the `*.hbm.xml` files which you will find in the main Confluence JAR file (`<CONFLUENCE-INSTALLATION>\confluence\WEB-INF\lib\confluence-5.1.1.jar`).
- The tables, columns and other attributes are likely to change with each major release of Confluence. To find the exact DDL of your Confluence site, please run a query after installation.

> **On this page:**
> - Database diagrams
> - Database tables and references
> - Authentication
> - Content
> - Clustering
> - System information
> - Spaces
> - Appearance
> - Miscellaneous

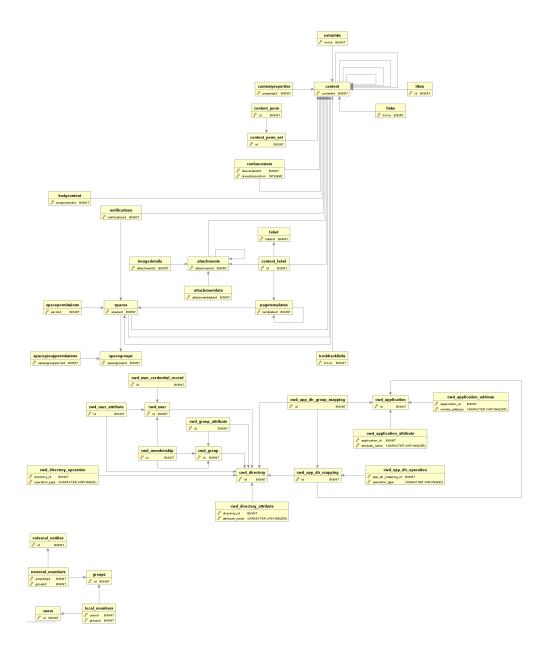### Database diagrams

Detailed diagrams

The following SVG images (Scalable Vector Graphics) include all the tables in the Confluence database. Click the links below to open the images in your browser, or download the SVG files for later use. You can use the browser's zoom (Ctrl++ or Cmd++) to see more detail in the diagrams:

- ConfluenceTables-KeysOnly.svg – Shows all tables, with primary keys only for each table.
- ConfluenceTables-AllColumns.svg – Shows all tables, and all columns for each table.

Overview diagram

This image shows the core tables. Note that the image is very large. You may need to download it (right-click on the image) and view it in an image viewer. Alternatively, use the SVG images linked in the previous section.

⌄ Click here to show/hide the image...

Database tables and references

Expand the link below to see a table of the primary and foreign keys for each table.

˅ Click here to show/hide the table...

Primary key table name	Primary key column name	Foreign key table name	Foreign key column name	Foreign
AO_9412A1_AOUSER	ID	AO_9412A1_USER_APP_LINK	USER_ID	fk_ao_94
attachments	attachmentid	attachmentdata	attachmentid	fk9dc3e3
attachments	attachmentid	attachments	prevver	fk54475f
attachments	attachmentid	content_label	attachmentid	fkf0e743
attachments	attachmentid	imagedetails	attachmentid	fka76804
content	contentid	attachments	pageid	fk54475f
content	contentid	bodycontent	contentid	fka898d4
content	contentid	confancestors	ancestorid	fk9494e2

content	contentid	confancestors	descendentid	fk9494e2
content	contentid	content	prevver	fk6382c0
content	contentid	content	parentid	fk6382c0
content	contentid	content	parentcommentid	fk6382c0
content	contentid	content	pageid	fk6382c0
content	contentid	content_label	contentid	fkf0e743
content	contentid	content_perm_set	content_id	fkbf45a7
content	contentid	contentproperties	contentid	fk984c5e
content	contentid	extrnlnks	contentid	fk97c10f
content	contentid	likes	contentid	fk4514b9
content	contentid	links	contentid	fk451579
content	contentid	notifications	pageid	fk594acc
content	contentid	spaces	homepage	fk922824
content	contentid	spaces	spacedescid	fk922824
content	contentid	trackbacklinks	contentid	fkf6977a
content_perm_set	id	content_perm	cps_id	fkbd74b3
cwd_app_dir_mapping	id	cwd_app_dir_group_mapping	app_dir_mapping_id	fk_app_c
cwd_app_dir_mapping	id	cwd_app_dir_operation	app_dir_mapping_id	fk_app_c
cwd_application	id	cwd_app_dir_group_mapping	application_id	fk_app_c
cwd_application	id	cwd_app_dir_mapping	application_id	fk52050e
cwd_application	id	cwd_application_address	application_id	fk_applic
cwd_application	id	cwd_application_attribute	application_id	fk_applic
cwd_directory	id	cwd_app_dir_group_mapping	directory_id	fk_app_c
cwd_directory	id	cwd_app_dir_mapping	directory_id	fk_app_c
cwd_directory	id	cwd_directory_attribute	directory_id	fk_direct
cwd_directory	id	cwd_directory_operation	directory_id	fk_direct
cwd_directory	id	cwd_group	directory_id	fk_direct
cwd_directory	id	cwd_group_attribute	directory_id	fk_group
cwd_directory	id	cwd_user	directory_id	fk_user_
cwd_directory	id	cwd_user_attribute	directory_id	fk_user_
cwd_group	id	cwd_group_attribute	group_id	fk_group
cwd_group	id	cwd_membership	parent_id	fk_paren
cwd_group	id	cwd_membership	child_group_id	fk_child_
cwd_user	id	cwd_membership	child_user_id	fk_child_

cwd_user	id	cwd_user_attribute	user_id	fk_user_
cwd_user	id	cwd_user_credential_record	user_id	fk76f874
external_entities	id	external_members	extentityid	fkd8c8d8
groups	id	external_members	groupid	fkd8c8d8
groups	id	local_members	groupid	fk6b8fb4
keystore	keyid	trustedapp	public_key_id	fkddb119
label	labelid	content_label	labelid	fkf0e743
os_group	id	os_user_group	group_id	fk932472
os_user	id	os_user_group	user_id	fk932472
pagetemplates	templateid	content_label	pagetemplateid	fkf0e743
pagetemplates	templateid	pagetemplates	prevver	fkbc7ce9
spacegroups	spacegroupid	spacegrouppermissions	spacegroupid	fk487d95
spacegroups	spacegroupid	spaces	spacegroupid	fk922824
spaces	spaceid	content	spaceid	fk6382c0
spaces	spaceid	notifications	spaceid	fk594acc
spaces	spaceid	pagetemplates	spaceid	fkbc7ce9
spaces	spaceid	spacepermissions	spaceid	fkd33f23
trustedapp	trustedappid	trustedapprestriction	trustedappid	fke8496b
users	id	local_members	userid	fk6b8fb4

The following sections describe the principal tables involved in each logical area of Confluence – authentication, content, system information, and so on.

Authentication

This section describes the tables involved in user authentication, which is implemented via the Atlassian Crowd framework embedded in Confluence.

Table	Description
cwd_user	Information for each user in Confluence.
cwd_group	The groups to which users can belong.
cwd_membership	Mapping the membership of users to groups.
cwd_directory	The user directories in your Confluence site. Examples of directories are the Confluence internal directory, or an LDAP directory.
cwd_application	The applications (JIRA, Confluence, and so on) defined in the authentication framework.

Content

This section describes the tables involved in storing content. Content is the information that Confluence users are storing and sharing.

Table	Description
`attachmentdata`	The binary data for attached files. This table is only used when Confluence is configured to store attachments in the database. Otherwise, attachments are stored in the local file system.
`attachments`	Metadata for the files attached to Confluence pages.
`bodycontent`	The content of Confluence pages. No version information or other metadata is stored here. That is all in the `content` table.
`content`	A persistence table for the `ContentEntityObject` class of objects. The subclass is indicated by the `contenttype` column.
`content_label`	Arbitrary text labels for content.
`label`	The other half of the `content_label` system.
`content_perm`	Content-level permissions objects.
`content_perm_set`	A one-to-many mapping for content items and their permissions, with added metadata.
`pagetemplates`	The back end of the templates feature.
`likes`	The pages and other content liked by a particular user.
`follow_connections`	A mapping of users who are following other users.

Clustering

The following table contains information about clustered Confluence sites.

Table	Description
`clustersafety`	Normally, this table only contains one row. The value of the `safetynumber` is what Confluence uses to find out whether another Confluence site is sharing its database without being part of the cluster.

System information

These tables store data related to the status and configuration of the Confluence site.

Table	Description
`confversion`	Used by the upgrade system to determine what to expect from the database, so as to negotiate upgrades.
`plugindata`	A record of the plugins that have been installed, and when. `data` is a blob of the actual plugin JAR file. This is principally cluster-related.

Spaces

This table is related to the management of spaces.

Table	Description
`spaces`	Information about the spaces themselves: key, human-friendly name and numeric ID.

Appearance

The following table contains information about the look and feel of your Confluence site.

Table	Description
`decorator`	The custom display templates used to customise Velocity layouts.

Miscellaneous

This section includes other tables worth commenting on.

Table	Description
`os_propertyentry`	Arbitrary association of entities and properties.
`bandana`	A catch-all persistence layer. This table contains things like user settings and space- and global-level configuration data, and is used as storage by plugins such as the Dynamic Task List plugin. Essentially, for storing arbitrary data that doesn't fit anywhere else.
`extrnlnks`	Referral links.
`hibernate_unique_key`	Used by the high/low ID generator – the subsystem which generates our primary keys. If you interfere with this table, you may not be able to create objects in Confluence.
`indexqueueentries`	Manages full-content indexing across the system. The table generally contains the last 12 hours (approximately) of updates, to allow re-syncing of cluster nodes after restarts.
`keystore`	Used by the trusted apps framework to store the server's private key, and other servers' public keys.
`links`	Tracks links within the server (that is, across and within spaces).
`notifications`	Stores page- and space-level watches.
`trackbacklinks`	Trackback links.
`confancestors`	Used to speed up permissions checks, by allowing quick lookup of all a page's ancestors.

Finding Unused Spaces

Sometimes, you want to know what is *not* being used. It's great to know what's getting most attention, but what about stagnant pages, or even entire spaces that are no longer active?

While viewing space activity can provide hints, it doesn't always provide enough detail. It is possible to find out this information directly from the database.

The following query identifies the last date on which content was modified in each space within a single Confluence instance:

```
SELECT spaces.spacename, MAX(content.lastmoddate)
FROM content, spaces
WHERE content.spaceid = spaces.spaceid
GROUP BY spaces.spacename;
```

It returns a list of space names, and the last date and time at which any content was added or changed.

Alternatively, this query identifies spaces where the content hasn't changed since a specified date:

```
SELECT spaces.spacename
FROM content, spaces
WHERE content.spaceid = spaces.spaceid
GROUP BY spaces.spacename
HAVING MAX(content.lastmoddate) < '2006-10-10';
```

The result is a simple list of space names.

Data Import and Export

Confluence administrators and users can import data into Confluence from a number of sources. The permissions required differ, depending on the scope of the import. See Import Content Into Confluence.

You can also export Confluence content to various formats. See Export Content to Word, PDF, HTML and XML.

> **Related pages:**
> * Managing Confluence Data
> * Confluence Administrator's Guide

Import a Text File

Confluence allows you to import text files from a directory on the Confluence server, and convert them into Confluence pages. Each file is imported as a separate Confluence page with the same name as the file.

> * The text file may contain plain text, HTML or Confluence Storage Format
> * You need to be part of the `confluence-administrators` group or a System Administrator to import text files
> * You can import pages from disk into site spaces, but not into personal spaces
>   Please see Spaces for information about differences between site spaces and personal spaces.

> **Related pages:**
> * Import Content Into Confluence
> * Site Backup and Restore

To make sure Confluence maintains the formatting of the text document, add <pre> to the beginning and </pre> to the end. This will let Confluence know that it should treat the text as pre-formatted.

If you're working in a Unix-like environment, you can add the opening and closing tags to all files in a particular directory by following these steps:

1. Go to the directory containing the files
2. Run the following command in the terminal:

```
for i in $(ls); do echo "<pre>" >> m$i; cat $i >> m$i; echo "</pre>" >>
m$i; mv m$i $i; done
```

To import text files:

1. Go to the space and choose **Space tools** > **Content Tools** from the bottom of the sidebar
2. Choose **Import**
3. Type the directory path into the **Import directory** box
4. Select **Trim file extensions** to remove file extensions from the page titles when converting the files to Confluence pages
   The Confluence pages will take their titles from the files' names (including their extensions). To avoid having page titles with a suffix like '.txt' check this box.
5. Select **Overwrite existing pages** if you want to replace existing Confluence pages with the same title with the one you're importing

6. Choose **Import**

If your space uses the Documentation theme:

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

   Note: The **Space Admin** option appears only if you have space admin permissions, or if you're part of the 'confluence-administrators' group.
2. Choose **Import** from the space administration options
3. Follow the steps above

Confluence will place the new pages at the root of the space.

Screenshot: Importing text files

You can use this action to import text files from a directory on the Confluence server.

These text files become pages in Confluence, with the following features:

- The page title is taken from the filename
- The content is the entire page body

Import directory [   ]

Trim file extensions ☐

Overwrite existing pages ☐

Import   Cancel

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 Con figuring Confluence instead.

Confluence is a J2EE web application. On the client side, users access Confluence primarily via a web browser.
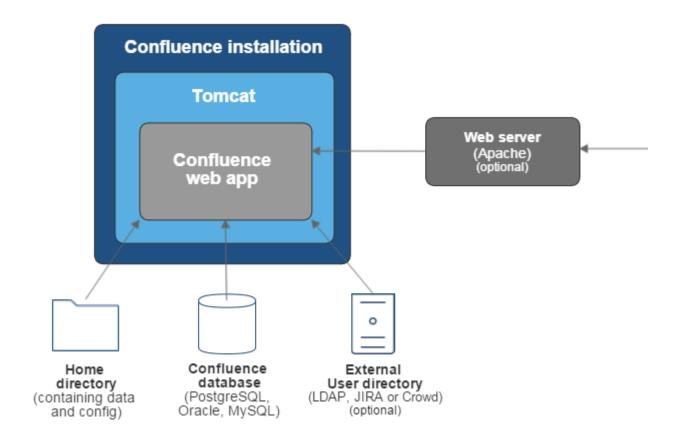
This section contains the following guidelines:

- Confluence Home and other important directories
- Application Server Configuration
- Web Server Configuration
- Starting Confluence Automatically on System Startup

Diagram: A Confluence installation

Related pages:

- Getting Started as Confluence Administrator
- Supported Platforms

Confluence Home and other important directories

Confluence installation directory

The 'Confluence Installation directory' is the directory where Confluence was installed. This directory is also sometimes called the 'Confluence Install directory'.

Important files in the installation directory:

- `bin/setenv.bat` or `bin/setenv.sh`
  This file is used to edit CATALINA_OPTS memory and garbage collection settings and define system properties.
- `confluence/WEB-INF/classes/conflu ence-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

Confluence home directory

The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. Another term for 'Home directory' would be 'data directory'.

Finding the home directory

The location of the Confluence home directory is defined when you install Confluence. This location is stored in the `confluence-init.properties` file, which is located in the `confluence/WEB-INF/classes` dire ctory of your Confluence Installation directory.

When Confluence is running you can find the location of the home directory in

> **General Configuration** > **System Information** > **Confluence Information - Confluence Home**.

If your Confluence instance is clustered, you will also have a shared home directory which will contain some data (such as attachments and backups) that would otherwise reside in the home directory.

Contents of the home directory

The Confluence home directory contains some of the configuration data used by Confluence. This section outlines the purpose of the files and directories in the Confluence home directory.

File or directory	Purpose
`confluence.cfg.xml`	This file contains all of the information necessary for Confluence to start up, such as: • Product license • Context path • Database details, such as location and connection pool settings • Paths to important directories
`attachments/`	This directory contains every version of each attachment stored in Confluence. You can specify an alternative directory for attachment storage by setting the `attachments.dir` property in `confluence.cfg.xml`.
`backups/`	Confluence will place its daily backup archives in this directory, as well as any manually generated backups. Backup files in this directory take the following form `daily-backup-YYYY_MM_DD.zip` You can specify an alternative directory for backups by setting the `daily.backup.dir` property in `confluence.cfg.xml`.
`bundled-plugins/`	Confluence 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.
`database/`	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.dir` property in `confluence.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 the `plugin-cache` directory. This directory is updated as plugins are installed and uninstalled from the system and is completely repopulated from the database every time Confluence is restarted. Removing plugins from this directory does not uninstall them.
`resources/`	The `resources` directory stores any space logos used in your Confluence instance. Space logos are stored in directories named with the spacekey.
`temp/`	The `temp` directory is used for runtime functions such as exporting, importing, file upload and indexing. Files in this directory are temporary and can be safely removed when Confluence is offline. A daily job within Confluence deletes files that are no longer needed. An alternative directory may be specified for temporary data by setting the `webwork.multipart.saveDir` property in `confluence.cfg.xml`.
`thumbnails/`	Stores temporary files for image thumbnails. This directory is essentially a thumbnail cache, and files deleted from this directory will be regenerated the next time the image is accessed.

Changing the location of the home directory

When Confluence first starts up, it reads the `confluence-init.properties` file to determine where to look for the Home directory.

To change the location of the home directory edit the `confluencee.home` property in the `confluence-init.properties` file as follows:

- **Windows**
  In Windows, the path `C:\confluence\data` would be written as:
  `confluence.home=C:/confluence/data`
  Note that all backslashes (\) are written as forward slashes (/)
- **Linux / Solaris**
  On any Linux-based system, the property is defined using the normal directory syntax:
  `confluence.home=/var/confluence/`

Symbolic links

There can be no symbolic links within the Confluence home directory. You must define an absolute path. If disk space is an issue, place the entire `confluence.home` directory on a disk partition where there is enough space. The absolute path of generated files (such as exports) is compared with the absolute path of the `confluence.home` directory when constructing URLs. When a sub-directory has a different path, the URL will be incorrect, and you may receive "Page not found" errors. These measures are in place to prevent "directory traversal" attacks.

Fixing the Confluence Configuration

The Confluence configuration file: `confluence-cfg.xml` inside the home directory may contain references to the original location of your Confluence home. You will need to edit this file to update these references to also point to the new location. The two properties in this file that need to change are:

- `daily.backup.dir` if you have not configured your backups to be placed elsewhere already
- `hibernate.connection.url` if you are using the embedded HSQL database.

Database

All other data, including page content, is kept in the database. If you installed Confluence as a trial, or chose to use the embedded HSQL database during setup, the database will store its files under `database/` in

the Confluence Home Directory. Otherwise, the database management system you are connecting to is responsible for where and how your remaining data is stored.

Temp directory

The temp directory is configured in the Java runtime and some Confluence components write temporary files or lockfiles into this directory.

The temp directory is located in the installation directory as /temp.

To change the location of this directory, start the Java Virtual Machine in which confluence is running with the argument:

`-Djava.io.tmpdir=/path/to/your/own/temp/directory`.

Application Server Configuration

The following pages contain information about configuring your application server for Confluence:
- Configuring URL Encoding on Tomcat Application Server
- Managing Application Server Memory Settings
- Switching to Apache Tomcat
- Java Policy Settings for Enterprise or Webhosting Environments

Configuring URL Encoding on Tomcat Application Server

Application servers may have different settings for character encodings. We strongly recommend **UTF-8** where possible.

By default, Tomcat uses ISO-8859-1 character encoding when decoding URLs received from a browser. This can cause problems when Confluence's encoding is UTF-8, and you are using international characters in the names of attachments or pages.

To configure the URL encoding in Tomcat:

1. Edit `conf/server.xml` and find the line where the `Coyote HTTP Connector` is defined. It will look something like this, possibly with more parameters:

   ```
   <Connector port="8090"/>
   ```

2. Add a `URIEncoding="UTF-8"` property to the connector:

   ```
   <Connector port="8090" URIEncoding="UTF-8"/>
   ```

3. Restart Tomcat

If you are using mod_jk

You should apply the same URIEncoding parameter as above to the AJP connector if you are using mod_jk, and add the following option to your Apache mod_jk configuration:

```
<Connector port="8009" protocol="AJP/1.3" URIEncoding="UTF-8"/>
```

```
JkOptions +ForwardURICompatUnparsed
```

More information using Apache with Tomcat

For comprehensive examples of how to use Tomcat and Apache with Confluence, see Running Confluence behind Apache.

Managing Application Server Memory Settings

The minimum and maximum JVM heap space allocated to the application server affects performance. Confluence administrators may wish to modify this value from the defaults depending on their server load. This document only provides guidelines rather than rules, so administrators optimising for performance should use this document as a starting point only.

> For a comprehensive overview of memory management, and memory tuning in Confluence under Sun JRE, please read Garbage Collector Performance Issues

Testing For Optimum Memory Settings

In the general case, both JIRA & Confluence users will benefit from setting the minimum and maximum values identical. In larger installations, there is benefit to memory tuning, if there is a perceived performance issue. If you are experiencing Out of Memory Heap errors, try increasing the -Xmx and -Xms values for your installation to see if this resolves or helps resolve your issue. It's best to increase in small increments (eg 512mb at a time), to avoid having too large a heap, which can cause different problems. If increasing the memory does not help, please lodge a support ticket as there may be other factors contributing.

Memory usage is most likely to be maximised under peak load, and when creating a site XML backup. In many cases, the backup can be the cause of the OOM, so increase -Xmx values and verify if a backup was occurring at the time of OOM. A quick rule of thumb for gauging the success of a memory adjustment is using simple anecdotal evidence from users. Is it snappier? The same? How does it handle while a backup is occurring?

> Atlassian recommends in normal use, to disable the XML backup and use a Production Backup Strategy .

- If you normally perform manual XML site backups on your server, test your maximum memory requirements by performing a site XML backup while the server is under maximum load
- If you do not create manual XML site backups, simply monitor the server while under maximum load

Applying Memory Settings

See How to fix out of memory errors by increasing available memory.

Related Topics

- Garbage Collector Performance Issues
- How to fix out of memory errors by increasing available memory
- Server Hardware Requirements Guide
- Performance Tuning
- Troubleshooting Slow Performance Using Page Request Profiling
- Tomcat JVM options and Modify the Default JVM Settings

Switching to Apache Tomcat

Apache Tomcat is the only application server supported for Confluence. To move Confluence from an application server (e.g. WebSphere) to Tomcat using the same database, follow the instructions below.

Follow these instructions:

- 1. Before You Start
- 2. Backing Up
- 3. Switching Application Servers
- 4. Applying Customisations
    - Confluence Server
    - Plugins
    - Look and Feel
    - Performance
    - Advanced Customisations
- 5. Testing Confluence

1. Before You Start

1. The following instructions will only work if you are running the same major version of Confluence on both application servers. If you are running different major versions of Confluence, you will need to upgrade Confluence before you can switch to Tomcat.
2. Note that you need current software maintenance, as the process for changing application servers involves installing Confluence.
3. If the environment (e.g. the database system, the operating system and so on) that you are running Confluence in has changed, please ensure it still complies with the Confluence System Requirements.
4. If you are using an external database, familiarise yourself with all known issues for your specific database. Also make sure the Confluence database connector principal (the database user login) has sufficient permissions to modify the database schema.
5. Note any customisations that you have made to Confluence, e.g. enabled/installed plugins, modified layouts, custom themes, etc. You will need to reapply these after you have switched to Tomcat. You can view the list of customisations in the Reapplying Customisations section below.
6. We recommend that you **do not** run any other applications in your Tomcat application server that is running Confluence, to prevent performance issues.

2. Backing Up

Before you switching to Tomcat, you must back up the following:

1. **Back up your Confluence Home directory.** The location of the Home directory is stored in a configuration file called `confluence-init.properties`, which is located inside the `confluence/WEB-INF/classes` directory in your Confluence Installation directory.
2. **Back up your database.** Perform a manual backup of your external database before proceeding with the upgrade and check that the backup was created properly. If you are not a database expert or unfamiliar with the backup-restore facilities of your database, you should try to restore the backup to a different system to ensure that the backup worked before proceeding. This recommendation is not specific to Confluence usage, but it is good practice to ensure that your database backup is not broken.
   ℹ The 'embedded database' is the HSQLDB database supplied with Confluence for evaluation purposes, you don't need to back it up since it is stored in the home directory. But you should not use this database for production systems anyway, so if you happen to accidentally still use HSQLDB in a production system, please migrate to a proper database **before** the upgrade.
3. **Back up your Confluence Installation directory**.

3. Switching Application Servers

1. Install Confluence on your new application server. We recommend that you download Confluence as a stand alone archive (rather than using the installer).
   ⚠ **Important:** At this stage, just unzip the standalone file to a location of your choice - this will be your new installation directory. You should not start Confluence until directed to in step 3 below.
2. Copy the following files from your old Confluence installation to your new one:
   - `<CONFLUENCE_INSTALL>\confluence\WEB-INF\classes\confluence-init.properties`
   - `<CONFLUENCE_INSTALL>\confluence\WEB-INF\classes\atlassian-user.xml`
   - `<CONFLUENCE_INSTALL>\confluence\WEB-INF\classes\osuser.xml` *(copy this over if you are using JIRA user management)*
   - `<CONFLUENCE_INSTALL>\confluence\WEB-INF\classes\seraph-config.xml` *(copy this over if you using custom SSO)*
   - `<CONFLUENCE_INSTALL>\confluence\WEB-INF\web.xml` *(copy this over if you have previously modified it, e.g. to configure a datasource)*
3. Start Confluence (make sure you shutdown the old server before you start up the new one)
4. If you are running the new application server on a different machine to the old one, carry out the following actions as soon as you start the new server:
   - Re-index your data.
   - Make sure that the attachments location is valid for the new server.
5. If you have applied special settings to their Confluence server and/or Confluence look and feel, you will need to reapply these customisations as described in below.

4. Applying Customisations

After switching to Tomcat, you need to review any customisations and other special configurations you previously used for your Confluence instance, and re-apply if necessary. This section also contains some Tomcat-specific customisations that you may wish to considering applying, if you haven't used Confluence with Tomcat before.

> **Before you apply customisations**
> Please ensure that your Confluence installation works correctly on Tomcat without any customisations before you apply any of customisations listed below. This will make it easier to identify problems, if you run into trouble during the switch to Tomcat.

Confluence Server

- For long-term use, we recommend that you configure Confluence to start automatically when the operating system restarts. For Windows servers, this means configuring Confluence to run as a Windows service.
- If you are using the Confluence edition and you have previously defined a **CATALINA_HOME environment variable**, please check that it points to the correct path for the new Confluence Tomcat server.
- If you were previously running **Confluence on a non-standard port**, edit your new `<Installation-Directory>\conf\server.xml` file as described in Change listen port for Confluence.

Plugins

- If you were previously using any **plugins**, install the latest compatible version and disable any plugins that are incompatible with your new instance of Confluence. The easiest way to do this is to use the Universal Plugin Manager in the Confluence Administration Console.

Look and Feel

- If you are using any **customised themes**, please check that they are displaying as expected. Some further customisation may be required to ensure compatibility with your new version of Confluence.
- If you had previously customised the **default site or space layouts**, you will need to reapply your changes to the new defaults as described here. Please do not just copy your VM (velocity) files across. Ensure that Confluence works without your custom layouts then apply the layout via the Confluence Administration console.

Performance

- If the load on your Confluence instance is high, you may need more simultaneous connections to the database. Read more about this in the Performance Tuning guide.
- If you had previously modified the **memory flags** (`Xms` and `Xmx`) in either the `<Installation-Directory>\bin\setenv.sh` or the `setenv.bat` file, you may want to make the modifications in your new installation. The parameters are specified in the `CATALINA_OPTS` variable. See How to fix out of memory errors by increasing available memory for more information.

Advanced Customisations

- If you were previously running **Confluence over SSL**, you will need to reapply your configuration as described in Running Confluence Over SSL or HTTPS.
- If you were using a custom SSO authenticator, change seraph-config.xml to the correct authenticator.
- If you had changed the Confluence interface text, you will need to copy over the ConfluenceActionSupport.properties file.
- If you had previously modified the Confluence **source code**, you will need to reapply your changes to the new version.

5. Testing Confluence

Make sure you **test Confluence on the new server** before deploying it in production.

The Working with Confluence Logs document contains the locations for the application logs, if you need to refer to them.

Java Policy Settings for Enterprise or Webhosting Environments

Confluence relies on a number of Java libraries. Some of these libraries make use of features of the Java

language that may be restricted by Java security policies.

This does not normally cause any problems. The default security configuration of most application servers will happily run Confluence. However, in some shared-hosting or enterprise environments, security settings may be such that Confluence cannot function.

> **Related pages:**
>
> - Application Server Configuration
> - Confluence Administrator's Guide

When you attempt to run Confluence, you may get the following error:

```
java.security.AccessControlException: access denied (java.lang.RuntimePermission
accessDeclaredMembers)
        at
java.security.AccessControlContext.checkPermission(AccessControlContext.java(Compil
ed Code))
        at
java.security.AccessController.checkPermission(AccessController.java(Compiled
Code))
        at java.lang.SecurityManager.checkPermission(SecurityManager.java(Compiled
Code))
```

The permissions required by Confluence to run are detailed in the sample policy file below. You may need to give this information to your systems administrator so that they can be deployed with the Confluence application.

```
grant codeBase "file:${catalina.home}/webapps/confluence/-" {
permission java.security.AllPermission;
};

grant {
permission java.lang.RuntimePermission "accessDeclaredMembers";
permission java.lang.reflect.ReflectPermission "suppressAccessChecks";
permission java.lang.RuntimePermission "defineCGLIBClassInJavaPackage";
};
```

Web Server Configuration

- Configuring Web Proxy Support for Confluence
- Running Confluence behind Apache
  - General Apache Configuration Notes
  - Using Apache with mod_proxy
  - Using Apache with virtual hosts and mod_proxy
  - Using Apache with mod_jk
  - Using mod_rewrite to Modify Confluence URLs
  - Configuring Apache to Cache Static Content via mod_disk_cache

Configuring Web Proxy Support for Confluence

Some of Confluence's macros, such as {rss} and {jiraissues} need to make web requests to remote servers in order to retrieve data. If Confluence is deployed within a data centre or DMZ, it may not be able to access the Internet directly to make these requests. If you find that the {rss} macro does not work, ask your network administrator if Confluence needs to access the Internet through a web proxy.

Configuring an outbound HTTP proxy in Confluence

Proxy support is configured by passing certain system properties to the Java Virtual Machine on startup. These properties follow the conventions defined by Oracle:

- http.proxyHost
- http.proxyPort (default: 80)
- http.nonProxyHosts (default: <none>)

- https.proxyHost
- https.proxyPort

At a minimum, you need to define `http.proxyHost` to configure an HTTP proxy, and `https.proxyHost` to configure an HTTPS proxy. System property configuration is described in the Configuring System Properties.

Properties `http.proxyHost` and `http.proxyPort` indicate the proxy server and port that the http protocol handler will use, and `https.proxyHost` and `https.proxyPort` indicate the same for the https protocol handler.

```
-Dhttp.proxyHost=proxy.example.org -Dhttp.proxyPort=8080
-Dhttps.proxyHost=proxy.example.org -Dhttps.proxyPort=8080
```

Property `http.nonProxyHosts` indicates the hosts which should be connected to directly and not through the proxy server. The value can be a list of hosts, each separated by a pipe character | . In addition, a wildcard character (asterisk) * can be used for matching. For example:

```
-Dhttp.nonProxyHosts=*.foo.com|localhost
```

Note: You may need to escape the pipe character | in some command-line environments.

If the `http.nonProxyHosts` property is not configured, all web requests will be sent to the proxy.

Please note that any command line parameters set are visible from the process list, and thus anyone who has the 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.proxyPort=yourProxyPort
http.proxyUser=yourUserName
http.proxyPassword=yourPassword
https.proxyHost=yourProxyURL
https.proxyPort=yourProxyPort
https.proxyUser=yourUserName
https.proxyPassword=yourPassword
```

Configuring HTTP proxy authentication

Proxy authentication is also configured by providing system properties to Java in your application server's configuration file. Specifically, the following two properties:

- http.proxyUser – username
- http.proxyPassword – secret

HTTP proxy (Microsoft ISA) NTLM authentication

Confluence supports NTLM authentication for outbound HTTP proxies when Confluence is running on a Windows server.

This means that the {rss} and {jiraissues} macro will be able to contact external websites if requests have to go through a proxy that requires Windows authentication. This support is not related to logging in Confluence users automatically with NTLM, for which there is a user-contributed authenticator available.

To configure NTLM authentication for your HTTP proxy, you need to define a domain system property, `http.au th.ntlm.domain`, in addition to the properties for host, port and username mentioned above:

```
-Dhttp.auth.ntlm.domain=MYDOMAIN
```

Configuring authentication order

Sometimes multiple authentication mechanisms are provided by an HTTP proxy. If you have proxy authentication failure messages, you should first check your username and password, then you can check for this problem by examining the HTTP headers in the proxy failure with a packet sniffer on the Confluence server. (Describing this is outside the scope of this document.)

To set the order for multiple authentication methods, you can set the system property `http.proxyAuth` to a comma-separated list of authentication methods. The available methods are: ntlm, digest and basic; this is also the default order for these methods.

For example, to attempt Basic authentication before NTLM authentication, and avoid Digest authentication entirely, you can set the `http.proxyAuth` property to this value:

```
-Dhttp.proxyAuth=basic,ntlm -Dhttps.proxyAuth=basic,ntlm
```

Troubleshooting

1. There's a diagnostic jsp file in CONF-9719 for assessing the connection parameters.
2. '**Status Code [407]**' errors are described in APR-160.
3. Autoproxies are not supported. See CONF-16941.

Running Confluence behind Apache

> This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

Introduction

For improved performance in high-load environments, you should run Confluence behind a web server. In general, web server caching and thread management is far superior to that provided by your application server's HTTP interface.

To run Confluence behind the Apache httpd web server, there are two main configuration options: **mod_jk** or **mod_proxy**.

Connection type	Unique features	Common features to both mod_proxy and mod_jk
mod_proxy (also known as reverse proxy)	• **recommended connection method** • simple HTTP proxy to application server • works with all application servers • if application paths are consistent, there is minimal load on the web server	• application paths must be consistent to avoid complex and slow URL rewriting • works with name-based virtual hosting, both on web server and app server • web server keeps a pool of connections to application server
mod_jk (also known as AJP)	• uses the AJP binary protocol • provides failover (and load balancing, which Confluence supports only with a clustered license) • only works with some application servers (typically Tomcat) • if application paths are consistent, there is some load on the web server to translate requests to AJP	

Configuration Guides

> Please choose one configuration. Trying to configure for both mod_proxy and mod_jk will only lead to confusion and tears.

- **Using Apache with mod_proxy**
- **Using Apache with mod_jk**
- **Using Apache with virtual hosts and mod_proxy**

Mod_jk2 not supported

The misleadingly-named mod_jk2 is an older method of connecting to Tomcat from Apache. Since mod_jk2 is no longer supported by the Apache Foundation, we do not support this configuration, and are not updating our mod_jk2 documentation. Mod_jk2 also has unresolved problems with Unicode URLs; you need to use either mod_proxy or mod_jk for international characters to work correctly in Confluence.

Caching static content via mod_disk_cache

To improve performance of a large Confluence site, we recommend that you move the caching of static content from the JVM into Apache. This will prevent the JVM from having a number of long running threads serving up static content. See Configuring Apache to Cache Static Content via mod_disk_cache.

Other related documentation

- Configuring Tomcat's URI encoding
- Running Confluence Over SSL or HTTPS

General Apache Configuration Notes

On this page:

- Prefer Apache mod_deflate to Confluence's built-in gzip implementation
- Ensure keepalive is enabled
- Enable keepalive for recent MSIE user agents

Prefer Apache mod_deflate to Confluence's built-in gzip implementation

1. Disable gzip in confluence. See Compressing an HTTP Response within Confluence.
2. Enable gzip compression in Apache. For RedHat distributions this can be achieved by adding the following lines:

```
AddOutputFilterByType DEFLATE text/html text/plain text/xml text/css
application/x-javascript

# ensure sensible defaults
DeflateBufferSize 8192
DeflateCompressionLevel 4
DeflateMemLevel 9
DeflateWindowSize 15
```

Ensure keepalive is enabled

```
KeepAlive On
```

Enable keepalive for recent MSIE user agents

The standard Apache SSL configuration is very conservative when it comes to MSIE and SSL. By default all keepalives are disabled when using HTTPS with MSIE. While MSIE will always be *special*, the issues with SSL

and MSIE have been solved since Service Pack 2 for Windows XP, released over 4 years go. For anyone using an XP machine SP2 or above, it is safe to allow keepalive for MSIE 6 and above.

Remove the following lines:

```
SetEnvIf User-Agent ".*MSIE.*" \
    nokeepalive ssl-unclean-shutdown \
    downgrade-1.0 force-response-1.0
```

Add these in their place:

```
BrowserMatch "MSIE [1-5]" nokeepalive ssl-unclean-shutdown downgrade-1.0
force-response-1.0
BrowserMatch "MSIE [6-9]" ssl-unclean-shutdown
```

RELATED TOPICS

Running Confluence behind Apache
Configuring Tomcat's URI encoding
Running Confluence Over SSL or HTTPS

Using Apache with mod_proxy

This page describes how to integrate Confluence into an Apache website using mod_proxy.

There are some common situations where you might use the configuration:

- You have an existing Apache-based website, and want to add Confluence to the mix (for example, **http://www.example.com/confluence**).
- You have two or more Java applications, each running in their own application server on different ports, for example, **http://example:8090/confluence** and **http://example:8080/jira**. By setting up Apache with mod_proxy, you can have both available on the regular HTTP port (80) – for example, at **http://www.example.com/confluence** and **http://www.example.com/jira**. This allows each application to be restarted, managed and debugged separately.

Note: This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

Base configuration

> In these examples, we use the following:
>
> **http://www.example.com/confluence** - your intended URL
>
> **http://example:8090** - the hostname and port Confluence is currently installed to
>
> **/confluence** - the intended context path (the part after hostname and port)
>
> Please substitute the examples below with your intended URL's in your own server. Copy/pasting these suggestions will not work on your server.

Set the context path

Set your Confluence application path (the part after hostname and port). To do this in Tomcat (bundled with Confluence), edit `conf/server.xml`, locate the "Context" definition:

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

and change it to:

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

Then restart Confluence, and ensure you can access it at **http://example:8090/confluence**

Set the URL for redirection

Set the URL for redirection. In the same `conf/server.xml` file, locate this code segment:

```
<Connector port="8090" maxHttpHeaderSize="8192"
           maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
           enableLookups="false" redirectPort="8443" acceptCount="100"
           connectionTimeout="20000" disableUploadTimeout="true" />
```

And append the last line:

```
<Connector port="8090" maxHttpHeaderSize="8192"
           maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
           enableLookups="false" redirectPort="8443" acceptCount="100"
           connectionTimeout="20000" disableUploadTimeout="true"
           proxyName="www.example.com" proxyPort="80" />
```

If this isn't working for you and you're using SSL, try adding a scheme attribute to your Connector tag: **scheme= "https"**.

Now we have two options:

- If you want a URL like **http://www.example.com/confluence**, follow the simple configuration.
- If you want a URL like **http://confluence.example.com**, go to the complex configuration.

Simple Configuration

Configure mod_proxy

Now enable mod_proxy in Apache, and proxy requests to the application server by adding the example below to your Apache httpd.conf (note: the files may be different on your system; See Integrating JIRA with Apache for the process for Ubuntu/Debian layout):

Apache 2.2

```
# Put this after the other LoadModule directives
LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so

# Put this in the main section of your configuration (or desired virtual host, if
using Apache virtual hosts)
ProxyRequests Off
ProxyPreserveHost On

<Proxy *>
    Order deny,allow
    Allow from all
</Proxy>

ProxyPass /confluence http://app-server.internal.example.com:8090/confluence
ProxyPassReverse /confluence http://app-server.internal.example.com:8090/confluence
<Location /confluence>
    Order allow,deny
    Allow from all
</Location>
```

Apache 2.4

```
# Put this after the other LoadModule directives
LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so

# Put this in the main section of your configuration (or desired virtual host, if
using Apache virtual hosts)
ProxyRequests Off
ProxyPreserveHost On

<Proxy *>
 # Auth changes in 2.4 - see
http://httpd.apache.org/docs/2.4/upgrading.html#run-time
    Require all granted
</Proxy>

ProxyPass /confluence http://app-server.internal.example.com:8090/confluence
ProxyPassReverse /confluence http://app-server.internal.example.com:8090/confluence
<Location /confluence>
 # Auth changes in 2.4 - see
http://httpd.apache.org/docs/2.4/upgrading.html#run-time
    Require all granted
</Location>
```

˅ Note to Windows Users

It is recommended that you specify the absolute path to the `mod_proxy.so` and `mod_proxy_http.so` files
.

Complex configuration

Complex configuration involves using the mod_proxy_html filter to modify the proxied content en-route. This is required if the Confluence path differs between Apache and the application server. For example:

Externally accessible (Apache) URL	http://confluence.example.com/
Application server URL	http://app-server.internal.example.com:8090/confluence/

Notice that the application path in the URL is different in each. On Apache, the path is /, and on the application server the path is /confluence.

> For this configuration, you need to install the mod_proxy_html module, which is not included in the standard Apache distribution.
>
> Alternative solutions are discussed below.

Apache 2.2

```
# Put this after the other LoadModule directives
LoadModule proxy_module modules/mod_proxy.so
LoadModule proxy_http_module modules/mod_proxy_http.so
LoadModule proxy_html_module modules/mod_proxy_html.so

<VirtualHost *>
    ServerName confluence.example.com

    # Put this in the main section of your configuration (or desired virtual host,
if using Apache virtual hosts)
    ProxyRequests Off
    ProxyPreserveHost On

    <Proxy *>
        Order deny,allow
        Allow from all
    </Proxy>

    ProxyPass / http://app-server.internal.example.com:8090/confluence
    ProxyPassReverse / http://app-server.internal.example.com:8090/confluence

    ProxyHTMLURLMap / /confluence/

    <Location />
        Order allow,deny
        Allow from all
    </Location>
</VirtualHost>
```

Apache 2.4

```
# Put this after the other LoadModule directives
LoadModule proxy_module modules/mod_proxy.so
LoadModule proxy_http_module modules/mod_proxy_http.so
LoadModule proxy_html_module modules/mod_proxy_html.so


<VirtualHost *>
    ServerName confluence.example.com

    # Put this in the main section of your configuration (or desired virtual host,
if using Apache virtual hosts)
    ProxyRequests Off
    ProxyPreserveHost On

    <Proxy *>
  # Auth changes in 2.4 - see
http://httpd.apache.org/docs/2.4/upgrading.html#run-time
     Require all granted
    </Proxy>

    ProxyPass / http://app-server.internal.example.com:8090/confluence
    ProxyPassReverse / http://app-server.internal.example.com:8090/confluence

    ProxyHTMLURLMap / /confluence/

    <Location />
  # Auth changes in 2.4 - see
http://httpd.apache.org/docs/2.4/upgrading.html#run-time
     Require all granted
    </Location>
</VirtualHost>
```

The ProxyHTMLURLMap configuration can become more complex if you have multiple applications running under this configuration. The mapping should also be placed in a Location block if the web server URL is a subdirectory and not on a virtual host. The Apache Week tutorial has more information how to do this.

Final Configuration Steps

Restart your Apache server

This is needed to pick up on the new configuration. This can be done by running the following on your command line/terminal/shell:

```
sudo apachectl graceful
```

Disable HTTP Compression

Having compression run on both the proxy and Tomcat can cause problems integrating with other Atlassian applications, such as JIRA. Please disable HTTP compression as per our Compressing an HTTP Response within Confluence docs.

Set the Confluence Base URL

The last stage is to set the Base URL to the address you're using within the proxy. In this example, it would be **http://www.example.com/confluence**

Adding SSL

If you're running Apache in front of Tomcat, it's a good idea to terminate your SSL configuration at Apache, then

forward the requests to Tomcat over HTTP. You can set up Apache to terminate the SSL connection and use the ProxyPass and ProxyPassReverse directives to pass the connection through to Tomcat (or the appropriate application server) which is running Confluence.

1. Create a new SSL host by creating a virtual host on 443
2. The standard http connection on apache could be used to redirect to https if you want or it could just be firewalled.
3. Within the VirtualHost definition:
   a. define the SSL options (SSLEngin and SSLCertificateFile)
   b. define the ProxyPass and ProxyPassReverse directives to pass through to Tomcat.

Most of the relevant Apache Config:

```
Listen 443

NameVirtualHost *:443
<VirtualHost *:443>
    SSLEngine On
    SSLCertificateFile /etc/apache2/ssl/apache.pem
    ProxyPass / http://localhost:8090/
    ProxyPassReverse / http://localhost:8090/
</VirtualHost>
```

Apart from the Apache configuration there are a couple of things you will need to do before you get your server working:

1. You will have to change your base URL to point to https addresses. See the documentation on configurin g the server base URL.
2. We need to set up the connector to use https. In your installation directory, edit the file server.xml and add this attributes to your connector:

```
proxyName="proxy.example.com" proxyPort="443" scheme="https" secure="true"
```

More information

- The mod_proxy_html site has documentation and examples on the use of this module in the complex configuration.
- Apache Week has a tutorial that deals with a complex situation involving two applications and ProxyHTMLURLMap.
- Using Apache with virtual hosts and mod_proxy shows how to configure the special case where you want JIRA and Confluence running on separate application servers on virtual host subdomains.

Alternatives

If Tomcat is your application server, you have two options:

- use mod_jk to send the requests to Tomcat
- use Tomcat's virtual hosts to make your Confluence application directory the same on the app server and the web server, removing the need for the URL mapping.

If your application server has an AJP connector, you can:

- use mod_jk to send the requests to your application server.

Using Apache with virtual hosts and mod_proxy

Note: This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

Introduction

The Apache web server is often used in front of an application server to improve performance in high-load environments. Mod_proxy simply redirects requests for certain URLs to another web server, so it typically requires no additional configuration on the application server.

This page documents a very common configuration request: configuring a JIRA application and Confluence on two Apache virtual hosts, running on different application servers. This is just a special case of mod_proxy configuration.

You can use virtual hosts in your application server if you want to run JIRA applications and Confluence on the same application server.

Apache configuration

For this configuration to work properly, the application paths must be the same on both the application servers and the web server. For both JIRA and Confluence below, this is /.

JIRA external URL	http://jira.example.com/
JIRA application server URL	http://jira-app-server.internal.example.com:8080/
Confluence external URL	http://confluence.example.com/
Confluence application server URL	http://confluence-app-server.internal.example.com:8090/

Add the following to your Apache httpd.conf:

```
# Put this after the other LoadModule directives
LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so

# Put this with your other VirtualHosts, or at the bottom of the file
NameVirtualHost *
<VirtualHost *>
    ServerName confluence.example.com

    ProxyRequests Off
    <Proxy *>
        Order deny,allow
        Allow from all
    </Proxy>

    ProxyPass / http://confluence-app-server.internal.example.com:8090/
    ProxyPassReverse / http://confluence-app-server.internal.example.com:8090/
    <Location />
        Order allow,deny
        Allow from all
    </Location>
</VirtualHost>
<VirtualHost *>
    ServerName jira.example.com

    ProxyRequests Off
    <Proxy *>
        Order deny,allow
        Allow from all
    </Proxy>

    ProxyPass / http://jira-app-server.internal.example.com:8080/
    ProxyPassReverse / http://jira-app-server.internal.example.com:8080/
    <Location />
        Order allow,deny
        Allow from all
    </Location>
</VirtualHost>
```

Points to note:

- ProxyPass and ProxyPassReverse directives send traffic from the web server to your application server.
- The application path is the same on the application server and on the web server (both are /).
- Because the above configuration uses name-based virtual hosting, you must configure your DNS server to point both names (jira.example.com, confluence.example.com) to your web server.

More information

For different ways to configure mod_proxy, see Using Apache with mod_proxy.

If you use Tomcat, mod_jk provides a different way of connecting Apache via AJP. You can also use the above configuration with just one application server if you use Tomcat's virtual hosts.

Using Apache with mod_jk

- The preferred configuration is Using Apache with mod_proxy. This works with any application server, and together with `mod_proxy_html` allows complex URL rewriting to deal with different application paths on the web server and the application server.
- This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you

> use it at your own risk.

Introduction

The Apache web server is often used in front of an application server to improve performance in high-load environments. Mod_jk allows request forwarding to an application via a protocol called AJP. Configuration of this involves enabling mod_jk in Apache, configuring a AJP connector in your application server, and directing Apache to forward certain paths to the application server via mod_jk.

Mod_jk is sometimes preferred to mod_proxy because AJP is a binary protocol, and because some site administrators are more familiar with it than with mod_proxy..

The scope of this documentation is limited to configuring the AJP connector in Tomcat 5.x. Other application servers may support AJP connectors; please consult your application server documentation for instructions on how to configure it.

The configuration below assumes your Confluence instance is accessible on the same path on the application server and the web server. For example:

Externally accessible (web server) URL	http://www.example.com/confluence/
Application server URL (HTTP)	http://app-server.internal.example.com:8090/confluence/

The AJP connection of the application server is set to: app-server.internal.example.com:8009.

Configuring mod_jk in Apache

The standard distribution of Apache does not include mod_jk. You need to download it from the JK homepage a nd put the mod_jk.so file in your Apache modules directory.

Next, add the following in **httpd.conf** directly or included from another file:

```
# Put this after the other LoadModule directives
LoadModule jk_module modules/mod_jk.so

# Put this in the main section of your configuration (or desired virtual host, if
using Apache virtual hosts)
JkWorkersFile conf/workers.properties
JkLogFile logs/mod_jk.log
JkLogLevel info

JkMount /confluence worker1
JkMount /confluence/* worker1
```

Configuring workers.properties

Create a new file called 'workers.properties', and put it in your Apache conf directory. (The path for workers.properties was one of the configuration settings above.)

```
worker.list=worker1

worker.worker1.host=app-server.internal.example.com
worker.worker1.port=8009
worker.worker1.type=ajp13
```

Tomcat 5.x configuration

In Tomcat 5, the AJP connector is enabled by default on port 8009. An absolutely minimal Tomcat server.xml is below for comparison. The relevant line is the Connector with port 8009 – make sure this is uncommented in

your server.xml.

```
<Server port="8000" shutdown="SHUTDOWN">
  <Service name="Catalina">

    <!-- Define a HTTP/1.1 Connector on port 8090 -->
    <Connector port="8090" />

    <!-- Define an AJP 1.3 Connector on port 8009 -->
    <Connector port="8009" protocol="AJP/1.3" />

    <Engine name="Catalina" defaultHost="localhost">
      <Host name="localhost" appBase="webapps">
 <Context path="/confluence" docBase="/opt/webapps/confluence-2.2/confluence"/>
        <Logger className="org.apache.catalina.logger.FileLogger"/>
      </Host>
    </Engine>
  </Service>
</Server>
```

Points to note:

- the Connector on port 8009 has protocol of "AJP/1.3". This is critical.
- the Context path of the Confluence application is "/confluence". This must match the path used to access Confluence on the web server.
- we recommend keeping your application Contexts outside the server.xml in Tomcat 5.x. The above example includes them for demonstration only.

Improving the performance of the mod_jk connector

The most important setting in high-load environments is the number of processor threads used by the Tomcat AJP connector. By default, this is 200, but you should increase it to match Apache's maxThreads setting (256 by default):

```
<Connector port="8009" minSpareThreads="5" maxThreads="256" protocol="AJP/1.3" />
```

All the configuration parameters for the AJP connector are covered in the Tomcat documentation.

Ensuring UTF-8 compatibility

If you have problems downloading attachments with non-ASCII characters in the filename, add the following to your Apache configuration:

```
JkOptions +ForwardURICompatUnparsed
```

And specify UTF-8 as the URIEncoding in the AJP connector configuration:

```
<Connector port="8009" protocol="AJP/1.3" URIEncoding="UTF-8" />
```

These settings are discussed further on Configuring Tomcat's URI encoding.

More information

The Tomcat JK website has complete documentation on workers.properties and Apache configuration. You can also find information there on how to use mod_jk with IIS.

Note: In IIS proxy server the maximum file upload is 30mb by default. Contact your server administrator if you need to upload a bigger file.

Using mod_rewrite to Modify Confluence URLs

Note: This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

Confluence requires URL rewriting for proper functionality, if Confluence is accessible via different domain names. If Confluence is configured for multiple domains *without* URL rewriting, you will experience an array of problems. See Various Issues Caused when Server Base URL Does Not Match the URL Used to Access Confluence.

An example of why you may want to access Confluence from different domains:

- From an internal network:
  http://wiki
- The externally visible domain:
  http://wiki.domain.com

Using URL rewriting to access Confluence over multiple domains

To configure Confluence over multiple domains:

1. Add a DNS entry mapping http://wiki to the externally visible IP address of the Confluence server.
2. Set Confluence's server base URL to http://wiki.domain.com.
3. Add Apache HTTP proxy, using the instructions from Running Confluence behind Apache.
4. Add the mod_rewrite module to change the URL.

Further information

You may be interested in the UrlRewriteFilter that is Java web filter that works in a similar way of the **Apache's mod_rewrite**.

Configuring Apache to Cache Static Content via mod_disk_cache

To improve performance of a large Confluence site, we recommend that you move the caching of static content from the JVM into Apache. This will prevent the JVM from having a number of long running threads serving up static content.

Static content in Confluence includes most JavaScript, CSS and image files which are included with the application or an installed plugin. This content will be cached by Apache in this configuration. User-provided content like space logos, attachments or embedded images are not considered static content and will not be cached.

Note: This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

Configuring Apache mod_disk_cache

To configure Apache to cache static Confluence content:

1. Add a `mod_disk_cache` stanza to the virtual host configuration:

   ```
   <IfModule mod_disk_cache.c>
     # "/s" is where Confluence serves "static" stuff. Instruct Apache
   to cache it:
     CacheEnable disk /s
     CacheIgnoreHeaders Set-Cookie
     CacheRoot "/var/cache/mod_proxy"
   </IfModule>
   ```

2. Configure Apache to load `mod_disk_cache`. For example, in our server configuration this is done in `/etc/httpd/conf/httpd.conf`:

```
LoadModule disk_cache_module modules/mod_disk_cache.so
```

3. Restart Apache after both modifications are complete.

Notes

- Please refer to the Apache documentation for mod_disk_cache.
- If you encounter problems where users are served stale content, you may need to purge the Apache cache directory (/var/cache/mod_proxy in the above configuration) after a Confluence or plugin upgrade. This is a simple 3 step process:
  - Shut down Apache.
  - Clear the cache directory. For example: sudo rm -r /var/cache/mod_proxy/*
  - Restart Apache.
- Ensure that you are running the htcacheclean daemon in order to prevent excessive use of disk space. In our situation we ran it like this:

```
sudo htcacheclean -d30 -n -t -p /var/cache/mod_proxy -l 512M
```

This will purge content once the cache reaches 512M every 30 minutes. See the Apache documentation for htcacheclean for details of the options.

Starting Confluence Automatically on System Startup

You can configure Confluence to start automatically on system startup, allowing it to recover automatically after a reboot.

Content by label

There is no content with the specified labels

Start Confluence Automatically on Linux

On Linux/Solaris, the best practice is to install, configure and run each service (including Confluence) as a dedicated user with only the permissions they require.

To install, configure and run Confluence automatically on Linux/Solaris:

1. Create a confluence user for instance, using the following command:

```
sudo useradd --create-home -c "Confluence role account" confluence
```

2. Create a directory to install Confluence into:

```
sudo mkdir /usr/local/confluence
sudo chown confluence: /usr/local/confluence
```

3. Log in as the confluence user to install Confluence:

```
sudo su - confluence
cd /usr/local/confluence/
tar zxvf /tmp/confluence-5.6.4.tar.gz
ln -s confluence-5.6.4/ current
```

4. Edit
   <<CONFLUENCE_INSTALL_DIRECTORY>>/confluence/WEB-INF/classes/confluence-init.properties file,
   and set confluence.home=/usr/local/confluence/<Confluence_Data_Home> (ensure you have removed
   the comment '#')
5. Then back as root, create the file /etc/init.d/confluence (code shown below), which will be
   responsible for starting up Confluence after a reboot (or when manually invoked).
   ⚠ If you are running Ubuntu Jaunty (or later) do not perform this step. Please use the instructions further
   down this page.

```
#!/bin/sh -e
# Confluence startup script
#chkconfig: 2345 80 05
#description: Confluence

# Define some variables
# Name of app ( JIRA, Confluence, etc )
APP=confluence
# Name of the user to run as
USER=confluence
# Location of application's bin directory
CATALINA_HOME=/usr/local/confluence/current
# Location of Java JDK
export JAVA_HOME=/usr/lib/jvm/java-7-oracle

case "$1" in
  # Start command
  start)
    echo "Starting $APP"
    /bin/su -m $USER -c "$CATALINA_HOME/bin/start-confluence.sh &> /dev/null"
    ;;
  # Stop command
  stop)
    echo "Stopping $APP"
    /bin/su -m $USER -c "$CATALINA_HOME/bin/stop-confluence.sh &> /dev/null"
    echo "$APP stopped successfully"
    ;;
   # Restart command
   restart)
        $0 stop
        sleep 5
        $0 start
        ;;
  *)
    echo "Usage: /etc/init.d/$APP {start|restart|stop}"
    exit 1
    ;;
esac

exit 0
```

6. Make this file executable:

```
sudo chmod +x /etc/init.d/confluence
```

7. Set this file to run at the appropriate runlevel. For example, use `sudo chkconfig --add confluence` on Redhat-based systems, `sudo update-rc.d confluence defaults` or `rcconf` on Debian-based systems.
8. You should now be able to start Confluence with the init script. A successful startup output typically looks like this:

```
$ sudo /etc/init.d/confluence start
Starting Confluence:
If you encounter issues starting up Confluence, please see the
Installation guide at
http://confluence.atlassian.com/display/DOC/Confluence+Installation
+Guide
Using CATALINA_BASE:   /usr/local/confluence/current
Using CATALINA_HOME:   /usr/local/confluence/current
Using CATALINA_TMPDIR: /usr/local/confluence/current/temp
Using JRE_HOME:        /usr/lib/jvm/java-1.7.0-oracle
done.
```

You should then see this running at http://<server>:8090/
ⓘ The port for this will be whatever is defined in your Confluence `server.xml` file.

Adding Confluence as a service for Ubuntu Jaunty (or later)

To continue configuring Confluence to start automatically as a service on Ubuntu Jaunty (or later):

1. After logging in as the `confluence` user to install Confluence, create `start` and `stop` scripts in `/usr/local/confluence`:

Example `start` script:

```
#!/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 `stop` script:

```
#!/bin/bash
export JAVA_HOME=/usr/lib/jvm/java-7-oracle-1.7.0.71/
export JDK_HOME=/usr/lib/jvm/java-7-oracle-1.6.0.71/
cd /usr/local/confluence/current/bin
./shutdown.sh
```

2. Make both of these scripts executable. For example, by issuing the command: `sudo chmod a+x /usr/local/confluence/start /usr/local/confluence/stop`.
3. Karmic and later: Create two text files in `/etc/init/` called `confluence-up.conf` and `confluence-down.conf`:

confluence-up:

```
start on runlevel [2345]

script

 date >> /tmp/confluence-startup.out
 exec sudo -u confluence /usr/local/confluence/start >>
/tmp/confluence-startup.out 2>&1

end script
```

confluence-down:

```
start on runlevel [16]

expect fork
respawn

exec sudo -u confluence /usr/local/confluence/stop >>
/tmp/confluence-shutdown.out 2>&1
```

... and make them readable to all users:
```
sudo chmod a+r /etc/init/confluence-up.conf /etc/init/confluence-down.conf
```

1. Jaunty, Intrepid: Create two text files in `/etc/event.d/` called `confluence-up` and `confluence-down`:

confluence-up:

```
start on runlevel 2
start on runlevel 3
start on runlevel 4
start on runlevel 5

exec sudo -u confluence /usr/local/confluence/start >>
/tmp/confluence-startup.out 2>&1
```

confluence-down:

```
start on runlevel 1
start on runlevel 6

exec sudo -u confluence /usr/local/confluence/stop >>
/tmp/confluence-shutdown.out 2>&1
```

... and make them readable to all users:

```
       sudo chmod a+r /etc/event.d/confluence-up /etc/event.d/confluence-down
```
RELATED TOPICS

[Starting Confluence Automatically on System Startup](#)

Start Confluence Automatically on Windows as a Service

For long-term use, we recommend that you configure Confluence to start automatically when the operating system restarts. For Windows servers, this means configuring Confluence to run as a Windows service.

There are two ways to install the Confluence distribution as a service: using the [Confluence installer](#) or manually as described below.

On this page:

- [Reasons for Starting Confluence as a Service](#)
- [Changing the User Running the Service](#)
- [Manually Installing the Confluence Distribution as a Service](#)
- [Managing Confluence as a Service](#)
- [Upgrading Confluence](#)
- [Troubleshooting Confluence while Running as a Windows Service](#)
- [Requesting Support](#)

> **Problem with 64-bit Windows**
> If you are running 64-bit Windows, please note that you may encounter problems with Apache Tomcat running as a Windows service if you are using a 64-bit JDK. Refer to our [knowledge base article](#) for more information.

Reasons for Starting Confluence as a Service

Installation as a Windows service offers these advantages:

- Reduced risk of shutting down Confluence by accident (If you start Confluence manually, a console window opens and there is a risk of someone accidentally shutting down Confluence by closing the window).
- Automated Confluence recovery after server restart.
- Improved troubleshooting through logging server output to file.

You can read more about Windows services in the [Microsoft Developer Network](#).

Changing the User Running the Service

If you wish to run the service as a non-administrator user for security, or if you are using network drives for backups, attachments or indexes, you can run the service as another user. To change users, open the Apache Tomcat Confluence properties, go to the 'Log On' tab and enter the required username and password. Go to your Windows Control Panel -> User Accounts and confirm that the user has write permissions for the [<CONFLUENCE-INSTALL>](#) and [<CONFLUENCE-HOME>](#) directories, and all subfolders. Note that any network drives must be specified by UNC and not letter mappings (eg. `\\backupserver\confluence` not `z:\confluence`).

For more detail, see [Creating a Dedicated User Account on the Operating System to Run Confluence](#).

Manually Installing the Confluence Distribution as a Service

From your Windows-based server:

1. Open a command prompt in the `<CONFLUENCE-INSTALL>/bin` directory.
2. Confirm that the JAVA_HOME variable is set to the JDK base directory with the command:

```
echo %JAVA_HOME%
```

Note that any directory in the path with spaces (eg. `C:\Program Files` must be converted to its eight-character equivalent (e.g. `C:\Progra~1`).

3. If you are installing Confluence on a Windows 2008 server, be sure to run the command prompt using 'run as administrator'. (Otherwise running 'service.bat', as described in the next step, will fail.)
4. Use the following command to install the service with default settings:

```
service.bat install Confluence
```

NB: This will create a service called **Apache Tomcat Confluence**.
5. Now, to have the service start automatically when the server starts, run:

```
tomcat7 //US//Confluence --Startup auto
```

(if you are using a different version of tomcat, specify your tomcat version, e.g. `tomcat6`)
6. If you have a less than a 1024 megabytes of memory, skip this step. For users with large Confluence installations, you can increase the maximum memory Confluence can use. (The default is 1024MB). For example, you can set the maximum memory to 2048 megs using:

```
tomcat7 //US//Confluence --JvmMx 2048
```

7. If you do not have any JVM parameters that you pass to your distribution of Confluence, you can skip this step. If you do, add them to the service using:

```
tomcat7 //US//Confluence ++JvmOptions="-Djust.an.example=True"
```

8. For further configuration options, please refer to the Tomcat Windows Service How-To guide
9. Go to your Windows Control Panel -> Administrative Tools -> Services -> Apache Tomcat Confluence and right-click on Properties to verify the settings are correct.

   ℹ️ Confluence is now installed as a service, but will not automatically start up until the next server reboot
10. Start the Confluence service with the command:

```
net start Confluence
```

Managing Confluence as a Service

You can manage the Confluence service from the command prompt.

- Stop Confluence with:

```
net stop Confluence
```

- Uninstall the Confluence service with:

```
service.bat remove Confluence
```

Upgrading Confluence

After upgrading Confluence, you can either uninstall and reinstall the Windows service or change the StartPath parameter to your new folder. Refer to the Tomcat documentation for help.

Troubleshooting Confluence while Running as a Windows Service

- Check the Knowledge Base articles:
  - Getting 'The image file tomcat6.exe is valid, but is for a machine type other than the current machine'
  - Confluence Does Not Start Due to Windows Firewall
  - Unable to start Confluence Windows service after allocating JVM memory
  - Unable to Configure Confluence to Run as a Service on Tomcat 5
  - Unable to Install Service on Windows Vista

- If none of the above solves your problem, please refer to the complete list of known issues in our Knowledge Base.

- When investigating memory issues or bugs, it may be useful to view information from Confluence's garbage collection. To turn on the verbose garbage collection, use the command:

```
tomcat6 //US//Confluence
++JvmOptions="-Xloggc:<CONFLUENCE-INSTALL>\logs\atlassian-gc.log"
```

- The Confluence 2.9 installer does not work when installed as service, due to a missing semi-colon in `service.bat`. Please refer to reported issue CONF-12785.

- You can use a Sysinternals tool called `Procmon.exe` from the The Microsoft Windows Sysinternals Team, to check that the error occurred at the specific time when the Confluence service started. You need to match the time when Tomcat failed, as captured by this tool, against the time in the Windows Event Viewer.

> **Note**
> We do not recommend that you run this tool for too long as it may disrupt other Atlassian applications. Once you have captured the required information you will need to press `Ctrl + E` to stop capturing.

Requesting Support

If, after following the troubleshooting guide above, you still cannot make Confluence run as a Windows Service or if there is an error when setting the JVM configuration for the service, you can create a support request.

Please provide the following information when creating your support request, because we will need it to assist you:

- Are you running a 32 bit or 64 bit Windows?
- Give us the result of running `java -version` from Windows command line console.
- A screen shot of your Windows Registry setting for Tomcat.
- If you have modified **service.bat**, please give us a copy of this file for review.
- What application server are you using? eg. Are you using the Confluence distribution?
- Your atlassian-confluence.log file.

RELATED TOPICS

Starting Confluence Automatically on System Startup
How to fix out of memory errors by increasing available memory

Configuring Confluence

This section focuses on settings and configurations within the Confluence application.

For guidelines on external configuration, see Configuring a Confluence Environment.

- Viewing System Information
- Configuring the Server Base URL
- Configuring the Confluence Search and Index
- Configuring Mail
- Configuring Character Encoding

Related pages:

- Customising your Confluence Site
- Confluence Administrator's Guide

- Other Settings
- Configuring System Properties
- Working with Confluence Logs
- Configuring Confluence Security
- Scheduled Jobs
- Configuring the Whitelist
- Configuring the Time Interval at which Drafts
  are Saved

Viewing System Information

The System Information screen provides information about Confluence's configuration, and the environment in which Confluence has been deployed.

To view your system information:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **System Information** in the left-hand panel.

Notes:

- The handy **memory graph** helps you keep track of Confluence's memory usage.
- Your system configuration information is helpful to Atlassian Support when diagnosing errors you may face using Confluence. When logging a **support request** or **bug report**, please provide as much detail as possible about your installation and environment.

> **Related pages:**
>
> - Cache Statistics
> - Live Monitoring Using the JMX Interface
> - Tracking Customisations Made to your
>   Confluence Installation

Live Monitoring Using the JMX Interface

With the JMX interface (introduced in Confluence 2.8), you can monitor the status of your Confluence instance in real time. This will provide you with useful data such as the resource usage of your instance and its database latency, allowing you to diagnose problems or performance issues. To read the JMX data, you will need to use a JMX client.

Disable JMX

> If you experience any problems during Confluence startup that are related to JMX, it is possible to disable the JMX registration process. Please place jmxContext.xml in your `<confluence-install>/confluence/WEB-INF/classes` folder to do so.

What is JMX?

JMX (Java Management eXtensions) is a technology for monitoring and managing Java applications. JMX uses objects called MBeans (Managed Beans) to expose data and resources from your application.

1. Enabling JMX Remote with Tomcat

By default, Confluence uses the Apache Tomcat web server. To use JMX, you must enable it on your Tomcat server, by carrying out the steps under the Apache Tomcat documentation, entitled Enabling JMX Remote. With those steps completed, restart your Tomcat server.

For the stand-alone, add the startup parameter -Dcom.sun.management.jmxremote to setenv.sh or setenv.bat. See instructions for Increasing JIRA application memory - enter it in the same place as PermGen Memory.

2. Selecting your JMX Client

You need to use a JMX client in order to view the JMX output from Confluence. JConsole is a readily available JMX client that is included with the supported Java Developer Kit (version 5 onwards). The full name is the '*Java*

Monitoring and Management Console', but we will refer to it as JConsole for the purposes of this document.

3. Adding the JMX Client to your Path

You must add the location of the JConsole binary file to your path environment variable. As JConsole resides in the 'bin' (binaries) folder under your Java directory, the path should resemble something like this:

```
JDK_HOME/bin/
```

In this example, replace 'JDK_HOME' with the full system path to your Java directory.

4. Configuring JConsole

To configure JConsole:

1. Run the JConsole application.
2. You will be prompted to create a new connection. Choose **remote process** and enter the hostname of your Confluence instance and a port of your choosing.

   > To connect easily, add the startup parameters to setenv.bat or setenv.sh:
   > -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=8086
   > -Dcom.sun.management.jmxremote.authenticate=false
   > Port 8086 is unlikely to be used. Then, connect remotely using port 8086.

   > JConsole, or any JMX client, will not see applications which are not owned by the same user. For example under Windows, if an application is started as a service, it is the System User which owns the process, and not the Current User.

3. Click **Connect**.

Note: Other JMX clients besides JConsole can read JMX information from Confluence.

What can I monitor with JMX?

The JMX interface allows you to see live internal information from your Confluence instance, via the following MBeans:

IndexingStatistics

This MBean shows information related to search indexing.

Property name	Function	Values
Flushing	Shows state of cache (i.e. flushing, or not).	True/False
LastElapsedMilliseconds	Time taken during last indexing.	Milliseconds
LastElapsedReindexing	Time taken during last re-indexing.	Milliseconds
TaskQueueLength	Shows number of tasks in the queue.	Integer

SystemInformation

This MBean shows information related to database latency. It also contains most of the information presented on the System Information page.

Property name	Function	Values
DatabaseExampleLatency	Shows the latency of an example query performed against the database.	Milliseconds

RequestMetrics

This MBean shows information related to system load and error pages served.

Property name	Function	Values
AverageExecutionTimeForLastTenRequests	Average execution time for the last ten requests.	Milliseconds
CurrentNumberOfRequestsBeingServed	Number of requests being served at this instant.	Integer
ErrorCount	Number of times the Confluence error page was served.	Integer
NumberOfRequestsInLastTenSeconds	Obviously, the Number Of Requests In the Last Ten Seconds.	Integer

MailServer-SMTPServer

This MBean shows information related to email dispatch attempts and failures. There will be an MBean for every SMTP Mailserver that has been configured in the Confluence instance.

Property name	Function	Values
EmailsAttempted	The number of email messages Confluence has tried to send.	Integer
EmailsSent	The number of email messages sent successfully.	Integer

MailTaskQueue

This MBean shows information related to the email workload.

Property name	Function	Values
ErrorQueueSize	Number of errors in the queue.	Integer
Flushing	Shows state (i.e. flushing, or not)	True/False
FlushStarted	Time that operation began.	Time
RetryCount	The number of retries that were performed.	Integer
TaskSize	Number of email messages queued for dispatch.	Integer

SchedulingStatistics

This MBean shows information related to current jobs, scheduled tasks and the time that they were last run.

High CPU consuming threads

For Java 1.6, add the Top Threads Plugin to monitor whether CPU is spiking. Download it to a directory and run JConsole like this:
JConsole -pluginpath /pathto/topthreads.jar

This works only with JDK 1.6, but that can be on the remote machine if the server is running a lower version.

> Please note, adding live monitoring to a production instance may itself have an impact on performance.

Tracking Customisations Made to your Confluence Installation

The '**Modification**' section of the Confluence '**System Information**' screen lists the files that have been changed since your Confluence application was installed. You will find this information particularly useful when upgrading Confluence to a new version, because you will need to re-apply all customisations after the upgrade.

To see the modifications made to files in your Confluence installation:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Select '**System Information**' in the 'Administration' section of the left-hand panel.
3. Scroll down to the section titled '**Modification**'.

Screenshot: Modifications tracker on the Confluence System Information screen

Modification	
Modified	decorators/main.vmd, pages/page-breadcrumbs.vm, template/includes/macros.vm, decorators/mail.vmd, decorators/space.vmd, template/includes/personal-sidebar.vm
Removed	No files removed

Notes

- The modification tracker does not detect changes to class files from the `confluence.jar` or other JAR files. If you modify classes, the Confluence modification detection does not report the modification. See issue CONF-20993.

Viewing Site Statistics

Note that the site activity information is **disabled by default**. See notes below.

If enabled, the global activity screen displays statistics on the activity in your Confluence site. These include:

- How many pages and blog posts have been viewed, added or updated over a given period.
- Which spaces are the most popular (most frequently viewed).
- Which spaces are the most active (most frequently edited).
- Which people are the most active contributors/editors of content.

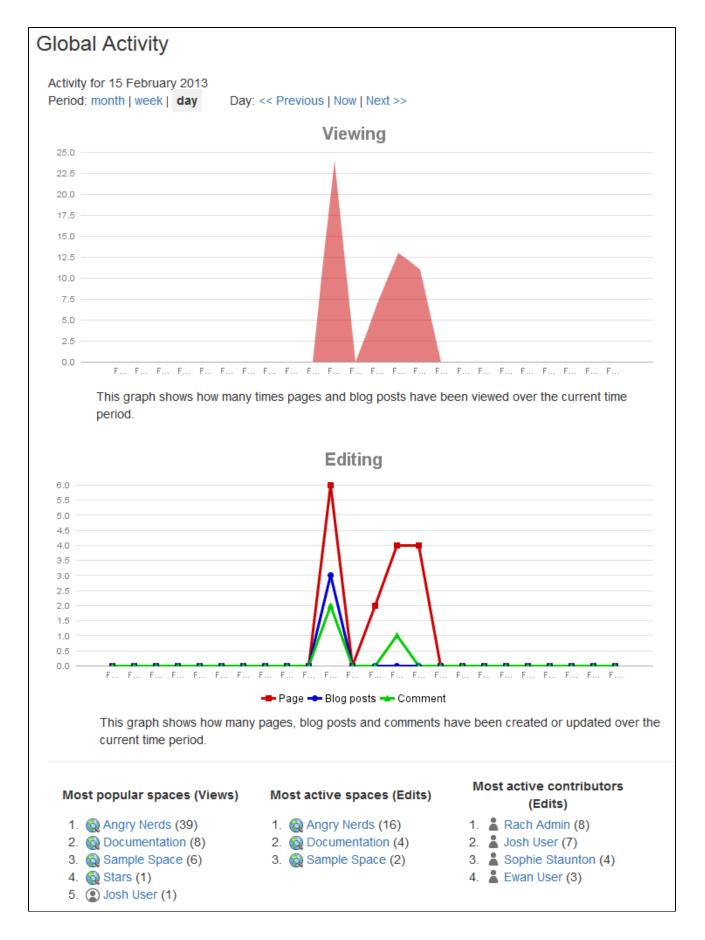**To view the activity on your site:**

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose '**Global Activity**' in the 'Administration' section of the left-hand panel (only appears if enabled - see below).

Screenshot: Global Activity

Related pages:

- How do I get more statistics from Confluence?
- Cache Statistics
- View Space Activity
- Live Monitoring Using the JMX Interface

Global Activity

Activity for 15 February 2013
Period: month | week | **day**   Day: << Previous | Now | Next >>

Viewing

This graph shows how many times pages and blog posts have been viewed over the current time period.

Editing

■ Page ● Blog posts ▲ Comment

This graph shows how many pages, blog posts and comments have been created or updated over the current time period.

Most popular spaces (Views)

1. Angry Nerds (39)
2. Documentation (8)
3. Sample Space (6)
4. Stars (1)
5. Josh User (1)

Most active spaces (Edits)

1. Angry Nerds (16)
2. Documentation (4)
3. Sample Space (2)

Most active contributors (Edits)

1. Rach Admin (8)
2. Josh User (7)
3. Sophie Staunton (4)
4. Ewan User (3)

The top ten most popular and most active pages and/or blog posts will be listed, with a link to each.

Notes

- The Confluence Usage Stats plugin, which provides the 'Global Activity' screen, is known to cause

performance problems on large installations. This plugin is **disabled by default**. A status report on the progress of the performance issues with this plugin is available in this issue: USGTRK-15.
- Your Confluence system administrator can enable the plugin, but please be aware of the possible impact upon your site's performance.
- The plugin is sometimes called 'Confluence Usage Tracking'.
- If your Confluence site is clustered, the global activity information will not be available.

Viewing System Properties

After adding memory, setting a proxy, or changing other Java options, it can be difficult to diagnose whether the system has picked them up. This page tells you how to view the system properties that your Confluence site is using.

You can see the expanded system properties on the 'System Information' screen of the Confluence Administration Console. You do not need to restart Confluence before viewing the information.

To see the system properties recognised by your Confluence installation:

1. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
2. Choose **System Information** in the left-hand panel.
3. Scroll down to the section titled **System Properties**.

Configuring the Server Base URL

The **Server Base URL** is the URL via which users access Confluence. The base URL **must** be set to the same URL by which browsers will be viewing your Confluence site.

Confluence will automatically detect the base URL during setup, but you may need to set it manually if your site's URL changes or if you set up Confluence from a different URL to the one that will be used to access it publicly.

You need to have System Administrator permissions in order to perform this function.

To configure the Server Base URL:

1. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
2. Choose **General Configuration** in the left-hand panel
3. Choose **Edit**
4. Enter the new URL in the **Server Base URL** text box
5. Choose **Save**

Example

If Confluence is installed to run in a non-root context path (that is, it has a context path), then the server base URL should include this context path. For example, if Confluence is running at:

```
http://www.foobar.com/confluence
```

then the server base URL should be:

```
http://www.foobar.com/confluence
```

Notes

- **Using different URLs.** If you configure a different base URL or if visitors use some other URL to access Confluence, it is possible that you may encounter errors while viewing some pages.

- **Changing the context path.** If you change the context path of your base URL, you may also need to edit the web server's `server.xml`file to reflect the new path:
    1. Stop the Confluence server.
    2. Go to your Confluence 'destination directory'. This is the directory where the Confluence installation files are stored. For example, `C:\Program Files\Atlassian\Confluence`. Let's call this directory '`{CONFLUENCE_INSTALLATION}`'.
    3. Edit the configuration file at `{CONFLUENCE_INSTALLATION}\conf\server.xml`.
    4. Change the value of the `path` attribute in the `Context` element to reflect the context path. For example, if Confluence is running at `http://www.foobar.com/confluence`, then your `path` attribute should look like this:

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

    5. Save the file.
- **Proxies**. If you are running behind a proxy, ensure that the proxy name matches the base URL. For example: `proxyName="foobar.com" proxyPort="443" scheme="https"`. This will make sure we are passing the information correctly.
- This information needs to be added in the `Connector` element at `{CONFLUENCE_INSTALLATION}\conf\server.xml`.

Configuring the Confluence Search and Index

Confluence administrators can adjust the behaviour of the Confluence search, and manage the index used by the search.

- Configuring Indexing Language
- Configuring Quick Navigation
- Content Index Administration
- Enabling OpenSearch
- Rebuilding the Ancestor Table
- Setting Up Confluence to Index External Sites
- Setting Up an External Search Tool to Index Confluence

Related pages:

- Search
- Confluence Administrator's Guide

Configuring Indexing Language

Changing the indexing language defined in Confluence may improve the accuracy of Confluence search results, if the majority of the content of your site is in some language other than English. Confluence supports content indexing in English (default), German, Russian, Chinese, CJK, Custom Japanese, French, Brazilian, Czech and Greek.

Related pages:

- Choosing a Default Language
- Installing a Language Pack
- Content Index Administration
- How to rebuild the content indexes from scratch

To configure the indexing language:

1. Choose the **cog icon**

    ⚙ ▾

    , then choose **General Configuration** under Confluence Administration
2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Select the **Indexing Language** from the dropdown list in the **Formatting and International Settings** section.
5. Choose **Save**.

Configuring Quick Navigation

When a user is searching Confluence (see Searchin g Confluence) the quick navigation aid automatically offers a dropdown list of pages and other items, matched by title to the search query. By default, this feature is enabled, with the maximum number of simultaneous quick navigation requests set to 40. These options can be modified as described below.

The maximum number of simultaneous quick navigation requests defines the maximum number of individuals who can use this feature simultaneously on the same Confluence server. If your Confluence server serves a large number of individuals who use this feature regularly, some of whom are being denied access to it, you may wish to increase this value.

To configure quick navigation:

1. Choose the **cog icon**
   ⚙ ▾
   , then choose **General Configuration** under Confluence Administration
2. Choose **Further Configuration** in the left-hand panel.
3. Choose **Edit**.
4. To disable Quick Navigation, deselect the **Quick Navigation** checkbox.
5. To modify the maximum number of simultaneous quick navigation requests, enter the appropriate number in the field beside **Max Simultaneous Requests**.
6. Choose **Save**.

Content Index Administration

The content index, also called the search index, supports Confluence's search functionality. It is also used for a number of related functions such as building email threads in the mail archive, the space activity feature, and lists of recently-updated content. The Gliffy plugin also uses the index for some of its functionality.

For reasons of efficiency, Confluence does not immediately add content to the index. New and modified Confluence content is first placed in a queue and the queue is processed once every five seconds (by default).

On this page:

- Viewing the content index summary
- Rebuilding the search index
- The 'Did You Mean' index is no longer relevant
- Slow reindexing
- Viewing the index browser
- More hints and tips

Viewing the content index summary

To see information about your Confluence site's content indexing:

1. Choose the **cog icon**
   ⚙ ▾
   , then choose **General Configuration** under Confluence Administration
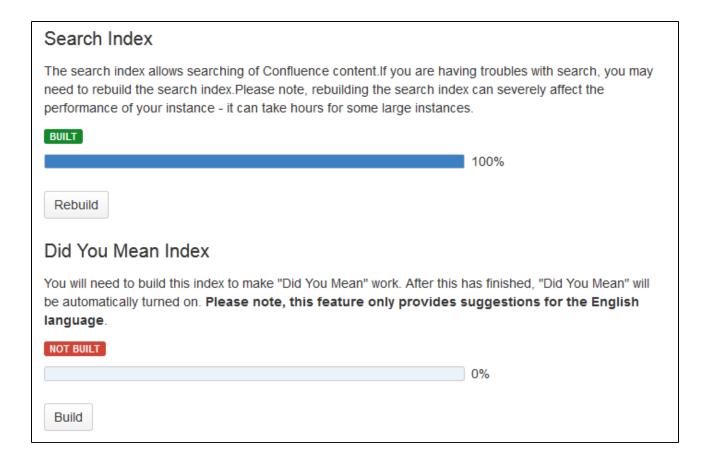2. Choose '**Content Indexing**' under the heading '**Administration**' in the left-hand panel.

Screenshot: Index summary

Search Index

The search index allows searching of Confluence content.If you are having troubles with search, you may need to rebuild the search index.Please note, rebuilding the search index can severely affect the performance of your instance – it can take hours for some large instances.

BUILT

100%

Rebuild

Did You Mean Index

You will need to build this index to make "Did You Mean" work. After this has finished, "Did You Mean" will be automatically turned on. **Please note, this feature only provides suggestions for the English language**.

NOT BUILT

0%

Build

Rebuilding the search index

The search index is maintained automatically, but you may need to rebuild it manually under circumstances such as these:

- Your searching and mail threading are malfunctioning.
- After an upgrade. If a re-index is required after an upgrade, it will be noted in an upgrade subsection of the relevant Confluence Release Notes.

To rebuild the search index:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose '**Content Indexing**' under the heading '**Administration**' in the left-hand panel.
3. Choose the '**Rebuild**' button in either the 'Search Index' section.
   (If the indexes has never been built, its button will indicate 'Build' instead of 'Rebuild.)

Screenshot: Content indexing

Search Index

The search index allows searching of Confluence content.If you are having troubles with search, you may need to rebuild the search index.Please note, rebuilding the search index can severely affect the performance of your instance - it can take hours for some large instances.

IN PROGRESS...

0%   (Elapsed Time: 00:00:00)

Rebuild

Did You Mean Index

You will need to build this index to make "Did You Mean" work. After this has finished, "Did You Mean" will be automatically turned on. **Please note, this feature only provides suggestions for the English language.**

'Did You Mean' index build temporarily disabled while Search Index build is in progress.

NOT BUILT

0%

The 'Did You Mean' index is no longer relevant

The 'Did You Mean' feature is no longer available in Confluence. This index is therefore redundant, and will be removed at some time in the future.

Slow reindexing

Does the reindexing take a long time to complete? The length of time depends on the following factors:

- Number of pages in your Confluence instance.
- Number, type and size of attachments.
- Amount of memory allocated to Confluence.
- Disk throughput.

It may help to increase the heap memory allocation of Confluence by following the instructions here Increasing JIRA application memory. The process is basically the same for Confluence or JIRA applications.

If you are running an older version of Confluence and find that the index rebuild is not progressing, you may need to shut down Confluence, and restart it with the following *Java system property* set: `bucket.indexing.threads.fixed=1`. This will cause the re-indexing to happen in a single thread and be much more stable (but slower).

Viewing the index browser

Confluence uses a search engine called Lucene. If you need to see more details of the indexed pages in your Confluence site, you can download and run Luke. Luke is a development and diagnostic tool that accesses existing Lucene indexes and allows you to display and modify their content in several ways.

Start Luke and use it to open the `index` directory, located in your Confluence Home directory. For example: `c:\confluence\data\confluence-home\index`.

Note: Confluence 5.2 (and later) use Lucene 4.3 (or later). If the Luke library has not been updated to support the latest version of Lucene, you can compile Luke yourself, from the fork on Github – please read the warnings and notes in the README file of that repository.

More hints and tips

- If you are still experiencing problems after performing the above rebuild, the next step might be to remove the index and rebuild it from scratch.
- ⚠ The space activity feature uses the index to store data. If you remove the index file, the existing

activity data will disappear.
- A tip for the development community: If you have the Confluence source, you can look for references to the SmartListManager to find the screens and lists that rely on the content index.

Enabling OpenSearch

With OpenSearch autodiscovery, you can add Confluence search to your Firefox or IE7 search box (see Searching Confluence from your Browser's Search Box). By default, OpenSearch autodiscovery is enabled. This feature can be enabled or disabled as described below.

To enable or disable OpenSearch autodiscovery:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **Further Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Select the **Open Search** checkbox to enable this feature (deselect to disable).
5. Choose **Save**.

> **Related pages:**
> - Search
> - Confluence Administrator's Guide

Rebuilding the Ancestor Table

> This process could take hours in a production instance and new content may not be save-able during this process. **It is highly recommended that you take a full backup of your database and then execute this process during a maintenance window**.

> **Related pages:**
> - Search
> - Confluence Administrator's Guide

In Confluence, the ancestor table defines what pages are ancestors or descendants of other pages (which can be used by search restrictions with the ancestorids restriction). Occasionally, the ancestor table will become out of sync. When this happens, you can rebuild the table to restore everything to normal.

Access this URL:

```
http://yoursite/admin/permissions/pagepermsad
min.action
```

After rebuilding the ancestor table, you'll need to flush the "Inherited Content Permissions" cache in Cache Statistics , otherwise the inherited permissions may not be applied immediately to all pages. You may also need to rebuild the content index so that the permissions take effect in search results.

Screenshot: Page level permissions

Page Level Permissions

[Rebuild Ancestor Table]

Note: There was a known issue with accessing this option from Confluence 5.4.1 to Confluence 5.5.6. See

CONF-32174 - Rebuild Ancestor Table is Inaccessible
RESOLVED
.

Setting Up Confluence to Index External Sites

Confluence cannot easily index external sites due to

technical reasons, but there are two alternatives:

1. Embed External Pages Into Confluence
2. Replace Confluence Search

Technical reasons

Confluence indexes pages using a customised Lucene search engine that returns matching pages, mail and blog posts for which the searcher has view permission. It would require significant source code modifications to enable Confluence to process search results from external pages, as the indexing process has been customised to utilise internal Confluence metadata. Note that users can still index content from new attachment filetypes.

Embedding external pages into Confluence

If you only have a small number of external sites to index, you may prefer to enable the HTML-include Macro and use it embed the external content inside normal Confluence pages.

Replacing the Confluence search

Use your own programmer resources to replace Confluence's internal search with a crawler that indexes both Confluence and external sites. This advanced option is easier than modifying the internal search engine. It requires removing Confluence internal search from all pages and replacing the internal results page with your own crawler front-end.

1. Setup a replacement federated search engine to index the Confluence site, as well as your other sites, and provide the results that way. You would need to host a web crawler, such as these open-source crawlers. Note that you can perform a search in Confluence via the remote API.
2. Replace references to the internal search by modifying the site layout so that it links to your search front-end
3. Host another site containing the search front-end. You may wish to insert it into a suitable context path in your application server so that it appears to be from a path under Confluence. Tomcat sets Confluence's paths from the Confluence install\confluence\WEBINF\web.xml file.

Setting Up an External Search Tool to Index Confluence

Any web crawler can be configured to index Confluence content, for example the Google Search Appliance or similar. If a login is required to view content that will be indexed, you should create a Confluence user specifically for the search crawler to use. Grant this user view rights to all content you wish to index, but deny that user all delete and administration rights. This ensures that an aggressive crawler will not be able to perform actions that could modify the site.

External applications can also use the search function in the Confluence remote API.

Configuring Mail

- Configuring a Server for Outgoing Mail
- Setting Up a Mail Session for the Confluence Distribution
- Configuring the Recommended Updates Email Notification
- The Mail Queue

- Customising the eMail Templates

Configuring a Server for Outgoing Mail

Configuring your Confluence server to send email messages allows your Confluence users to:

- Receive emailed notifications and daily reports of updates.

- Send a page via email.

You can personalise email notifications by configuring the 'From' field to include the name and email address of the Confluence user who made the change.

You need System Administrator permissions in order to configure Confluence's email server settings.

On this page:

- Configuring Confluence to send email messages
- Testing the email settings

Related pages:

- The Mail Queue
- Setting Up a Mail Session for the Confluence Distribution

Configuring Confluence to send email messages

To configure Confluence to send outgoing mail:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Select **Mail Servers** under **Configuration** in the left-hand panel. This will list all currently configured SMTP servers.
3. Click **Add New SMTP Server** (or **edit** an existing server).
4. Edit the following fields as required:
   - **Name**: By default, this is simply 'SMTP Server'.
   - **From Address**: Enter the email address that will be displayed in the 'from' field for email messages originating from this server.
     This field is mandatory. You will not be able to complete the Confluence mail server configuration until this field has been specified.
   - **From Name**: Enter the name that will be displayed in the 'from' field for email messages originating from this server. This is the text which appears before the user's registered email address (in square brackets).
     This field accepts the following variables, which reference specific details defined in the relevant Confluence user's profile:

Variable	Description
${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.
5. Enter your **Hostname, Port, User name** and **Password** details.
   If your SMTP host uses the Transport Layer Security (TLS) protocol select **Use TLS**.
     OR
   Specify the **JNDI location** of a mail session configured in your application server. For more information on how to set up a JNDI mail session, see Setting Up a Mail Session for the Confluence Distribution.

Testing the email settings

A Confluence administrator can test the email server as follows:

1. Set up a mail server as described above.

2. Click **Send Test Email** to check that the server is working. Check that you get the test email in your inbox.
3. You can flush the email queue to send the email message immediately. Go to **Mail Queue**, and click **Flus h Mail Queue**. See The Mail Queue.

A user can test that notifications are working as follows:

1. Go to your user profile (using the **Settings** link) and edit your email preferences. See Email Notifications.
2. Enable **Notify On My Actions**. (By default, Confluence does not send you notifications for your own changes.)
3. Go to a page you wish to get notifications about.
4. Choose **Watch** at the top-right of the page. See Watch Pages, Spaces and Blogs.
5. Edit the page, make a change, and save the page.
6. Check your email inbox. You may need to wait a while for the email message to arrive.

Setting Up a Mail Session for the Confluence Distribution

The simplest way to set up a mail server through the Confluence Administration console. See Configuring a Server for Outgoing Mail.

If you want to add different options or parameters you can also set up a mail session for the Confluence distribution. In the example below we'll set up Gmail.

To set up a mail session for the Confluence distribution:

1. Stop Confluence.
2. Move (don't copy) mail-x.x.x.jar from <confluence-install>/confluence/WEB-INF/lib to <confluence-install>/lib.
   Note: where x.x.x. represents the version numbers on the jar files in your installation.
   Do not leave a renamed backup of the jar files in `/confluence/WEB-INF/lib`. Even with a different file name, the files will still be loaded as long as it remains in the directory
3. Add the following to your server.xml file found in <confluence-install>/conf/ (add it just before the </Context> tag):

For Confluence 3.5.x

```
<Resource name="mail/GmailSMTPServer"
    auth="Container"
    type="javax.mail.Session"
    mail.smtp.host="smtp.gmail.com"
    mail.smtp.port="465"
    mail.smtp.auth="true"
    mail.smtp.user="yourEmailAddress@gmail.com"
    password="yourPassword"
    mail.smtp.starttls.enable="true"
    mail.transport.protocol="smtps"
    mail.smtp.socketFactory.class="javax.net.ssl.SSLSocketFactory"
/>
```

4. Restart Confluence.
5. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
6. Choose **Mail Servers**.
7. Choose either **Edit an existing configuration**, or **Add a new SMTP mail server**.
8. Edit the server settings as necessary, and set the **JNDI Location** as:

```
java:comp/env/mail/GmailSMTPServer
```

Note that the JNDI Location is case sensitive and must match the resource name specified in server.xml.

9. Save your changes and send a test email.

Configuring the Recommended Updates Email Notification

Confluence sends a regular email report to subscribers, containing the top content that is relevant to the person receiving the message, from spaces they have permission to view. This is called the 'Recommended Updates' notification.

If you have Confluence Administrator or System Administrator permissions, you can configure the default settings that determine how often the Recommended Updates notification is sent. When new users are added to Confluence, the default settings will be applied to their user profiles.

Confluence users can choose their personal settings, which will override the defaults. See Email Notifications.

Initial settings of the defaults

When you install Confluence, the initial values of the default settings are as follows:

- The default frequency is weekly.
- If your Confluence site has public signup enabled, the Recommended Updates notification is disabled by default. If public signup is not enabled, the notification is enabled by default.

You can change the above settings, specifying a different default value for the site.

Notes:

- The Recommended Updates notification is sent only to people who have a user profile in Confluence. If your Confluence site uses external user management, such as LDAP, then people will receive the report only after they have logged in for the first time. (The first login creates their user profile.)
- The daily email message is sent at 1 p.m. in the user's configured time zone.
- The weekly email message is sent at 1 p.m. on Thursdays in the user's configured time zone.

> **On this page:**
>
> - Initial settings of the defaults
> - Configuring the Recommended Updates notification
> - Disabling the Recommended Updates notification for the entire site
>
> **Related pages:**
>
> - Email Notifications
> - Confluence Administrator's Guide

Configuring the Recommended Updates notification

You can set the the default send option (send / do not send) and the default schedule (daily or weekly).

To configure the Recommended Updates email notification:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Click **Recommended Updates Email** in the left-hand panel.

Disabling the Recommended Updates notification for the entire site

You can also turn off the **recommended updates** notification for the entire site, by disabling the 'Confluence daily summary email' plugin. See Disabling and enabling add-ons.

The Mail Queue

Email messages waiting to be sent are queued in a mail queue and periodically flushed from Confluence once a minute. A Confluence administrator can also manually flush messages from the mail queue.

If there is an error sending messages, the failed email messages are sent to an error queue from which you can either try to resend them or delete them.

To view the mail queue:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **Mail Queue** in the left-hand panel. This will display the email messages currently in the queue.
3. Choose **Flush Mail Queue** to send all email messages immediately.
4. Choose **Error Queue** to view failed email messages. You can try to **Resend** the messages, which will flush the mails back to the mail queue, or you can **Delete** them from here.

> **Related pages:**
> - Configuring a Server for Outgoing Mail
> - Setting Up a Mail Session for the Confluence Distribution

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

Configuring Character Encoding

This page explains the encoding settings that are applicable in Confluence and how they relate to application behaviour.

To avoid problems with character encoding, make sure the encoding used across the different components of your system are the same. In general, **always set all character encodings to UTF-8**:

- Confluence character encoding – described b elow.

- Database – see Configuring Database Character Encoding.
- Application server – see Configuring URL Encoding on Tomcat Application Server

> **On this page:**
> - Configuring the Confluence character encoding
> - More details about character encoding
>   - Java character encoding
>   - Confluence character encoding
>   - Database encoding
>   - Filesystem encoding
> - Problems with character encodings
> - Notes
>
> **Related pages:**
> - Configuring Confluence
> - Confluence Administrator's Guide
>
> - Application Server Configuration
> - Database Configuration

Configuring the Confluence character encoding

By default, Confluence uses UTF-8 character encoding to deliver its pages.

Note: While it is possible to change the character encoding, we recommend that you leave this as it is unless you are certain of what you are doing.

In summary: Changing the Confluence character encoding will change your HTTP request and response encoding and your filesystem encoding as used by exports and Velocity templates.

To change the Confluence character encoding via the UI:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Enter the new character encoding of your choice in the text box next to **Encoding**.
5. Choose **Save**.

Note: At runtime, the character encoding is available in `Settings.defaultEncoding`.

More details about character encoding

There are three places where character encoding matters to Confluence:

1. **Database encoding** - usually the most important; it is where almost all user data is stored.
2. **Filesystem encoding** - important for attachment storage (pre-2.2), reading Velocity templates and writing exported files.
3. **HTTP request and response encoding** - important for form parsing, correct rendering by the browser and browser interpretation of encoded URLs.

Problems generally arise when Confluence thinks one of the above encoding is different to what it actually is. For example, Confluence might believe the database is using ISO-8859-1 encoding, when in fact it is UTF-8 encoded.

In certain cases (for example, Microsoft Windows), it might not be possible to use a fully Unicode filesystem (that is, a default Windows installation does not support Unicode filenames properly). If so, keep UTF-8 for the other two and be aware that your operating system might have limitations around international attachments (pre-2.2), backup and restore of international data, etc.

Java character encoding

Java *always* uses the multibyte UTF-16 character encoding for all `String` data*. This means that each of the encodings above defines how, at that particular point, characters are converted to and from Java's native UTF-16 format into some other format that the browser, filesystem or database might understand.

So when a request comes in to Confluence, we convert it from the request encoding to UTF-16. Then we store that data into the database, converting from UTF-16 to the database's encoding. Retrieving information from the database and sending it back to the browser is the same process in the opposite direction.

*A `char` represents single Unicode code point from the Base Multilingual Plane (BMP), encoded as UTF-16. Multiple `char`s are used as surrogate pairs for characters beyond U+FFFF.

Confluence character encoding

The Confluence character encoding is used in the following parts of the system:

- ConfluenceWebWorkConfiguration sets `webwork.i18n.encoding` to the this encoding, which WebWork uses in the response Content-Type header.
- AbstractEncodingFilter sets the HTTP request encoding to this encoding. This seems unnecessary, since the Content-Type header from the client should include the encoding used. This affects form submissions and file uploads.
- VelocityUtils reads in Velocity templates using this encoding when reading templates from disk.
- AbstractXmlExporter creates its output using this encoding.
- GeneralUtil uses this encoding when doing URLEncode and URLDecode. Different browsers have different support for character sets in URLs, so it's uncertain how much benefit this provides.

See Configuring Confluence Character Encoding (described above.)

Database encoding

The database encoding is the responsibility of your JDBC drivers. The drivers are responsible for reading and writing from the database in its native encoding and translating this data to and from Java Strings (which are UTF-16). For some drivers, such as MySQL, you must set Unicode encoding explicitly in the JDBC URL. For others, the driver is smart enough to determine the database encoding automatically.

Ideally, your database itself should be in a Unicode encoding (and we recommend doing this for the simplest configuration), but that is not necessary as long as:

- the database encoding supports all the characters you want to store in Confluence
- your JDBC drivers can properly convert from the database encoding to UTF-16 and vice-versa.

See Configuring Database Character Encoding.

Filesystem encoding

The filesystem encoding is mostly ignored by Confluence, except for the cases where the above

configuration setting above plays a part (exports, velocity). When attachments are uploaded, they are written as a stream of bytes directly to the filesystem. It is the same when they are downloaded: the bytes from the file InputStream are written directly to the HTTP response.

In some places in Confluence, we use the *default filesystem encoding* as determined by the JVM and stored in the `file.encoding` system property (it can be overridden by setting this property at startup). This encoding is used by the Java InputStreamReader and InputStreamWriter classes by default. This encoding should probably never be used; for consistent results across all filesystem access we should be using the encoding set in the General Configuration.

In certain cases we explicitly hard-code the encoding used to read or write data to the filesystem. Two important examples are:

- importing Mbox mailboxes which are known to be ISO-8859-1
- Confluence Bandana config files are always stored as UTF-8.

Some application servers, Tomcat for example, have an encoding setting that modifies Confluence URLs before they reach the application. This can prevent access to international pages and attachments (really anything with international characters in the URL). See configuring your Application Server URL encoding.

Problems with character encodings

If Confluence has the wrong idea about encoding for one of the above, it manifests itself in different ways:

1. Incorrect database encoding - user data is corrupted between saving and restoring from the database. This often happens after a delay, as we cache data as it is written to the database and only later retrieve the corrupted copy from the database.
2. Incorrect/non-Unicode filesystem encoding - international filenames break attachment download/upload/removal (pre-2.2); exports break with international content or attachments.
3. Incorrect HTTP encoding - incorrect encoding selected by browser, resulting in incorrect rendering of characters. Changing browser encoding causes page to render properly. Broken URLs when linking to pages or attachments with non-ASCII characters.

See Troubleshooting Character Encodings.

Notes

- Mac users please note that **MacRoman** encoding is compatible with UTF-8. You do not need to change your encoding settings if you are already using MacRoman.
- This is a good article by Joel Spolsky: The Absolute Minimum Every Software Developer Absolutely, Positively Must Know About Unicode and Character Sets (No Excuses!)

Troubleshooting Character Encodings

Often users may have problems with certain characters in a Confluence instance. Symptoms may include:

- Non-ASCII characters appearing as question marks (?)
- Page links with non-ASCII characters not working
- Single characters being displayed as two characters
- Garbled text appearing

In most cases, it is due to a mis-configuration in one of the components that Confluence uses.

Follow these steps to diagnose the problem.

1. Run the encoding test

Confluence includes an encoding test that can reveal problems with your configuration.

To perform the test, access the Encoding Test page via the `<confluence base-url>/admin/encodingtest.action` page on your Confluence instance. You will be required to copy and paste a line of text and submit a form. The test will take the text and pass it through Confluence, the application server and the database, and return the results.

You should also test pasting some sample text (Japanese for example) if you are experiencing problems with a

specific language.

Example:

```
http://confluence.atlassian.com/admin/encodingtest.action
```

or

```
http://<host address>:<port>/admin/encodingtest.action
```

> If the text displayed in the encoding test is different to what was entered, then there are problems with your character encoding settings.

A successful test looks like the following:

Screenshot: Successful encoding test

Character Encoding Test Results

The encoding test has now been run. Below, you can compare the raw text delivered from Confluence against the text returned by your browser in web forms, and the text as it appears after a round-trip through the database. All the test results should appear identical.

Iñtërnâtiônàlizætiøn This image is how the sample text below *should* appear. If it does not, please file a support request at http://support.atlassian.com, including a screenshot of this page, and all of your System Information.

Test 1: Raw text

This is the test string generated in Confluence

Iñtërnâtiônàlizætiøn

Test 2: Form submission

This is the test string pasted by you into the web form and submitted back to Confluence

Iñtërnâtiônàlizætiøn

Test 3: Database round-trip (select as lower-case)

This is the string from Test 2 after being stored in the database and then retrieved as lower-case

iñtërnâtiônàlizætiøn

Expected result (converting Java string to lowercase)

iñtërnâtiônàlizætiøn

Test 4: Database round-trip (select as upper-case)

This is the string from Test 2 after being stored in the database and then retrieved as upper-case

IÑTËRNÂTIÔNÀLIZÆTIØN

Expected result (converting Java string to uppercase)

IÑTËRNÂTIÔNÀLIZÆTIØN

Test 5: International file name support

Try to write a file to the confluence home directory with the test string as the file name

✅ **File was written successfully**

Test 6: Detect international file name mangling

Detect whether the file system is mangling the file name when it is saved

✅ **The file name has been preserved**

MySQL 3.x
MySQL 3.x is known to have some problems with the upper- and lower-casing of some characters, and may fail the last two tests. For more information, see MySQL 3.x Character Encoding Problems.

2. Ensure the same encoding is used across all components

As mentioned in the Configuring Encoding document, the same character encoding should be used across the database, application server and web application (Confluence).

- To change the character encoding used in **Confluence**, see Configuring Character Encoding.
- To change the character encoding used in the **application server**, please ensure you set the Application Server URL encoding and view your application server's documentation on any other settings required to enable your encoding.
- To change the character encoding used in the **database**, see Configuring Database Character Encoding.

3. Requesting support

If there are still problems with character encoding after following the above steps, create a support request, and our support staff will aid in solving your problem.

Entering in the following details will help us to identify your problem:

- Attach screenshots of the problem
- Attach the results of the encoding test (above)
- Select which application server (and version) you are using
- Select which database (and version) you are using
- Copy the contents of the System Information page into the 'Description' field

"€" Euro character not displaying properly

The € (euro) symbol is a three byte character, with byte values in file (UTF-8) of 0xE2, 0x82, 0xAC.

Sometimes, if the character encoding is not set consistently among all participating entities of the system, Confluence, server and the database, one may experience strange behaviour.

> ...
> I write a page with a Euro sign in it (€). All is well, the Euro sign shows up in the wiki markup text-box, and the preview, and the display of the saved page.
> One day later, the Euro sign has changed into a question mark upside down!
> ...
> What is going on? Why does the Euro sign mysteriously change? How do I prevent it?

Interestingly enough the character encoding test passes with no problems, demonstrating that Confluence and the connected Database both recognise the € symbol.

There are two potential reasons for this behaviour:
Database and Confluence is using utf-8 encoding. The connection is not.

When data transferred to it via the connection which does not use utf-8 encoding gets encoded incorrectly. Hence, updating the connection encoding may resolve this problem from now on, yet it probably would not affect already existing data.
Database is not using utf-8. Confluence and your connection are.

If your Database encoding is not set to UTF-8, yet is using some other encoding such as *latin1*, it could be one of the potential reasons why you lose the "€" characters at some stage. It could be occurring due to **caching**. When Confluence saves data to the database, it may also keep a local cached copy. If the database encoding is set incorrectly, the Euro character may not be correctly recorded in the database, but Confluence will continue to use its cached copy of that data (which is encoded correctly). The encoding error will only be noticed when the cache expires, and the incorrectly encoded data is fetched from the database.

> For instance the *latin1* encoding would store and display all 2-byte UTF8 characters correctly except for the euro character which is replaced by '?' before being stored. As Confluence's encoding was set to UTF-8, the 2-byte UTF-8 characters were stored in *latin1* database assuming that they were two *la tin1* 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)

Solution

Upgrade to a newer version of MySQL. (4.1 is confirmed to work.)

Other Settings

- Configuring a WebDAV client for Confluence
- Configuring HTTP Timeout Settings
- Configuring Number Formats
- Configuring Shortcut Links
- Configuring Time and Date Formats
- Enabling the Remote API
- Enabling Threaded Comments
- Installing a Language Pack
- Installing Patched Class Files

Configuring a WebDAV client for Confluence

WebDAV allows users to access Confluence content via a WebDAV client, such as 'My Network Places' in Microsoft Windows. Provided that the user has permission, they will be able to read and write to spaces, pages and attachments in Confluence. Users will be asked to log in and the standard Confluence content access permissions will apply to the equivalent content available through the WebDAV client.

> Mapping a Confluence WebDAV network drive requires a set of specific criteria to be met. For specific information, please see Windows Network Drive Requirements.

Introduction to Confluence's WebDAV Client Integration

By default, all WebDAV clients have permission to write to Confluence. Write permissions include the ability for a WebDAV client to create, edit, move or delete content associated with spaces, pages and attachments in a Confluence installation.

On this page:

- Introduction to Confluence's WebDAV Client Integration
- Using a WebDAV Client to Work with Pages
- Restricting WebDAV Client Write Access to Confluence
- Disabling Strict Path Checking
- Virtual Files and Folders
- Known Issues

Related pages:

- Global Permissions Overview
- Disabling and enabling add-ons
- Disabling or Enabling Confluence Add-ons
- Attachment Storage Configuration

On the '**WebDAV Configuration**' screen in the Confluence Administration Console, you can:

- Deny a WebDAV client write permissions to a Confluence installation using a regular expression (regex).
- Disable or enable strict path checking.
- Enable or disable access to specific virtual files/folders.

Note:

- The 'WebDav Configuration' page is only available if the WebDAV plugin has been enabled. Note that this plugin is bundled with Confluence, and can be enabled or disabled by the System Administrator.
- The settings on the 'WebDav Configuration' page do **not** apply to external attachment storage configuration.

Using a WebDAV Client to Work with Pages

The following sections tell you how to set up a WebDAV client natively for a range of different operating systems. WebDAV clients typically appear as drives in your operating system's file browser application, such as Windows Explorer in Microsoft Windows, or Konqueror in Linux.

Accessing Confluence in Finder on Mac OSX

To use Finder to view and manage Confluence content:

1. In Finder choose **Go** > **Connect to Server**.
2. Enter your Confluence URL in this format:

   ```
   http://<confluenceURL>/plugins/servlet/confluence/default
   ```

   For example if your Confluence URL is http://ourconfluence.atlassian.net/wiki you would enter:

   ```
   http://ourconfluence.atlassian.net/wiki/plugins/servlet/confluence/default
   ```

   You must use HTTP, you'll be able to connect with HTTPS but you wont see any content.
3. Enter your **Confluence** username and password and click **Connect**.
   Note: use your username (jsmith), not your email address, unless your email address is your username.
4. Confluence will appear as a shared drive in Finder.

You can use the same URL to connect using a third party WebDav client, like CyberDuck.

Accessing Confluence in Explorer in Microsoft Windows

This section covers the two methods for configuring a WebDAV client natively in Microsoft Windows:

- As a network drive
- As a web folder

If possible, use the network drive method as this will enable more comprehensive WebDAV client interaction with Confluence than that provided by a web folder. However, your Confluence instance must meet several environmental constraints if you use this method. If you cannot configure your instance to meet these requirements, then use the web folder method or third-party WebDAV client software.

If you run into any problems with the procedures in this section, please refer to the WebDAV Troubleshooting page.

Windows Network Drive

To map a Confluence WebDAV client network drive, your Confluence instance must be configured so that *all* of the following criteria is met:

- Uses HTTP (not HTTPS)

- Listens on port 80 (not 8090, the default for Confluence distributions)
- Has no context root
- There is an issue (WBDV-208) that can prevent Network Drives from being mapped. Please use the Network Folders steps below as a workaround.

The reason for these restrictions results from limitations in Microsoft's Mini-Redirector component. For more information, please refer to Microsoft's server discovery issue.

To map a Confluence WebDAV client network drive in Microsoft Windows:

1. In Windows XP, go to **My Computer -> Tools menu -> Map Network Drive**.
   In Windows Vista, go to **Computer -> Map Network Drive**.
   The 'Map Network Drive' dialog box opens.
2. Specify the following input to map the WebDAV client as a network drive:
   - **Drive:** `<Any drive letter>` (for example, `Z:`)
   - **Folder:** `\\<hostname>\webdav` (for example, `\\localhost\webdav`)
3. Click '**Finish**'.
   When prompted for login credentials, specify your Confluence username and password.

Windows Web Folder

To map a Confluence WebDAV client web folder in Windows XP:

1. Go to **My Network Places** and choose '**Add a network place**'. The 'Add Network Place Wizard' opens.
2. Click '**Next**', ensure that '**Choose another network location**' is selected and then click '**Next**' again.
3. In the 'Internet or network address' field, enter the URL for the Confluence WebDAV location (for example, `http://<confluence server url>/confluence/plugins/servlet/confluence/default` or `http://<confluence server url>/plugins/servlet/confluence/default`) and then click '**Next**'.
   When prompted for login credentials, specify your Confluence username and password.
4. Provide a meaningful name for your web folder and proceed with the remainder of the wizard.
5. Click '**Finish**'.

Screenshot: A Confluence WebDAV Client Web Folder in Windows XP

To map a Confluence WebDAV client web folder in Windows Vista:
This procedure is very similar to the one for Windows XP. However, the following procedure includes the slight interface differences that are specific to Windows Vista.

1. Open the 'Map Network Drive' dialog box (refer to first step of the procedure above for mapping a

network drive) and choose '**Connect to a Web site that you can use to store your documents and pictures**'. The 'Add Network Location' wizard opens.

2. Click '**Next**', ensure that '**Choose a custom network location**' is selected and then click '**Next**' again.
3. In the 'Internet or network address' field, enter the URL for the Confluence WebDAV location (for example, `http://<confluence server url>/confluence/plugins/servlet/confluence/default` or `http://<confluence server url>/plugins/servlet/confluence/default`) and then click '**Next**'.
   When prompted for login credentials, specify your Confluence username and password.
4. Provide a meaningful name for your network location/web folder and proceed with the remainder of the wizard.
5. Click '**Finish**'.

Setting up a WebDAV client in Linux or Solaris

There are many tools and mechanisms available for configuring WebDAV clients in these operating systems. Therefore, we have chosen to demonstrate this using the file manager Konqueror, which is part of the Linux K Desktop Environment.

To set up a Confluence WebDAV client in Konqueror:

1. Open Konqueror.
2. In the 'Location' field, enter the URL for the Confluence WebDAV location using the 'protocol' `webdavs` (for example, `webdavs://<confluence server url>/confluence/plugins/servlet/confluence/default` or `webdavs://<confluence server url>/plugins/servlet/confluence/default`) and press `Enter`.
   If prompted for login credentials, specify your Confluence username and password.
   You should be able to click to load many, but not all files. In practice, you would normally save a modified file locally, then drag it to the Konqueror window to upload it to Confluence.

Restricting WebDAV Client Write Access to Confluence

In earlier versions of the WebDAV plugin, separate options for restricting a WebDAV client's write permissions (that is, create/move, edit and delete actions), were available. However, in the current version of this plugin, they have been simplified and combined into a general write permission restriction that covers all of these actions.

WebDAV clients are now denied write permission to your Confluence installation by setting a regex that matches specific content within the WebDAV client's user agent header. Upon setting a regex, it will be added to a list of restricted WebDAV clients. Any WebDAV clients whose user agent header matches a regex in this list will be denied write permission to your Confluence installation.

Example: A PROPFIND method header generated by a Microsoft Web Folder WebDAV client, showing the user agent header field:

```
PROPFIND /plugins/servlet/confluence/default HTTP/1.1
Content-Language: en-us
Accept-Language: en-us
Content-Type: text/xml
Translate: f
Depth: 1
Content-Length: 489
User-Agent: Microsoft Data Access Internet Publishing Provider DAV
Host: 127.0.0.1:8082
Connection: Keep-Alive
```

Note: Unlike earlier versions of the WebDAV plugin which could only restrict write permissions for **all** WebDAV clients, the current version of this plugin allows you to restrict write permissions to specific WebDAV clients selectively.

To restrict a WebDAV client's write access permissions to your Confluence installation:

1. Choose the **cog icon**

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

2. Choose '**WebDav Configuration**' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Enter a regex that matches a specific component of the user agent header sent by the WebDAV client you want to restrict.
4. Click the '**Add new regex**' button. The regex is added to the list of restricted WebDAV clients.
   You can repeat steps 3 and 4 to add a regex for each additional WebDAV client you want to restrict.
5. Click the '**Save**' button to save the configuration changes.

To restore one or more restricted WebDAV client's write access permissions to your Confluence installation:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Select the regex(es) from the list that match(es) the user agent header sent by the restricted WebDAV client(s) you want to restore.
4. Click the '**Remove selected regexes**' button. The regexes you had selected are removed from the list of restricted WebDAV clients.
5. Click the '**Save**' button to save the configuration changes.

Screenshot: WebDAV configuration

Disabling Strict Path Checking

If you observe any idiosyncrasies with your WebDAV client, such as a folder that does exist on your Confluence site but is missing from the client, you can disable the WebDAV plugin's strict path checking option, which may minimise these problems.

To disable the WebDAV plugin's strict path checking option:

1. Choose the **cog icon**

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

2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Clear the 'Disable strict path check' check box.
   ℹ️ You can re-enable this option at a later point in time by simply selecting this check box.
4. Click the '**Save**' button to save this configuration change.
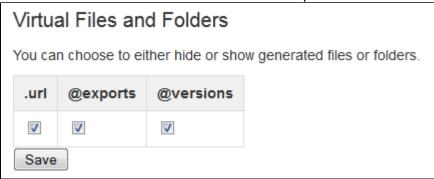
Virtual Files and Folders

In the unlikely event that you observe any problems with the WebDAV client's performance or stability, you can enable access to automatically generated (that is, virtual) files and folders.

Note:

By default, these options are hidden on the 'WebDAV Configuration' page. To make them visible, you must append the parameter `?hiddenOptionsEnabled=true` to the end of your URL and reload the page. For example:

```
<Confluence base
URL>/admin/plugins/webdav/config.action?hiddenOptionsEnabled=true
```

Screenshot: The Hidden Virtual Files and Folders Option

To enable or disable access to virtual files and folders:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Amend your URL as described in the note above and reload the 'WebDav Configuration' page.
4. Select or clear the check box options in the 'Virtual Files and Folders' section as required.
5. Click the '**Save**' button to save the configuration changes.

Known Issues

Please refer to the WebDAV plugin documentation for a description of the known issues and suggested workarounds.

Configuring HTTP Timeout Settings

When macros such as the RSS Macro make HTTP requests to servers which are down, a long timeout value is used. You can set this timeout value through a system parameter to avoid this.

To configure the HTTP Timeout Settings:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Select '**General Configuration**' under the 'Configuration' heading in the left-hand panel.
3. Find the '**Connection Timeouts**' section in the lower portion of the screen.

4. Click '**Edit**' to adjust the settings:
    - **Adjust External connections enabled**: This setting allows system administrators to disable external connections so macros like the RSS Macro wont be allowed to make connections to an external server. It provides protection against external servers providing insecure HTML, timing out or causing performance problems. The default setting is 'true'.
    - **Connection Timeout (milliseconds)**: Sets the maximum time for a connection to be established. A value of zero means the timeout is not used. The default setting is ten seconds (10000).
    - **Socket Timeout (milliseconds)**: Sets the default socket timeout (SO_TIMEOUT) in milliseconds, which is the maximum time Confluence will wait for data. A timeout value of zero is interpreted as an infinite timeout. The default setting is ten seconds (10000).

Configuring Number Formats

There are two number format settings in Confluence:

- Long number format. For example: ##############
- Decimal number format. For example: ##############.##########

Confluence uses the guidelines in this Java document from Oracle: Class NumberFormat.

To change the number formats in Confluence:

1. Choose

    ⚙️ ▾

    > **General Configuration**
2. Choose **Edit**
3. Update the **Long Number Format** and **Decimal Number Format** to suit your requirements
4. Choose **Save**

Configuring Shortcut Links

Shortcut links provide a quick way of linking to resources that are frequently referenced from Confluence. When you create a shortcut link, you assign a key to an URL so that, when editing, a user can type just the key instead of the complete URL.

Example: Creating a shortcut to Google

Most Google searches look like this: `http://www.google.com/search?q=`. If you create a shortcut for this search with the key 'google', every time a user needs to use `http://www.google.com/search?q=`**searchterms**, they can just type [**search terms**@google] instead.

On this page:
- Creating shortcut links
- Using shortcut links
- Deleting shortcut links

Related pages:

- Links
- Confluence Administrator's Guide

Here is a screenshot showing the shortcuts currently defined on http://confluence.atlassian.com:

Key	Expanded Value	Default Alias	Operations
cache	http://www.google.com/search?q=cache:		Remove
imdb	http://us.imdb.com/Title?		Remove
jira	http://jira.atlassian.com/secure/QuickSearch.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=		Remove

Shortcut links are added and maintained by Confluence administrators from the **Administration Console**.

Creating shortcut links

To create a shortcut link:

1. Choose the **cog icon**

    ⚙️ ▾

, then choose **General Configuration** under Confluence Administration
2. Choose **Shortcut Links** in the left-hand panel.
3. Enter a **Key** for your shortcut. This is the shortcut name a user will use to reference the URL.
4. Enter the **Expanded Value**. This is the URL for the link. You can use '%s' in the URL to specify where the user's input is inserted. If there is no '%s' in the URL, the user's input will be put at the end.
5. Enter a **Default Alias**. This is the text of the link which will be displayed on the page where the shortcut is used, with the user's text being substituted for '%s'.
6. Choose **Submit**.

Using shortcut links

Enter a shortcut link on the **Advanced** tab of the Insert Link dialog. See Links for details.

Specify in the link what should be appended to the end of the shortcut URL, followed by an at-sign (@) and the key of the shortcut. Shortcut names are case-insensitive. So, for example, using the keys shown in the above screenshot:

To link to...	Type this	Resulting URL
a JIRA issue	CONF-1000@JIRA	http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=CONF-1000
a Google search	Atlassian Confluence@Google	http://www.google.com/search?q=Atlassian+Confluence

Deleting shortcut links

Shortcut links are listed on the **Shortcut Links** tab of the Administration Console. Click **Remove** to delete the shortcut.

Configuring Time and Date Formats

You can localise the formats that Confluence uses to display dates and times within the web interface. The settings use the syntax of Java's SimpleDateFormat class, as described in this document: Java SimpleDateFormat.

There are three time and date format settings:

- Time format: Used when displaying only the time of day. For example, when a blog post is published. Example of configuration: `h:mm a`
- Date time format: Used when displaying both the date and the time of day. For example, in historical versions of pages. Example of configuration: `MMM dd, yyyy HH:mm`
- Date format: Used when displaying only the date. For example, the creation and most recent modification dates of pages. Example of configuration: `MMM dd, yyyy`

To change the time and date formats:

1. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
2. Choose **General Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Enter the values for **Time Format**, **Date Time Format** and **Date Format**, to suit your requirements.
5. Choose **Save**.

Related pages:

- Choosing a Default Language
- Installing a Language Pack

- Confluence Administrator's Guide

Enabling the Remote API

Confluence provides XML-RPC and SOAP remote APIs (application programming interfaces). You need to enable the APIs from the **Administration Console** before you can access Confluence remotely.

You need System Administrator permissions in order to perform this function.

To enable the remote API:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Click **Further Configuration** in the left-hand panel.
3. Click **Edit**.
4. Click the check box next to **Remote API (XML-RPC & SOAP)**.
5. Click **Save**.

Related pages:

- Confluence Remote API Reference

Enabling Threaded Comments

Comments on pages or blog posts are displayed in one of two views:

- **Threaded**: Shows the comments in a hierarchy of responses. Each reply to a comment is indented to indicate the relationships between the comments.
- **Flat**: Displays all the comments in one single list and does not indicate the relationships between comments.

By default, comments are displayed in **threaded** mode. A Confluence Administrator (see Global Permissions Overview) can enable or disable the threaded view for the entire Confluence site.

To enable or disable the threaded view:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Select **Further Configuration** in the left-hand panel
3. Choose **Edit**
4. Select or unselect the **Threaded Comments** checkbox to enable or disable threaded mode
5. Choose **Save**

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. You can make additional languages available for selection by installing language packs. Please note, you must be a Confluence administrator to install a language pack.

Language packs are plugins. The process of installing a language pack is the same as installing a new plugin.

Related pages:

- Choosing a Default Language
- Configuring Indexing Language
- Installing add-ons

Installing a Language Pack using the Universal Plugin Manager

To install a language pack using the Universal Plugin Manager:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **Find New Add-ons** in the left-hand panel.

3. Find the language pack on the Atlassian Marketplace.
4. Choose **Install** to install the language pack.

Installing a Language Pack Manually

To install a language pack manually, you will need to upload the language pack plugin as described below. The language pack plugin will be enabled by default once you have installed it.

Plugins are distributed as JAR or OBR (OSGi Bundle Repository) files. To install a plugin:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Manage Add-ons**.
3. Choose **Upload Plugin**.
4. Choose **Browse** to find the plugin file you wish to install from your hard drive and select it, or enter a network location by URL.
5. Choose **Upload**.
   The plugin will be uploaded to Confluence and will be automatically installed.
6. Check the list of user-installed plugins to ensure that the add-on is available.
7. Enable the plugin if necessary. (Some plugins will be enabled by default when they are installed. Others will have to be manually enabled from the 'Manage Add-ons' page.)

Finding more Language Packs

- You can download official language packs from the Atlassian Marketplace. You can also download language packs developed by the Confluence user community from the Language Pack Translations page.

Showing User Interface Key Names for Translation

This feature is useful if you are working on creating translations of the Confluence user interface. After opening the Confluence dashboard, you can add this text to the end of your Confluence URL:

```
?i18ntranslate=on
```

Then press Enter.

This will cause each element of the user interface to display its special **key name**. This makes it easier to find the context for each key within the user interface. You can then search for the key on http://translations.atlassian.com where you can enter an appropriate translation for your custom language pack.

The key names are displayed with a 'lightning bolt' graphic. For example:

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

```
?i18ntranslate=off
```

Installing Patched Class Files

Atlassian support or the Atlassian bug-fixing team may occasionally provide patches for critical issues that have been resolved but have not yet made it into a release. Those patches will be class files which are attached to the relevant issue in our JIRA bug-tracking system.

Installation Instructions for the Confluence Distribution

Follow these steps to install a patched class file:

1. Shut down your confluence instance.
2. Copy the supplied class files to `<installation-directory>/confluence/WEB-INF/classes/<subdirectories>`, where:
   - `<installation-directory>` must be replaced with your Confluence Installation directory. (If you need more information, read about the Confluence Installation Directory.)
   - `<subdirectories>` must be replaced by the value specified in the relevant JIRA issue. This value will be different for different issues. In some cases, the subdirectories will not exist and you will need to create them before copying the class files. Some issues will contain the patch in the form of a ZIP file which will contain the desired directory structure.
3. Restart your Confluence instance for the changes to become effective.

ⓘ Class files in the `/WEB-INF/classes` directory of a web application will be loaded before classes located in JAR files in the `/WEB-INF/lib` directory. Therefore, classes in the first directory will effectively replace classes of the same name and package which would otherwise be loaded from the JAR files.

Reverting the patch

To revert the patch, simply remove the class files from the `<installation-directory>/confluence/WEB-INF/classes/` folder (taking care to only remove those that apply to the patch you wish to revert), then restart the instance.

Once the issue that the patch relates to is resolved, you should upgrade to the version of Confluence that contains the fix, and revert the patch. Patches are often naive and untested and may not solve the problem in the most efficient way. As such, an official fix should be preferred in all cases.

Configuring System Properties

This page describes how to set Java properties and options on startup for Confluence.

See How to fix out of memory errors by increasing available memory for specific instructions for OutOfMemory Errors.

On this page:
- Linux
- Windows (starting from .bat file)
- Windows Service

Related pages:
- Recognised System Properties
- How to fix out of memory errors by increasing available memory

Linux

To configure System Properties in Linux installations:

1. From `<confluence-install>/bin` (Stand-alone) or `<Tomcat-home>/bin` (EAR-WAR installation), open **setenv.sh**.
2. Find the section **CATALINA_OPTS=**
   (this is JAVA_OPTS= in Confluence 5.5 and earlier)
3. Refer to the list of parameters below.

ⓘ Add all parameters in a space-separated list, inside the quotations.

Windows (starting from .bat file)

To Configure System Properties in Windows Installations When Starting from the .bat File:

1. From `<confluence-install>/bin` (Stand-alone) or `<Tomcat-home>/bin` (EAR-WAR installation), open **setenv.bat**.
2. Find the section **set CATALINA_OPTS=%CATALINA_OPTS%**
   (this is set JAVA_OPTS=%JAVA_OPTS% in Confluence 5.5 and earlier)
3. Refer to the list of parameters below.

ⓘ Add all parameters in a space-separated list. Make sure to keep the string %CATALINA_OPTS% in
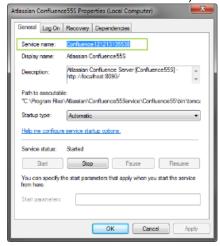
place.

Windows Service

There are two ways to configure system properties when you Start Confluence Automatically on Windows as a Service, either via command line or in the Windows Registry

Setting Properties for Windows Services via Command Line

To set properties for Windows Services via a command line:

1. Identify the name of the service that Confluence is installed as in Windows (Go to **Control Panel** > **Ad ministrative Tools** > **Services**):
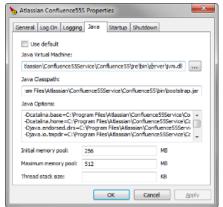

   In the above example, the service name is *Confluence121213135538*.
2. Open the command window (Choose **Start** > **cmd.exe**)
3. `cd` to the `bin` directory of your Confluence instance (or the `bin` directory of your Tomcat installation if you are running Confluence EAR/WAR).
4. Run the following command:

```
tomcat7w //ES//<SERVICENAME>
```

   In the above example, it would be `tomcat7w //ES//Confluence121213135538`
   The Tomcat version number may be different if you are using an earlier version of Confluence.
5. Click on the `Java` tab to see the list of current start-up options:

6. Append any new option on its own new line by adding to the end of the existing Java Options. Refer to the list of parameters below.
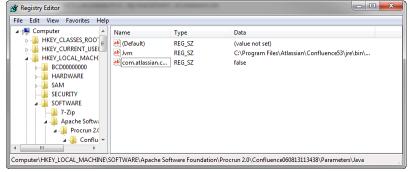
Setting Properties for Windows Services via the Windows Registry

In some versions of Windows, there is no option to add Java variables to the service. In these cases, you must add the properties by viewing the option list in the registry.

1. Go to the Registry Editor (**Start** > **regedit.exe**).
2. Find the Services entry:
   32-bit: `HKEY_LOCAL_MACHINE >> SOFTWARE >> Apache Software Foundation >> Procrun 2.0 >> Confluence`
   64-bit: `HKEY_LOCAL_MACHINE >> SOFTWARE >> Wow6432Node >> Apache Software Foundation >> Procrun 2.0 >> Confluence service name`

3. To change existing properties double-click the appropriate value.
4. To change additional properties, double-click options.
5. Refer to the list of parameters below. Enter each on a separate line.

Verifying Your Settings

To see what Confluence is using, check Viewing System Properties.

Recognised System Properties

Since	Default Value	Module...	Effect
atlassian.forceSchemaUpdate			
1.0	`false`	atlassian-config	By default, Confluence will only run its database schema update when it detects that it has been upgraded. This flag will force Confluence to perform the schema update on system startup.
confluence.home			
1.0	Any filesystem path	Confluence and atlassian-config	If this system property is set, Confluence will ignore the contents of the `confluence-init.properties` file, and use this property as the setting for the Confluence Home directory.
confluence.dev.mode			
1.0	`false`	Confluence	Enables additional debugging options that may be of use to Confluence developers (additionally it changes spring bean creation to use lazy initialization by default to decrease startup time). Do not enable this flag on a production system.
confluence.disable.mailpolling			

2.4	`false`	Confluence	If set to "true", will prevent Confluence from retrieving mail for archiving within spaces. Manually triggering "check for new mail" via the web UI will still work. This property has no effect on outgoing mail
confluence.i18n.reloadbundles			
1.0	`true`	Confluence	Setting this property will cause Confluence to reload its i18n resource bundles every time an internationalised string is looked up. This can be useful when testing translations, but will make Confluence run *insanely slowly*.
confluence.ignore.debug.logging			
1.0	`true`	Confluence	Confluence will normally log a severe error message if it detects that DEBUG level logging is enabled (as DEBUG logging generally causes a significant degradation in system performance). Setting this property will suppress the error message.
confluence.jmx.disabled			
3.0	`false`	Confluence	If set to "true", will disable Confluence's JMX monitoring. This has the same effect as setting the "enabled" property to false in `WEB-INF/classes/jmxContext.xml`
confluence.optimize.index.modulo			
2.2	`20`	Confluence	Number of index queue flushes before the index is optimised.
confluence.plugins.bundled.disable			
2.9	`false`	Confluence	Starts confluence without bundled plugins. May be useful in a development environment to make Confluence start quicker, but since bundled plugins are necessary for some of Confluence's core functionality, this property should not be set on a production system.
atlassian.indexing.contentbody.maxsize			
3.0		atlassian	Confluence stores the body of attachments in the index so the content of attachments are searchable. Set this property to set the maximum indexable size (in bytes). Limiting the size of attachment indexing decreases the size of the index when large attachments are present and prevents excerpts of large attachments being displayed in search results. It does not limit the size of attachments that can be uploaded.

atlassian.mail.fetchdisabled			
3.5	`false`	Confluence	Disables mail fetching services for IMAP and POP
atlassian.mail.senddisabled			
3.5	`false`	Confluence and atlassian-mail	Disables sending of mail
atlassian.disable.caches			
2.4	`true`	atlassian-plugins, atlassian-cache-servlet	Setting this property will disable conditional get and expires: headers on some web resources. This will significantly slow down the user experience, but is useful in devlopment if you are frequently changing static resources and don't want to continually flush your browser cache.
confluence.html.encode.automatic			
2.9		Confluence	Setting this property forces the antixss encoding on or off, overriding the behaviour dictated by settings. The default behaviour differs between Confluence versions.
org.osgi.framework.bootdelegation			
2.10	empty	atlassian-plugins	Comma-separated list of package names to provide from application for OSGi plugins. Typically required when profiling Confluence. For example: "com.jprofiler.**,com.yourkit.**".
confluence.diff.pool.size			
3.1	`20`	Confluence	Maximum number of concurrent diffs. When that number is exceeded, additional attempts by RSS feeds to create diffs are ignored and logged. (The RSS requests succeed, they are just missing diffs).
confluence.diff.timeout			
3.1	`1000`	Confluence	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.
confluence.html.diff.timeout			
4.0	`10000`	Confluence	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.
atlassian.user.experimentalMapping			

2.10	`false`	Confluence	Setting this property changes the relationship between local users and local groups to reduce performance degradation when adding a local user to a local group with a large number of users. Please note, setting this property can slow down other user management functions. We recommend that you set it only if you are experiencing performance problems when adding local users to large local groups. Please refer to CONF-12319, fixed in Confluence 3.1.1.
confluence.import.use-experimental-importer			
3.2	`false`	Confluence	Setting this property changes Confluence to use the Experimental XML Importer. It is designed to be a more stable implementation but, at the time of the release of 3.2, the importer is largely untested and thus not supported.
atlassian.webresource.disable.minification			
3.3	`false`	atlassian-plugins	Disables automatic minification of JavaScript and CSS resources served by Confluence.
index.queue.thread.count			
3.3	*See "Effect"*	Confluence	Sets the number of threads to be used for the reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. There is no default value, i.e.

Under the "Effect" column for *index.queue.thread.count*:

- 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.queue.batch.size			
3.3	1500	Confluence	Size of batches used by the indexer. Reducing this value will reduce the load that the indexer puts on the system, but indexing takes longer. Increasing this value will cause indexing to be completed faster, but puts a higher load on the system. Normally this setting does not need tuning.

password.confirmation.disabled			
3.4	`false`	Confluence	This property disables the password confirmation functionality that Confluence uses as an additional security measure. With this property set, Confluence will *not* require password confirmation for the following actions: administrative actions, change of email address and Captcha for failed logins. Disabling password confirmations is useful if you are using a custom authenticator.

confluence.browser.language.enabled			
3.5	`true`	Confluence	Setting this property to "false" disables the detection of browser language headers, effectively restoring Confluence behaviour to that of earlier releases. Setting this property to "true" enables the detection of the language headers sent by the browser. Confluenc will change the UI language based on the browser headers. See documentation on how users can choose a language preference.

upm.pac.disable			
Universal Plugin Manager 1.5	`false`	Universal Plugin Manager (UPM)	When this property is set to true, then UPM will not try to access the The Atlassian Marketplace. This is useful for application servers that do not have access to the Internet. See the UPM documentation.

confluence.reindex.documents.to.pop			
3.5.9	`20`	Confluence	Indicates how many objects each indexing thread should process at a time during a full re-index. Please note that this number does not include attachments

confluence.reindex.attachments.to.pop			
3.5.9	`10`	Confluence	Indicates how many attachments each indexing thread should process at a time during a full re-index.

confluence.upgrade.active.directory			
3.5.11	`false`	Confluence	Forces Confluence to treat any LDAP directories it migrates as Active Directory, rather than relying on looking for sAMAccountName in the username attribute. This is necessary if you are upgrading from before Confluence 3.5, and need to use an attribute other than sAMAccountName to identify your users and are seeing `LDAP: error code 4 - Sizelimit Exceeded` exceptions in your logs. For more details, see Unable to Log In with Confluence 3.5 or Later Due to 'LDAP error code 4 - Sizelimit Exceeded'

confluence.context.batching.disable			
4.0	`false`	Confluence	Disables batching for web resources in contexts (e.g. editor, main, admin). Useful for diagnosing the source of javascript or CSS errors.

com.atlassian.logout.disable.session.invalidation			
4.0	`false`	Confluence	Disables the session invalidation on log out. As of 4.0 the default behaviour is to invalidate the JSession assigned to a client when they log out. If this is set to true the session is kept active (but the user logged out). This may be valuable when using external authentication systems, but should generally not be needed.

officeconnector.spreadsheet.xlsxmaxsize			
4.0.5	2097152	Office Connector	Indicates the maximum size in bytes of an Excel file that can be viewed using the `viewxls` macro. If empty, the maximum size defaults to 2Mb. See CONF-21043 for more details.

com.atlassian.confluence.extra.calendar3.display.events.calendar.maxpercalendar			
	200	Team Calendars	Specifies the maximum number of events per calendar. This property is effective only if the Team Calendars plugin is installed on your Confluence site.

com.atlassian.confluence.allow.downgrade			
4.3.2, 5.0-OD-10	`false`	Confluence	Allows Confluence to start up against the home directory of a newer version of Confluence. Note that running Confluence like that is unsupported. You should only turn this on if you know what you are doing. See After Downgrading, Confluence Will No Longer Run for details.

reindex.thread.count			
5.2		Confluence	Sets the number of threads to be used for a one-off reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. There is no default value. This system property does not affect the incremental indexing that Confluence does.

reindex.attachments.thread.count			
5.2	4	Confluence	Sets the number of threads to be used for reindexing attachments specifically, and allows you to reduce the concurrency for these more memory intensive index items.

confluence.mbox.directory

5.4.1		Confluence	Setting this property defines the directory on your Confluence Server where mailboxes can be imported from (for use with the Confluence Mail Archiving add-on). Mailboxes are not able to be imported from any other location. We recommend administrators create a directory in the Confluence Home directory specifically for this purpose. Mail cannot be imported from the server until this system property is set.

confluence.search.max.results

5.5	1000	Confluence	Setting this property changes the maximum number of items Confluence Search will return.

confluence.upgrade.recovery.file.enabled

5.5	`true`	Confluence	By default, Confluence creates an upgrade recovery file before and after an upgrade. The operation can take long time on large databases and can be safely turned off if there is a process to back up database and verify the backup before performing an upgrade. Setting this property to false will disable upgrade recovery file creation.

confluence.junit.report.directory

5.5		Confluence	Setting this property defines the directory on your Confluence Server where JUnit Reports can be imported from (for use in the JUnit Report Macro). No other locations are permitted. We recommend administrators create a directory in the Confluence Home directory specifically for this purpose. JUnit Test result files cannot be imported from the server until this system property is set.

officeconnector.textextract.word.docxmaxsize

5.5.3	16777216	Confluence	This property defines the maximum amount of memory that can be allocated to extracting text from a DOCX file (for the purposes of search etc). The value refers to the amount of text inside the file (in bytes), ignoring any images.

cluster.login.rememberme.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.

confluence.cluster.hazelcast.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.
confluence.document.conversion.threads			
5.7		Confluence	The number of threads allocated to the file conversion service is calculated dynamically based on the amount of memory assigned to the instance and the number of CPU cores (usually 4 to 6 threads). This property can be used to change the number of threads. Decrease threads to resolve OOME issues, increase threads to resolve problems with documents spending too long in the queue.
confluence.document.conversion.threads.wait			
5.7	1000	Confluence	Set this property to change the maximum number of items that can be queued for conversion.Any file conversion requests that are made when this maximum limit has been reached are aborted.
confluence.cluster.node.name			
5.7		Confluence	Set this property to give each node in your Data Center cluster a human readable name (displayed in email notifications and in the footer). If left unset, Confluence will assign a node identifier to each node.
confluence.document.conversion.fontpath			
5.8.7		Confluence	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).
confluence.document.conversion.words.defaultfontname			
5.8.7		Confluence	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).
confluence.document.conversion.slides.defaultfontname.regular			

5.8.7		Confluence	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).
confluence.document.conversion.slides.defaultfontname.asian			
5.8.7	TakaoPGothic	Confluence	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).
confluence.document.conversion.slides.defaultfontname.symbol			
5.8.7		Confluence	Set this property to change the default font for rendering symbols in Powerpoint (`.ppt` and `.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).

Recognised System Properties

Confluence supports some configuration and debugging settings that can be enabled through Java system properties. System properties are usually set by passing the `-D` flag to the Java virtual machine in which Confluence is running. See the full instructions: Configuring System Properties.

Since	Default Value	Module...	Effect
atlassian.forceSchemaUpdate			
1.0	`false`	atlassian-config	By default, Confluence will only run its database schema update when it detects that it has been upgraded. This flag will force Confluence to perform the schema update on system startup.
confluence.home			
1.0	Any filesystem path	Confluence and atlassian-config	If this system property is set, Confluence will ignore the contents of the `confluence-init.properties` file, and use this property as the setting for the Confluence Home directory.
confluence.dev.mode			
1.0	`false`	Confluence	Enables additional debugging options that may be of use to Confluence developers (additionally it changes spring bean creation to use lazy initialization by default to decrease startup time). Do not enable this flag on a production system.
confluence.disable.mailpolling			

| 2.4 | `false` | Confluence | If set to "true", will prevent Confluence from retrieving mail for archiving within spaces. Manually triggering "check for new mail" via the web UI will still work. This property has no effect on outgoing mail |

confluence.i18n.reloadbundles

| 1.0 | `true` | Confluence | Setting this property will cause Confluence to reload its i18n resource bundles every time an internationalised string is looked up. This can be useful when testing translations, but will make Confluence run *insanely slowly*. |

confluence.ignore.debug.logging

| 1.0 | `true` | Confluence | Confluence will normally log a severe error message if it detects that DEBUG level logging is enabled (as DEBUG logging generally causes a significant degradation in system performance). Setting this property will suppress the error message. |

confluence.jmx.disabled

| 3.0 | `false` | Confluence | If set to "true", will disable Confluence's JMX monitoring. This has the same effect as setting the "enabled" property to false in `WEB-INF/classes/jmx Context.xml` |

confluence.optimize.index.modulo

| 2.2 | `20` | Confluence | Number of index queue flushes before the index is optimised. |

confluence.plugins.bundled.disable

| 2.9 | `false` | Confluence | Starts confluence without bundled plugins. May be useful in a development environment to make Confluence start quicker, but since bundled plugins are necessary for some of Confluence's core functionality, this property should not be set on a production system. |

atlassian.indexing.contentbody.maxsize

| 3.0 | | atlassian | Confluence stores the body of attachments in the index so the content of attachments are searchable. Set this property to set the maximum indexable size (in bytes). Limiting the size of attachment indexing decreases the size of the index when large attachments are present and prevents excerpts of large attachments being displayed in search results. It does not limit the size of attachments that can be uploaded. |

atlassian.mail.fetchdisabled

| 3.5 | `false` | Confluence | Disables mail fetching services for IMAP and POP |

atlassian.mail.senddisabled

| 3.5 | `false` | Confluence and atlassian-mail | Disables sending of mail |

atlassian.disable.caches

2.4	`true`	atlassian-plugins, atlassian-cache-servlet	Setting this property will disable conditional get and expires: headers on some web resources. This will significantly slow down the user experience, but is useful in devlopment if you are frequently changing static resources and don't want to continually flush your browser cache.
confluence.html.encode.automatic			
2.9		Confluence	Setting this property forces the antixss encoding on or off, overriding the behaviour dictated by settings. The default behaviour differs between Confluence versions.
org.osgi.framework.bootdelegation			
2.10	empty	atlassian-plugins	Comma-separated list of package names to provide from application for OSGi plugins. Typically required when profiling Confluence. For example: "com.jprofiler., **com.yourkit.**".
confluence.diff.pool.size			
3.1	`20`	Confluence	Maximum number of concurrent diffs. When that number is exceeded, additional attempts by RSS feeds to create diffs are ignored and logged. (The RSS requests succeed, they are just missing diffs).
confluence.diff.timeout			
3.1	`1000`	Confluence	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.
confluence.html.diff.timeout			
4.0	`10000`	Confluence	Number of milliseconds to wait for a diff operation (comparing two page versions) to complete before aborting with an error message.
atlassian.user.experimentalMapping			
2.10	`false`	Confluence	Setting this property changes the relationship between local users and local groups to reduce performance degradation when adding a local user to a local group with a large number of users. Please note, setting this property can slow down other user management functions. We recommend that you set it only if you are experiencing performance problems when adding local users to large local groups. Please refer to CONF-1231 9, fixed in Confluence 3.1.1.
confluence.import.use-experimental-importer			
3.2	`false`	Confluence	Setting this property changes Confluence to use the Experimental XML Importer. It is designed to be a more stable implementation but, at the time of the release of 3.2, the importer is largely untested and thus not supported.
atlassian.webresource.disable.minification			
3.3	`false`	atlassian-plugins	Disables automatic minification of JavaScript and CSS resources served by Confluence.

index.queue.thread.count			
3.3	*See "Effect"*	Confluence	Sets the number of threads to be used for the reindex job. The value has to be in the range of 1 to 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.queue.batch.size			
3.3	1500	Confluence	Size of batches used by the indexer. Reducing this value will reduce the load that the indexer puts on the system, but indexing takes longer. Increasing this value will cause indexing to be completed faster, but puts a higher load on the system. Normally this setting does not need tuning.
password.confirmation.disabled			
3.4	`false`	Confluence	This property disables the password confirmation functionality that Confluence uses as an additional security measure. With this property set, Confluence will *not* require password confirmation for the following actions: administrative actions, change of email address and Captcha for failed logins. Disabling password confirmations is useful if you are using a custom authenticator.
confluence.browser.language.enabled			
3.5	`true`	Confluence	Setting this property to "false" disables the detection of browser language headers, effectively restoring Confluence behaviour to that of earlier releases. Setting this property to "true" enables the detection of the language headers sent by the browser. Confluenc will change the UI language based on the browser headers. See documentation on how users can choose a language preference.
upm.pac.disable			
Universal Plugin Manager 1.5	`false`	Universal Plugin Manager (UPM)	When this property is set to true, then UPM will not try to access the The Atlassian Marketplace. This is useful for application servers that do not have access to the Internet. See the UPM documentation.
confluence.reindex.documents.to.pop			

3.5.9	20	Confluence	Indicates how many objects each indexing thread should process at a time during a full re-index. Please note that this number does not include attachments
confluence.reindex.attachments.to.pop			
3.5.9	10	Confluence	Indicates how many attachments each indexing thread should process at a time during a full re-index.
confluence.upgrade.active.directory			
3.5.11	false	Confluence	Forces Confluence to treat any LDAP directories it migrates as Active Directory, rather than relying on looking for sAMAccountName in the username attribute. This is necessary if you are upgrading from before Confluence 3.5, and need to use an attribute other than sAMAccountName to identify your users and are seeing LDAP: error code 4 - Sizelimit Exceeded exceptions in your logs. For more details, see Unable to Log In with Confluence 3.5 or Later Due to 'LDAP error code 4 - Sizelimit Exceeded'
confluence.context.batching.disable			
4.0	false	Confluence	Disables batching for web resources in contexts (e.g. editor, main, admin). Useful for diagnosing the source of javascript or CSS errors.
com.atlassian.logout.disable.session.invalidation			
4.0	false	Confluence	Disables the session invalidation on log out. As of 4.0 the default behaviour is to invalidate the JSession assigned to a client when they log out. If this is set to true the session is kept active (but the user logged out). This may be valuable when using external authentication systems, but should generally not be needed.
officeconnector.spreadsheet.xlsxmaxsize			
4.0.5	2097152	Office Connector	Indicates the maximum size in bytes of an Excel file that can be viewed using the viewxls macro. If empty, the maximum size defaults to 2Mb. See CONF-21043 f or more details.
com.atlassian.confluence.extra.calendar3.display.events.calendar.maxpercalendar			
	200	Team Calendars	Specifies the maximum number of events per calendar. This property is effective only if the Team Calendars plu gin is installed on your Confluence site.
com.atlassian.confluence.allow.downgrade			
4.3.2, 5.0-OD-10	false	Confluence	Allows Confluence to start up against the home directory of a newer version of Confluence. Note that running Confluence like that is unsupported. You should only turn this on if you know what you are doing. See After Downgrading, Confluence Will No Longer Run for details.
reindex.thread.count			

5.2		Confluence	Sets the number of threads to be used for a one-off reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. There is no default value. This system property does not affect the incremental indexing that Confluence does.
reindex.attachments.thread.count			
5.2	4	Confluence	Sets the number of threads to be used for reindexing attachments specifically, and allows you to reduce the concurrency for these more memory intensive index items.
confluence.mbox.directory			
5.4.1		Confluence	Setting this property defines the directory on your Confluence Server where mailboxes can be imported from (for use with the Confluence Mail Archiving add-on). Mailboxes are not able to be imported from any other location. We recommend administrators create a directory in the Confluence Home directory specifically for this purpose. Mail cannot be imported from the server until this system property is set.
confluence.search.max.results			
5.5	1000	Confluence	Setting this property changes the maximum number of items Confluence Search will return.
confluence.upgrade.recovery.file.enabled			
5.5	`true`	Confluence	By default, Confluence creates an upgrade recovery file before and after an upgrade. The operation can take long time on large databases and can be safely turned off if there is a process to back up database and verify the backup before performing an upgrade. Setting this property to false will disable upgrade recovery file creation.
confluence.junit.report.directory			
5.5		Confluence	Setting this property defines the directory on your Confluence Server where JUnit Reports can be imported from (for use in the JUnit Report Macro). No other locations are permitted. We recommend administrators create a directory in the Confluence Home directory specifically for this purpose. JUnit Test result files cannot be imported from the server until this system property is set.
officeconnector.textextract.word.docxmaxsize			
5.5.3	16777216	Confluence	This property defines the maximum amount of memory that can be allocated to extracting text from a DOCX file (for the purposes of search etc). The value refers to the amount of text inside the file (in bytes), ignoring any images.
cluster.login.rememberme.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.

confluence.cluster.hazelcast.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.

confluence.document.conversion.threads

5.7		Confluence	The number of threads allocated to the file conversion service is calculated dynamically based on the amount of memory assigned to the instance and the number of CPU cores (usually 4 to 6 threads). This property can be used to change the number of threads. Decrease threads to resolve OOME issues, increase threads to resolve problems with documents spending too long in the queue.

confluence.document.conversion.threads.wait

5.7	1000	Confluence	Set this property to change the maximum number of items that can be queued for conversion.Any file conversion requests that are made when this maximum limit has been reached are aborted.

confluence.cluster.node.name

5.7		Confluence	Set this property to give each node in your Data Center cluster a human readable name (displayed in email notifications and in the footer). If left unset, Confluence will assign a node identifier to each node.

confluence.document.conversion.fontpath

5.8.7		Confluence	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).

confluence.document.conversion.words.defaultfontname

5.8.7		Confluence	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).

confluence.document.conversion.slides.defaultfontname.regular

5.8.7		Confluence	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).

confluence.document.conversion.slides.defaultfontname.asian

5.8.7	TakaoPGothic	Confluence	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).

confluence.document.conversion.slides.defaultfontname.symbol

5.8.7		Confluence	Set this property to change the default font for rendering symbols in Powerpoint (`.ppt` and `.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).

Working with Confluence Logs

Confluence uses Apache's log4j logging service. This allows a developer or administrator to control the logging behavior and the log output file by editing a configuration file, without touching the application binary. There are six known log4j logging levels.

If you request help from Atlassian Support, we will almost always ask for the `atlassian-confluence.log` from the `confluence-home/logs` directory. To access the logs, go to go to

⚙ ▾

> **General Configuration** > **Atlassian Support Tools**.

If you cannot access logs via Support Tools, check the properties file at `<confluence-installation>/confluence/WEB-INF/classes/confluence-init.properties`, look for the `confluence.home` setting in that file, then find the logs in the Confluence home directory.

On this page:

- Finding the Confluence Log Files
- Finding the Log Configuration File
- Changing the Destination of the Log Files
- Changing the Logging Levels
- Using Some Specific Confluence Logging Options
- Scanning Log Files for Known Problems
- Notes

Related pages:

- Enabling Detailed SQL Logging
- Enabling user access logging
- Generating a Thread Dump
- Enabling Page Request Profiling

Finding the Confluence Log Files

This section describes Confluence's default logging behaviour, assuming that you have not changed the destination of the logs. In order to unify logging across different application servers, Confluence uses the `atlassian-confluence.log` as its primary log, not the application server log.

Both the Confluence and Confluence EAR/WAR distributions follow the same default behaviour:

- When you start Confluence, log entries will be sent to the application server logs until Confluence has completed its initial bootstrap. Any log entries written to the console will be repeated into the log in the Confluence home directory as described below.
- Once the initial startup sequence is complete, all logging will be to `<confluence-home>/logs/atlassian-confluence.log`. For example: `c:/confluence/data/logs/atlassian-confluence.log`.

Note that the default location is the Confluence **home directory**, not the application server's log file. The

home directory is specified in `<confluence-installation>/confluence/WEB-INF/classes/confluence-init.properties`.

Finding the Log Configuration File

Confluence's logging behaviour is defined in the following properties file:
`<CONFLUENCE-INSTALL>/confluence/WEB-INF/classes/log4j.properties`

This file is a standard log4j configuration file, as described in the Apache log4j documentation.

Changing the Destination of the Log Files

Terminology: In log4j, an output destination is called an 'appender'.

To change the destination of the log files, you need to stop Confluence and then change the settings in the '**Logging Location and Appender**' section of the `log4j.properties` file. The location of this file is described above.

In the standard properties file, you will find entries for two appenders:

- `com.atlassian.confluence.logging.ConfluenceHomeLogAppender` – This is a custom appender which controls the default logging destination described above. This appender allows the following settings:
    - MaxFileSize
    - MaxBackupIndex
- `org.apache.log4j.RollingFileAppender` – If you want to log to a different location, uncomment the `RollingFileAppender` line and change the destination file in the line below it. Comment out the previous lines referring to the `ConfluenceHomeLogAppender`.

Confluence ships with the full suite of appenders offered by log4j. Read more about appenders in the log4j documentation.

Changing the Logging Levels

See Configuring Logging for instructions on how to change the logging configuration of Confluence.

Using Some Specific Confluence Logging Options

This section contains some pointers to specific log configurations you may need.

Log the Details of SQL Requests made to the Database

You may want to increase Confluence's logging so that it records individual SQL requests sent to the database. This is useful for troubleshooting specific problems.

You can enable detailed SQL logging in two ways:

- At runtime – see instructions above.
- Via the logging properties file – see the detailed instructions.

Log the Details of Users Viewing/Accessing each Confluence Page

You can configure the log to show which users are accessing which pages in Confluence. This can only be done via the logging properties file – see the detailed instructions.

Scanning Log Files for Known Problems

Confluence provides an inbuilt log scanner that will check your Confluence logs for errors and attempt to match them against known issues in our knowledge base and bug tracker. See Troubleshooting Problems and Requesting Technical Support.

Notes

- **Finding the thread dumps.** Thread dumps are logged to the application server log file.

Configuring Logging

We recommend that you configure Confluence's logging to your own requirements. You can change the log settings in two ways:

- Configure logging in Confluence Administration – Your changes will be in effect only until you next restart Confluence.
- Edit the properties file – Your changes will take effect next time you start Confluence, and for all subsequent sessions.

Both methods are described below. In some rare circumstances you may *also* need to configure the logging.properties file.

Terminology: In log4j, a 'logger' is a named entity. Logger names are case-sensitive and they follow a hierarchical naming standard. For example, the logger named com.foo is a parent of the logger named com.foo.Bar.

Configure logging in Confluence Administration

You can change some of Confluence's logging behaviour via the **Administration Console** while Confluence is running. Any changes made in this way will apply only to the currently-running Confluence lifetime. The changes are not written to the log4j.properties file and are therefore discarded when you next stop Confluence.

Not all logging behaviour can be changed via the Administration Console. For logging configuration not mentioned below, you will need to stop Confluence and then edit the logging properties file instead.

The '**Logging and Profiling**' screen shows a list of all currently defined loggers. On this screen you can:

- Turn page profiling on or off.
- Turn detailed SQL logging on or off.
- Add a new logger for a class/package name.
- Remove a logger for a class/package name.
- Set the logging level (INFO, WARN, FATAL, ERROR or DEBUG) for each class or package name.
- Reset all logging levels to a predefined profile.

Changing the logging configuration

1. Choose the **cog icon**

   ⚙️ ▾

   , then choose **General Configuration** under Confluence Administration
2. Select '**Logging and Profiling**' in the '**Administration**' section of the left-hand panel.
   ℹ️ You need to have System Administrator permissions in order to perform this function.
3. The '**Logging and Profiling**' screen appears, as shown below. Use the following guidelines to change the logging behaviour while Confluence is running:
   - '**Performance Profiling**' — See Page Request Profiling.
   - '**SQL Logging**' — Click the '**Enable SQL Logging**' button to log the details of SQL requests made to the database.
     ℹ️ If you need to enable logging of SQL parameter values, you will need to change the setting in the properties file. This option is not available via the Administration Console.
   - '**Log4j Logging**' — Click one of the profile buttons to reset all your loggers to the predefined profiles:
     - The '**Production**' profile is a fairly standard profile, recommended for normal production conditions.
     - The '**Diagnostic**' profile gives more information, useful for troubleshooting and debugging. It results in slower performance and fills the log files more quickly.
   - '**Add New Entry**' — Type a class or package name into the text box and click the '**Add Entry**' button. The new logger will appear in the list of '**Existing Levels**' in the lower part of the screen.
   - '**Existing Levels**' - These are the loggers currently in action for your Confluence instance.
     - You can change the logging level by selecting a value from the '**New Level**' dropdown list. Read the Apache documentation for a definition of each level.

- Click the '**Remove**' link to stop logging for the selected class/package name.
4. Click the '**Save**' button to save any changes you have made in the '**Existing Levels**' section.

Screenshot: Changing Log Levels and Profiling

Logging and Profiling

Performance Profiling
Profiling is currently OFF.

[Enable Profiling]

SQL Logging

[Enable SQL Logging]

Log4j Logging

loglevel.profileDesc

[Production] [Diagnostic]

OR:

Customise specific logging settings

Add New Entry

Class/Package Name	New Level	
	INFO ▾	Add entry

Existing Levels

Class/Package Name	Current Level	New Level	
atlassian.plugin	INFO	INFO ▾	Remove
com.atlassian.confluence.admin.actions	INFO	INFO ▾	Remove

Editing the Properties File

To configure the logging levels and other settings on a permanent basis, you need to stop Confluence and then change the settings in the `log4j.properties` file, described above.

The properties file contains a number of entries for different loggers that can be uncommented if you are interested in logging from particular components. Read more in the Apache log4j documentation.

See Working with Confluence Logs for some guidelines on specific configuration options you may find useful.

Configuring Levels for java.util.logging in logging.properties

A few libraries used by Confluence use java.util.logging rather than log4j or slf4j. These libraries include:

- com.sun.jersey
- org.apache.shindig
- net.sf.ehcache

Confluence's `logging.properties` file is set to redirect java.util.logging at specific levels to log4j via slf4j.

To increase logging levels for these libraries you must first configure the `logging.properties` file in `<CONFL UENCE-INSTALL>/confluence/WEB-INF/classes/`. The logging levels are different from log4j and are listed here.
For example, to increase logging for shindig change the following line in the `logging.properties` file:

```
org.apache.shindig.level = INFO
```

to

```
org.apache.shindig.level = FINE
```

And then use one of the methods above **as well** to configure the log4j level.

log4j Logging Levels

Logging Levels

- **DEBUG** - designates fine-grained informational events that are most useful to debug an application (*what is going on*)
- **INFO** - announcements about the normal operation of the system - scheduled jobs running, services starting and stopping, user-triggered processes and actions
- **WARN** - any condition that, while not an error in itself, may indicate that the system is running sub-optimally
- **ERROR** - a condition that indicates something has gone wrong with the system
- **FATAL** - a condition that indicates something has gone wrong so badly that the system can not recover
- **TRACE** - n/a within confluence

> There are two ways to modify the logging levels, as described in Working with Confluence Logs.
>
> 1. Modifying the runtime log levels via the **Administration Console**.
> 2. Manually modifying the `<Confluence-Install>\confluence\WEB-INF\classes\log4j.properties` file.

Default Log Level

The standard Confluence log level **WARN** is a way for Confluence to communicate with the server administrator. Logging at WARN level and higher should be reserved for situations that require some kind of attention from the server administrator, and for which corrective action is possible.

See log4j manual for more information.

Troubleshooting SQL Exceptions

If you get an exception similar to those shown below, it is a good idea to increase the logging levels of your Confluence instance. If you request Atlassian support, this additional logging will help us work out the cause of the error.

Increased logging levels will enable us to diagnose errors like these:

```
org.springframework.dao.DataIntegrityViolationException: (HibernateTemplate): data
integrity violated by SQL ''; nested exception is java.sql.BatchUpdateException:
Duplicate entry '1234' for key 1
 at
org.springframework.jdbc.support.SQLStateSQLExceptionTranslator.translate(SQLStateS
QLExceptionTranslator.java:88)
caused by: java.sql.BatchUpdateException: Duplicate entry '1234' for key 1
 at
com.mysql.jdbc.ServerPreparedStatement.executeBatch(ServerPreparedStatement.java:64
7)
```

or

```
(HibernateTemplate): data integrity violated by SQL ''; nested exception is
java.sql.BatchUpdateException: ORA-00001: unique constraint
(CONFLUENCE.SYS_C0012345) violated
```

This document outlines the steps to take to increasing logging on your system.

> **Changing the logging levels via the Administration Console**
> With Confluence 2.7 and later, you can adjust logging levels at runtime via the Administration Console
> — read the instructions. Below we tell you how to edit the log4j files directly.

1. Open `confluence/WEB-INF/classes/log4j.properties` and uncomment the following lines. The double ## lines are comments, leave them intact.

   ```
   ## log hibernate prepared statements/SQL queries (equivalent to setting
   'hibernate.show_sql' to 'true')
   #log4j.logger.net.sf.hibernate.SQL=DEBUG

   ## log hibernate prepared statement parameter values
   #log4j.logger.net.sf.hibernate.type=DEBUG
   ```

   ⓘ If you can not locate these lines in your `log4j.properties` file, please add them to the end of it.
2. Restart Confluence.
3. Redo the steps that led to the error.
4. Zip up your logs directory and attach it your support ticket.
5. If you are using Oracle and received a **constraint error**, please ask your database administrator which *table* and *column* the constraint (that is, `CONFLUENCE.SYS_C0012345`) refers to and add that information to your support ticket.
6. Open `confluence/WEB-INF/classes/log4j.properties` again and remove the 4 lines you added in step 1. (The additional logging will impact performance and should be disabled once you have completed this procedure.)

RELATED TOPICS

Enabling Detailed SQL Logging
Working with Confluence Logs
Troubleshooting failed XML site backups

Configuring Confluence Security

This section gives guidelines on configuring the security of your Confluence site:

- Confluence Security Overview and Advisories

- Confluence Cookies
- Configuring Secure Administrator Sessions
- Using Fail2Ban to limit login attempts
- Securing Confluence with Apache
- Trackback and External Referrers
- Best Practices for Configuring Confluence Security
- Hiding the People Directory
- Configuring Captcha for Spam Prevention
- Hiding External Links From Search Engines
- Configuring Captcha for Failed Logins
- Configuring XSRF Protection
- User Email Visibility
- Anonymous Access to Remote API
- Running Confluence Over SSL or HTTPS
- Connecting to LDAP or JIRA applications or Other Services via SSL
- Configuring RSS Feeds
- Preventing and Cleaning Up Spam

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 Crowd before being stored in the database. There is no mechanism within Confluence to retrieve a user's password – when password recovery is performed, a reset password link is generated and mailed to the user's registered address.

When external user management is enabled, password storage is delegated to the external system.

On this page:

- Application Security Overview
- Finding and Reporting a Security Vulnerability
- Publication of Confluence Security Advisories
- Severity Levels
- Our Security Bugfix Policy
- Published Security Advisories

Buffer Overflows

Confluence is a 100% pure Java application with no native components. As such it is highly resistant to buffer overflow vulnerabilities – possible buffer overruns are limited to those that are bugs in the Java Runtime Environment itself.

SQL Injection

Confluence interacts with the database through the Hibernate Object-Relational mapper. Database queries are generated using standard APIs for parameter replacement rather than string concatenation. As such, Confluence is highly resistant to SQL injection attacks.

Script Injection

Confluence is a self-contained Java application and does not launch external processes. As such, it is highly resistant to script injection attacks.

Cross-Site Scripting

As a content-management system that allows user-generated content to be posted on the web, precautions have been taken within the application to prevent cross-site scripting attacks:

- The wiki markup language in Confluence does not support dangerous HTML markup
- Macros allowing the insertion of raw HTML are disabled by default
- HTML uploaded as a file attachment is served with a content-type requesting the file be downloaded, rather than being displayed inline
- Only system administrators can make HTML-level customisations of the application

When cross-site scripting vulnerabilities are found in the Confluence web application, we endeavour to fix them as quickly as possible.

Transport Layer Security

Confluence does not directly support SSL/TLS. Administrators who are concerned about transport-layer security should set up SSL/TLS at the level of the Java web application server, or the HTTP proxy in front of the Confluence application.

For more information on configuring Confluence for SSL, see: Running Confluence Over SSL or HTTPS

Session Management

Confluence delegates session management to the Java application server in which it is deployed. We are not aware of any viable session-hijacking attacks against the Tomcat application server shipped with Confluence. If you are deploying Confluence in some other application server, you should ensure that it is not vulnerable to session hijacking.

Plugin Security

Administrators install third party plugins **at their own risk**. Plugins run in the same virtual machine as the Confluence server, and have access to the Java runtime environment, and the Confluence server API.

Administrators should always be aware of the source of the plugins they are installing, and whether they trust those plugins.

Administrator Trust Model

Confluence is written under the assumption that anyone given System Administrator privileges is trusted. System administrators are able, either directly or by installing plugins, to perform any operation that the Confluence application is capable of.

As with any application, you should not run Confluence as the root/Administrator user. If you want Confluence to listen on a privileged network port, you should set up port forwarding or proxying rather than run Confluence with additional privileges. The extra-careful may consider running Confluence inside a `chroot` jail.

Stack Traces

To help debug support cases and provide legendary support, Confluence provides stack traces through the web interface when an error occurs. These stack traces include information about what Confluence was doing at the time, and some information about your deployment server.

Only non-personal information is supplied such as operating system and version and Java version. With proper network security, this is not enough information to be considered dangerous. No usernames or passwords are included.

Finding and Reporting a Security Vulnerability

Atlassian's approach to reporting security vulnerabilities is detailed in How to Report a Security Issue.

Publication of Confluence Security Advisories

Atlassian's approach to releasing security advisories is detailed in Security Advisory Publishing Policy.

Severity Levels

Atlassian's approach to ranking security issues is detailed in Severity Levels for Security Issues.

Our Security Bugfix Policy

Our approach to releasing patches for security issues is detailed in our Security Bugfix Policy.

Published Security Advisories

Confluence Cookies

This page lists cookies stored in Confluence users' browsers which are generated by Confluence itself. This page does not list cookies that may originate from 3rd-party Confluence plugins.

Authentication cookies

Confluence uses Seraph, an open source framework, for HTTP cookie authentication. Confluence uses two types of cookies for user authentication:

- The JSESSIONID cookie is created by the application server and used for session tracking purposes. This cookie contains a random string and the cookie expires at the end of every session or when the browser is closed.
- The 'remember me' cookie, `seraph.confluence`, is generated by Confluence when the user selects the **Remember me** check box on the login page.

ⓘ You can read about cookies on the Wikipedia page about HTTP cookies.

On this page:

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 **Rememb er me** check box on the login page.
Cookie key and contents

By default, the cookie key is `seraph.confluence`, which is defined by the `login.cookie.key` parameter in the `CONFLUENCE-INSTALLATION/confluence/WEB-INF/classes/seraph-config.xml` file.

The cookie contains a unique identifier plus a securely-generated random string (i.e. token). This token is generated by Confluence and is also stored for the user in the Confluence database.
Use of cookie for authentication

When a user requests a web page, if the request is not already authenticated via session-based authentication or otherwise, Confluence will match the 'remember me' cookie (if present) against the token (also if present), which is stored for the user in the Confluence database.

If the token in the cookie matches the token stored in the database and the cookie has not expired, the user is authenticated.
Life of 'remember me' cookies

You can configure the maximum age of the cookie. To do that you will need to modify the `CONFLUENCE-INSTA LLATION/confluence/WEB-INF/classes/seraph-config.xml` file and insert the following lines below the other `init-param` elements:

```
<init-param>
 <param-name>autologin.cookie.age</param-name>
 <param-value>2592000</param-value><!-- 30 days in seconds -->
</init-param>
```

Automatic cleanup of 'remember me' tokens

Every cookie issued by Confluence has a corresponding record in the database. A scheduled job runs on the 20th of every month to clean up expired tokens. The name of the trigger is `clearExpiredRememberMeToken sTrigger`.

Note: The only purpose of this job is to prevent the database table from growing too big. For authentication purposes, Confluence will ignore expired tokens even if they still exist in the database.
Is it possible to disable the 'remember me' feature?

Confluence does not offer an option for disabling the 'Remember Me' feature. See the workaround.

Other Confluence cookies

There are several cookies that Confluence uses to store basic 'product presentation' states. Confluence users' authentication details are not stored by these cookies.

Cookie Key	Purpose	Cookie Contents	Expiry
doc-sidebar	Remembers the user's preference for the width of the navigation sidebar in the Confluence documentation theme.	The width of the sidebar in pixels. For example, 300px	One year from the date it was set or was last updated.

confluence.list.pages.cookie	Remembers the user's last chosen tab in the "list pages" section.	The name of the last selected tab. For example, list-content-tree	One year from the date it was set or was last updated.
confluence.browse.space.cookie	Remembers the user's last chosen tab in the "browse space" section	The name of the last selected tab. For example, space-pages	One year from the date it was set or was last updated.
confluence-language	Remembers the user's language chosen on the login page. This cookie relates to a feature that allows a user to change Confluence's language from (and including) the login page, when the language presented to the user prior to logging in is not appropriate.	A locale relating to the chosen language. For example, de_DE	360 days from the date it was set or was last updated.
AJS.conglomerate.cookie	Tracks which general tabs were last used or expansion elements were last opened or closed.	One or more key-value strings which indicate the states of your last general tab views or expansion elements.	One year from the date it is set or was last updated.

Notes

- The *autocomplete* feature in browser text fields (which are typically noticeable when a user logs in to Confluence) is a browser-specific feature, not a Confluence one. Confluence cannot enable or disable this autocompletion, which is typically set through a browser's settings.

Configuring Secure Administrator Sessions

Confluence protects access to its administrative functions by requiring a secure administration session to use the Confluence administration console or administer a space. When a Confluence administrator (who is logged into Confluence) attempts to access an administration function, they are prompted to log in again. This logs the administrator into a temporary secure session that grants access to the Confluence/space administration console.

The temporary secure session has a rolling timeout (defaulted to 10 minutes). If there is no activity by the administrator in the Confluence/space administration console for a period of time that exceeds the timeout, then the administrator will be logged out of the secure administrator session (note, they will remain logged into Confluence). If the administrator does click an administration function, the timeout will reset.

To configure secure administrator sessions:

1. Choose the **cog icon**

   ⚙️ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **Security Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Configure the setting as follows:

- To disable secure administrator sessions, uncheck the **Enable** check box next to **Secure administrator sessions**. When this setting is disabled, administrators will no longer be required to log into a secure session to access the administration console.
- To change the timeout for secure administrator sessions, update the value next to **minutes before invalidation**. The default timeout for a secure administration session is 10 minutes.

5. Choose **Save**.

Notes

- **Disabling password confirmation.** Confluence installations that use a custom authentication mechanism may run into problems with the Confluence security measure that requires password confirmation. If necessary, you can set the `password.confirmation.disabled` system property to disable the password confirmation functionality. See Recognised System Properties. See issue CONF-20 958 "Confluence features that require password confirmation (websudo, captcha) do not work with custom authentication".
- **WebSudo.** The feature that provides secure administrator sessions is also called 'WebSudo'.
- **Manually ending a secure session.** An administrator can choose to manually end their secure session by clicking the '**drop access**' link in the banner displayed at the top of their screen. For example:

> ⚠ You have temporary access to administrative functions. Drop access if you no longer require it. For more information, refer to the documentation.

- **Note for developers.** Secure administrator sessions can cause exceptions when developing against Confluence or deploying a plugin. Please read this FAQ: How do I develop against Confluence with Secure Administrator Sessions? Note: The Confluence XML-RPC and REST APIs are not affected by secure administration sessions.

Using Fail2Ban to limit login attempts

What is Fail2Ban?

We need a means of defending sites against brute-force login attempts. Fail2Ban is a Python application which trails logfiles, looks for regular expressions and works with Shorewall (or directly with iptables) to apply temporary blacklists against addresses that match a pattern too often. This can be used to limit the rate at which a given machine hits login URLs for Confluence.

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

Prerequisites

- Requires Python 2.4 or higher to be installed
- Needs a specific file to follow, which means your Apache instance needs to log your Confluence access to a known logfile. You **should adjust the configuration below** appropriately.

How to set it up

This list is a skeletal version of the instructions

- There's an RPM available for RHEL on the download page, but you can also download the source and set it up manually
- Its configuration files go into `/etc/fail2ban`
- The generic, default configuration goes into `.conf` files (`fail2ban.conf` and `jail.conf`). Don't change these, as it makes upgrading difficult.
- Overrides to the generic configuration go into `.local` files corresponding to the `.conf` files. These only need to contain the specific settings you want overridden, which helps maintainability.
- Filters go into `filter.d` — this is where you define regexps, each going into its own file
- Actions go into `action.d` — you probably won't need to add one, but it's handy to know what's available
- "jails" are a configuration unit that specify one regexp to check, and one or more actions to trigger when the threshold is reached, plus the threshold settings (e.g. more than 3 matches in 60 seconds causes that address to be blocked for 600 seconds)
- Jails are defined in `jail.conf` and `jail.local`. Don't forget the `enabled` setting for each one — it can be as bad to have the wrong ones enabled as to have the right ones disabled.

Running Fail2Ban

- Use `/etc/init.d/fail2ban {start|stop|status}` for the obvious operations
- Use `fail2ban-client -d` to get it to dump its current configuration to STDOUT. Very useful for troubleshooting.
- Mind the CPU usage; it can soak up resources pretty quickly on a busy site, even with simple regexp
- It can log either to syslog or a file, whichever suits your needs better

Common Configuration

jail.local

```
# The DEFAULT allows a global definition of the options. They can be
override
# in each jail afterwards.

[DEFAULT]

# "ignoreip" can be an IP address, a CIDR mask or a DNS host. Fail2ban
will not
# ban a host which matches an address in this list. Several addresses
can be
# defined using space separator.
# ignoreip = <space-separated list of IPs>

# "bantime" is the number of seconds that a host is banned.
bantime  = 600

# A host is banned if it has generated "maxretry" during the last
"findtime"
# seconds.
findtime   = 60

# "maxretry" is the number of failures before a host get banned.
maxretry = 3


[ssh-iptables]

enabled  = false


[apache-shorewall]

enabled  = true
filter   = cac-login
action   = shorewall
logpath = /var/log/httpd/confluence-access.log
bantime = 600
maxretry = 3
findtime = 60
backend = polling
```

Configuring for Confluence

The following is an example only, and you should adjust it for your site.

filter.d/confluence-login.conf

```
[Definition]

failregex = <HOST>.*"GET /login.action

ignoreregex =
```

Securing Confluence with Apache

The following outlines some basic techniques to secure a Confluence instance using Apache. These instructions are basic to-do lists and should not be considered comprehensive. For more advanced security topics see the "Further Information" section below.

- Using Apache to limit access to the Confluence administration interface
- Using Fail2Ban to limit login attempts

Further Information

Running Confluence behind Apache

Using Apache to limit access to the Confluence administration interface

Limiting administration to specific IP addresses

The Confluence administration interface is a critical part of the application; anyone with access to it can potentially compromise not only the Confluence instance but the entire machine. As well as limiting access to users who really need it, and using strong passwords, you should consider limiting access to it to certain machines on the network or internet. If you are using an Apache web server, this can be done with Apache's **Location** functionality as follows:

1. Create a file that defines permission settings

This file can be in the Apache configuration directory or in a system-wide directory. For this example we'll call it "sysadmin_ips_only.conf". The file should contain the following:

```
Order Deny,Allow
 Deny from All

 # Mark the Sysadmin's workstation
 Allow from 192.168.12.42
```

2. Add the file to your Virtual Host

In your Apache Virtual Host, add the following lines to restrict the administration actions to the Systems Administrator:

> This configuration assumes you've installed Confluence under '/confluence'. If you have installed under '/' or elsewhere, adjust the paths accordingly.

```
<Location /confluence/admin>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/consumers/list>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/view-consumer-info>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/service-providers/list>
```

```
    Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/service-providers/add>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/consumers/add>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/consumers/add-manually>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/oauth/update-consumer-info>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/pages/templates/listpagetemplates.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/pages/templates/createpagetemplate.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/spacepermissions.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/pages/listpermissionpages.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/removespace.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/importmbox.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/viewmailaccounts.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/addmailaccount.action?>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/importpages.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/flyingpdf/flyingpdf.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/exportspacehtml.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/spaces/exportspacexml.action>
  Include sysadmin_ips_only.conf
</Location>
<Location /confluence/plugins/servlet/embedded-crowd>
  Include sysadmin_ips_only.conf
</Location>
```

```
<Location /confluence/plugins/servlet/upm>
   Include sysadmin_ips_only.conf
</Location>
```

Trackback and External Referrers

Trackback

When Trackback is enabled, any time you link to an external webpage that supports Trackback Autodiscovery, Confluence will send a trackback ping to that page to inform it that it has been linked to.

Confluence pages also support Trackback Autodiscovery and when Trackback is enabled, can receive trackback pings from other sites.

To enable trackback:

1. Go to
   ⚙ ▾
   > **General Configuration** > **Further Configuration**
2. Choose **Edit**
3. Select the **Trackback** checkbox then **Save**

On this page:

- Trackback
- External referrers
  - Excluding external referrers
  - Hiding external referrers
  - Ignoring external referrers

External referrers

An external referrer is any site that links to your Confluence instance. Each time someone clicks on the external link, your Confluence site can record the click as a referral.

By default, external referrers for a page are listed under '**Hot Referrers**' on the '**Info**' screen of the page. Confluence shows a maximum of 10 referrers. If there are more than 10, confluence shows the 10 with the highest number of hits.

Note that you do *not* need to enable trackback in order to have external referrers enabled.

Screenshot: hot referrers on the page information screen.

Page Information

Title:	The new logo
Author:	Sophie Staunton Feb 15, 2013
Last Changed by:	Sophie Staunton Feb 15, 2013
Tiny Link: (useful for email)	http://localhost:8090/x/HwA8
Export As:	Word · PDF
Operations:	Copy

Labels

There are no labels assigned to this page.

Hot Referrers

(2) https://mail.google.com/mail/u/1/   [ban]
(1) https://atlassian.hipchat.com/chat   [ban]

Excluding external referrers

Confluence Administrators can exclude external referrers to prevent them from being recorded or displayed anywhere on your site. Once you have specified your list of blocked URLs, any incoming links from URLs that match the list will no longer be recorded.

Referrer URLs are blocked if they start with any of the URLs in the exclusion list. So http://evilspamsite.blogspot.com will also match http://evilspamsite.blogspot.com/nastypage.html

You might want to do this if:

1. **You're running a Confluence installation that is open to public**
   In a site that is open to public, one unfortunate problem is that malicious sites can spam the display of a page's incoming links statistics. This is usually done to get the site's URL to appear in the sidebar. By adding these sites to the 'excluded referrers' list, you can prevent them from being listed on your site.
2. **Confluence is installed on a server with multiple domain names or IP addresses**
   Confluence will consider any URL originating from the domain name where Confluence is installed as an internal link. However, if Confluence is installed on a server with multiple domain names or IP addresses, you will need to add the other domain name prefixes to this list to let Confluence know that any links from these domains should not be considered external links.

To add a URL to the excluded referrers list:

1. Go to
   ⚙ ▾
   > **General Configuration** > **Manage Referrers**
2. Enter the URL in the **Referrer URL Prefix** field (you must include http://)
3. Choose **Add**

You can add multiple URLs to the list.

Exclude External Referrers

Referrer URL Prefix [   ]

Add a URL Prefix that will no longer be recorded in the External Referrers. All
URLs that start with this prefix will be excluded.

[Add]

Excluded Referrer URL Prefixes

URL Prefix	Operations
	Purge All
http://evilspamsite.blogspot.com	Delete · Purge

Hiding external referrers

By default, Confluence lists the external referrers as '**Hot Referrers**' on the page information screen for a
page. If you turn this option off, external referrers will not be listed on the page.

To hide external referrers:

1. Go to
   ⚙ ▾
   > **General Configuration** > **Manage Referrers**.
2. Deselect **Show Referrers in Page Info**.

Screenshot: Managing external referrers

External Referrer Settings

Record External Referrers ☑

Allow Confluence to record the external URLs that link to Confluence pages.
More about External Referrers

Show Referrers in Page Info ☐

Show the top 10 external links pointing to that page on the 'Info' screen.

[Submit] Cancel

Ignoring external referrers

An external referrer is any site that links to your Confluence instance. Each time someone clicks on the
external link, your Confluence site can record the click as a referral. By default, Confluence records the
number of hits made to a page from any link on an external site. If you turn this option off, Confluence will not
record the hits.

To ignore external referrers:

1. Go to
   ⚙ ▾
   > **General Configuration** > **Manage Referrers**
2. Deselect **Record External Referrers**

Screenshot: Managing external referrers

External Referrer Settings

Record External Referrers ☐

Allow Confluence to record the external URLs that link to Confluence pages.
More about External Referrers

Show Referrers in Page Info ☐

Show the top 10 external links pointing to that page on the 'Info' screen.

Submit   Cancel

Best Practices for Configuring Confluence Security

The best way to harden a system is to look at each of the involved systems individually. Contact your company's security officer or department to find out what security policies you should be using. There are many things to consider, such as the configuration of your underlying operating systems, application servers, database servers, network, firewall, routers, etc. It would be impossible to outline all of them here.

This page contains guidelines on good security practices, to the best of our knowledge.

Configuring the Web Server

Please refer to the following guides for system administrators:

- How to configure Apache to lock down the administration interface to those people who really need it: Using Apache to limit access to the Confluence administration interface.
- How to reduce the risk of brute force attacks: Using Fail2Ban to limit login attempts.

Configuring the Application Server

See the following system administrator guide for general hints on the application server level:

- Tomcat security best practices

Configuring the Application

The way you set up Confluence roles, permissions and processes makes a big difference in the security of your Confluence site.

Below are some more Confluence-specific items to consider. None of these provides 100% security. They are measures to reduce impact and to slow down an intruder in case your system does become compromised.

- Keep the number of Confluence administrators extremely low. For example, 3 system administrator accounts should be the maximum.
- Similarly, restrict the number of users with powerful roles or group memberships. If only one department should have access to particularly sensitive data, then do restrict access to the data to those users. Do not let convenience over-rule security. Do not give all staff access to sensitive data when there is no need.
- The administrators should have separate Confluence accounts for their administrative roles and for their day to day roles. If John Doe is an administrator, he should have a regular user account without administrator access to do his day to day work (such as writing pages in the wiki). This could be a 'john.doe' account. In addition, he should have an entirely separate account (that cannot be guessed by an outsider and that does not even use his proper name) for administrative work. This account could be 'jane smith' – using a username that is so obscure or fake that no outsider could guess it. This way, even if an attacker singles out the actual person John Doe and gets hold of his password, the stolen account would most likely be John's regular user account, and the attacker cannot perform administrative actions with that account.

- Lock down administrative actions as much as you can. If there is no need for your administrators to perform administrative actions from outside the office, then lock down access to those actions to known IP adresses, for example. See Using Apache to limit access to the Confluence administration interface.
- Put documented procedures in place for the case of employees leaving the company.
- Perform security audits regularly. Know who can help in case a security breach occurs. Perform 'what if' planning exercises. ('What is the worst thing that could happen if a privileged user's password were stolen while he's on vacation? What can we do to minimise damage?').
- Make sure the Confluence database user (and all datasource database users) only has the amount of database privileges it really needs.
- Monitor your binaries. If an attacker compromises an account on your system, he will usually try to gain access to more accounts. This is sometimes done by adding malicious code, such as by modifying files on the system. Run routine scripts that regularly verify that no malicious change has been made.

As another precaution:

- Regularly monitor the above requirements. There are many things that could start out well, but deteriorate over time:
    - A system may start out with just 3 administrators, but over the course of a year this could grow to 30 administrators if no one prevents expansion.
    - Apache administration restrictions may be in place at the start of the year, but when the application server is migrated after a few months, people may forget to apply the rules to the new system.

Again, keep in mind that the above steps may only be a fraction of what could apply to you, depending on your security requirements. Also, keep in mind that none of the above rules can guarantee anything. They just make it harder for an intruder to move quickly.

Hiding the People Directory

The People Directory provides a list of all users in your Confluence system.

If you need to disable the People Directory set the following system properties on your application server command line:

- **To disable the People Directory for anonymous users:**

```
-Dconfluence.disable.peopledirectory.anonymous=true
```

- **To disable the People Directory entirely:**

```
-Dconfluence.disable.peopledirectory.all=true
```

This workaround will prevent the People directory from appearing on the dashboard, but if you navigate to the profile of a user, and then click on the "People" in the breadcrumb link (Dashboard >> **People** >> FullName >> Profile) or you go to the URL directly `<CONFLUENCE_INSTALL>/browsepeople.action`, you will be able to access the people directory.

To workaround this, set up your Apache webserver in front of Confluence and redirect requests to this URL.

Configuring Captcha for Spam Prevention

If your Confluence site is open to the public (you allow anonymous users to add comments, create pages etc) you may find that automated spam is being added, in the form of comments or new pages.

Related pages:

- Configuring Confluence Security

You can configure Confluence to deter automated spam by asking users to prove that they are human before they are allowed to:

- Sign up for an account.
- Add a comment.
- Create a page.

- Edit a page.
- Send a request to the Confluence administrators.

Captcha is a test that can distinguish a human being from an automated agent such as a web spider or robot. When Captcha is switched on, users will see a distorted picture of a word, and must enter it in a text field before they can proceed.

Screenshot: Example of a Captcha test

By default, Captcha is disabled. When enabled, the default is that only anonymous users will have to perform the Captcha test when creating comments or editing pages. You can also choose to enforce Captcha for all users or members of particular groups.

You need System Administrator permissions to configure Captcha for spam prevention in Confluence.

To enable Captcha for spam prevention in Confluence:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Spam Prevention** in the left-hand panel
3. Choose **ON** to turn on Captcha
4. If you want to disable Captcha for certain groups:
   - Select **No one** if you want everyone to see Captchas.
   - Select **Signed in users** if you want only anonymous users to see Captchas.
   - If you want everyone to see Captchas except members of specific groups, select **Members of the following groups** and enter the group names in the text box.
     You can click the magnifying-glass icon to search for groups. Search for all or part of a group name and click the **Select Groups** button to add one or more groups to the list.
   - To remove a group from the list, delete the group name
5. Choose **Save**

Hiding External Links From Search Engines

Hiding external links from search engines helps to discourage spammers from posting links on your site. If you turn this option on, any URLs inserted in pages and comments will be given the 'nofollow' attribute, which prevents search engines from following them.

ℹ️ Shortcut links (e.g. CONF-2622@JIRA) and internal links to other pages within Confluence are not tagged.

To hide external links from search engines:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Click '**Security Configuration**' in the left panel.
3. This will display the '**Security Configuration**' screen. Click '**Edit**'.
4. Check the '**Hide External Links From Search Engines**' checkbox.
5. Click the '**Save**' button.

> **Background to the nofollow attribute**
> As part of the effort to combat the spamming of wikis and blogs (Confluence being both), Google came up with some markup which instructs search engines not to follow links. By removing the main benefit of wiki-spamming it's hoped that the practice will stop being cost-effective and eventually die out.

Configuring Captcha for Failed Logins

If you have confluence administrator permissions, you can configure Confluence to impose a maximum number of repeated login attempts. After a given number of failed login attempts (the default is three) Confluence will display a Captcha form asking the user to enter a given word when attempting to log in again. This will prevent brute force attacks on the Confluence login screen.

Similarly, after three failed login attempts via the XML-RPC or SOAP API, an error message will be returned instructing the user to log in via the web interface. Captcha will automatically be activated when they attempt this login.

'Captcha' is a test that can distinguish a human being from an automated agent such as a web spider or robot. When Captcha is activated, users will need to recognise a distorted picture of a word, and must type the word into a text field. This is easy for humans to do, but very difficult for computers.

Screenshot: example of a Captcha test

Enabling, Disabling and Configuring Captcha for Failed Logins

By default, Captcha for failed logins is enabled and the number of failed login attempts is set to three.

To enable, disable and configure Captcha for failed logins:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose '**Security Configuration**' from the left menu.
3. Choose '**Edit**'.
4. To enable Captcha:
     - Select the '**Enable**' checkbox next to '**CAPTCHA on login**'.
     - Set the maximum number of failed logins next to '**Maximum Authentication Attempts Allowed**'. You must enter a number greater than zero.
5. To **disable** Captcha, deselect the '**Enable**' checkbox.
6. Choose '**Save**'.

Screenshot: Configuring Captcha for failed logins

Security and Privacy

Settings for user management, site security and user privacy.

☐ External user management

Delegate user management to JIRA. More about User Management

☑ Append wildcards to user and group searches

☑ Hide External Links From Search Engines

This helps discourage spammers from posting malicious links by preventing search engines to follow the site. More about External Links

☐ Anonymous Access to Remote API

Enabling this will allow 'anonymous' to access Confluence remotely.

☐ Custom Stylesheets for Spaces

☐ Show system information on the 500 page

User email visibility | public ▾

Maximum RSS Items | 200

Limit the maximum number of items an RSS Feed can request.

RSS timeout | 60 | seconds

The time in seconds allowed to create each RSS Feed. Any items rendered within the timeout will still be returned.

Page timeout | 120 | seconds

The time in seconds allowed to render the content of each wiki Page. Pages taking longer to render will display a timeout error to the user. The default is 120 seconds.

CAPTCHA on login | ☑ Enable

3 | maximum authentication attempts allowed

Secure administrator sessions | ☑ Enable

10 | minutes before automatic invalidation

Requires Confluence Administrator login to access administration functions. More about Secure Administrator Sessions

Notes

- **Disabling all password confirmation requests, including Captcha on login.** Confluence installations that use a custom authentication mechanism may run into problems with the Confluence security measure that requires password confirmation. If necessary, you can set the `password.conf irmation.disabled` system property to disable the password confirmation functionality on administ rative actions, change of email address and Captcha for failed logins. See Recognised System Properties.

Configuring XSRF Protection

Confluence requires an XSRF token to be present on comment creation, to prevent users being tricked into unintentionally submitting malicious data. All the themes bundled with Confluence have been designed to use this feature. However, if you are using a custom theme that does not support this security feature, you can disable it.

⚠ **Please carefully consider the security risks before you disable XSRF protection for comments in your Confluence installation.**

Read more about XSRF (Cross Site Request Forgery) at cgisecurity.com.

To configure XSRF protection for comments:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose **Security Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Uncheck the **Adding Comments** checkbox in the **XSRF Protection** section, to disable XSRF protection.
5. Choose **Save**.

> **Related pages:**
>
> - Configuring Confluence Security
> - Confluence Administrator's Guide
> - Developer documentation on XSRF protection in Confluence

User Email Visibility

Confluence provides three options for email address privacy which can be configured by a Confluence administrator from the **Administration Console**:

- **Public**: email addresses are displayed publicly.
- **Masked**: email addresses are still displayed publicly, but masked in such a way to make it harder for spam-bots to harvest them.
- **Only visible to site administrators**: only Confluence administrators can see the email addresses. Note that, if you select this option, email addresses will not be available in the 'User Search' popup (e.g. when setting Page Restrictions).

To configure user email visibility:

1. Choose the **cog icon**

   ⚙ ▾

   , then choose **General Configuration** under Confluence Administration
2. Choose '**Security Configuration**'.
3. Choose '**Edit**'. The fields on the '**Security Configuration**' screen will be editable.
4. Select one of the options from the '**User email visibility**' dropdown: '**public**', '**masked**', or '**only visible to site administrators**'.
5. Choose '**Save**'.

Screenshot: Email Visibility

Anonymous Access to Remote API

Administrators may wish to disable anonymous access to the Confluence remote API. to make it harder for malicious users to write 'bots' that perform bulk changes to the site.

To disable anonymous access to the remote API:

1. Choose the **cog icon**

   , then choose **General Configuration** under Confluence Administration
2. Choose **Security Configuration** in the left-hand panel. The **Security Configuration** screen will appear.
3. Choose **Edit**.
4. Uncheck the **Anonymous Access to API** check box.
5. Choose **Save**.

Notes

This page is about access to the remote API. If you are looking for information about preventing anonymous users from accessing Confluence, see Global Permissions Overview.

Running Confluence Over SSL or HTTPS

Atlassian applications allow the use of SSL within our products, however Atlassian Support does not provide assistance for configuring it. Consequently, Atlassian **cannot guarantee providing any support for it**.

- If assistance with conversions of certificates is required, please consult with the vendor who provided the certificate.
- If assistance with configuration is required, please raise a question on Atlassian Answers.

This document tells you how to configure Confluence to enable access via HTTPS (HTTP over SSL), so that your Confluence logins and data are encrypted during transport to and from Confluence. SSL encryption is a good way to safeguard your Confluence data and user logins from being intercepted and read by outsiders.

These instructions apply to the following platforms:

- **Confluence** - Apache Tomcat is the application server shipped with Confluence, and is the only supported application server. If you are using a different application server or Apache HTTP Server ("httpd"), see the page on Apache with mod_proxy for instructions on how to terminate an SSL connection at the Apache web server.
- **Java 8** - JDK 1.8 is the supported Java version for Confluence. Note that you need the JDK, since it includes the `keytool` utility used in the instructions below. The JRE is not enough.

ℹ The default connector port for Confluence is 8090.

On this page:

- Step 1. Create or Request a New SSL Certificate
- Step 2. Modify the Server Configuration File in your Confluence Installation
- Step 3. Specify the Location of your Certificate
- Step 4. Change your Confluence Base URL to HTTPS
- Step 5. Add a Security Constraint to Redirect All URLs to HTTPS
- Notes
- Troubleshooting

Related Topics

- SSL Configuration HOW-TO in the Apache Tomcat 8.0 documentation
- keytool - Key and Certificate Management Tool in the Java SE documentation
- Connecting to LDAP or JIRA applications or Other Services via SSL

Step 1. Create or Request a New SSL Certificate

You will need a valid SSL certificate before you can enable HTTPS. If you already have a certificate prepared, skip to step 2 below.

You can choose to create a self-signed certificate or to use a certificate issued by a certificate authority (CA, sometimes also called a 'certification authority'). We described both options below.

Certificate Option 1 – Create a Self-Signed Certificate

Self-signed certificates are useful if you require encryption but do not need to verify the identity of the requesting website. In general, you might use a self-signed certificate on a test environment and on internal

corporate networks (intranets).

Because the certificate is not signed by a certificate authority (CA), users may receive a message that the site is not trusted and may have to perform several steps to accept the certificate before they can access the site. This usually will only occur the first time they access the site.

Follow the steps below to generate a certificate using Java's `keytool` utility. This tool is included in the JDK.

1. Use Java's `keytool`utility to generate the certificate:

   > Many SSL issuers (including but not limited to GoDaddy and RapidSSL) are now requiring a 2048-bit key size. To generate a key with 2048-bit encryption, add '-keysize 2048' to these queries.

   - On Windows, run the following command at the command prompt:

     ```
     "%JAVA_HOME%\bin\keytool" -genkeypair -keysize 2048 -alias tomcat
     -keyalg RSA -sigalg SHA256withRSA
     ```

   - On OS X or UNIX-based systems, run the following command at the command prompt:

     ```
     $JAVA_HOME/bin/keytool -genkeypair -keysize 2048 -alias tomcat
     -keyalg RSA -sigalg SHA256withRSA
     ```

2. When asked for a **password**:
   - Specify the password you want to use for the certificate (private key). Note that the password text will not appear as you type it.
   - Make a note of the password you choose, because you will need it in the next step when editing the configuration file.
   - The default password is '`changeit`'.

     > Tomcat has a known issue with passwords containing special characters. You should use a password that only contains alphanumeric characters.

3. Follow the prompts to specify your name, organisation and location. This information is used to construct the X.500 Distinguished Name (DN) of the entity. The CN ("What is your first and last name?") must match the fully-qualified hostname of the server running Confluence, otherwise Tomcat will not be able to use the certificate for SSL. For example for a Confluence running on a server named "confluence.example.com":
   `CN=confluence.example.com, OU=Java Software Division, O=Sun Microsystems Inc, C=US`
4. Enter '`y`' to confirm the details.
5. When asked for the **password** for '`tomcat`' (the alias you entered in the keytool command above), press the 'Enter' key. This specifies that your keystore entry will have the **same password** as your private key. You MUST use the same password here as was used for the keystore password itself. This is a restriction of the Tomcat implementation.
6. You certificate is now ready. Go to step 2 below.

Certificate Option 2 – Use a Certificate Issued by a Certificate Authority

When running Confluence in a production environment, you will need a certificate issued by a certificate authority (CA, sometimes also called a 'certification authority') such as VeriSign, Thawte or TrustCenter. The instructions below are adapted from the Tomcat documentation.

First you will generate a local certificate and create a 'certificate signing request' (CSR) based on that certificate. You will submit the CSR to your chosen certificate authority. The CA will use that CSR to generate a certificate for you.

1. Use Java's `keytool` utility to generate a local certificate, as described in the previous section.

2. Use the `keytool` utility to generate a CSR, replacing the text `<MY_KEYSTORE_FILENAME>` with the path to and file name of the `.keystore`file generated for your local certificate:

```
keytool -certreq -keyalg RSA -alias tomcat -file certreq.csr -keystore
<MY_KEYSTORE_FILENAME>
```

3. Submit the generated file called `certreq.csr` to your chosen certificate authority. **Refer to the documentation on the CA's website to find out how to do this.**
4. The CA will send you a certificate.
5. Import the new certificate into your local keystore:

```
keytool -importcert -alias tomcat -keystore <MY_KEYSTORE_FILENAME> -file
<MY_CERTIFICATE_FILENAME>
```

Please note that some CAs require you to install an intermediate certificate before importing your certificate. Please refer to your CA documentation to successfully install your certificate.

If you receive an error, and you use Verisign or GoDaddy, you may need to export the certificate to PKCS12 format along with the private key.

1. First, remove the certificate added above from the keystore:

```
keytool -delete -alias tomcat -keystore <MY_KEYSTORE_FILENAME>
```

2. Then export to PKCS12 format:

```
openssl pkcs12 -export -in <MY_CERTIFICATE_NAME> -inkey
<MY_PRIVATEKEY_NAME> -out <MY_PKC12_KEYSTORE_NAME> -name tomcat
-CAfile
<MY_ROOTCERTIFICATE_NAME-alsoCalledBundleCertificateInGoDaddy>
-caname root
```

3. Then import from PKCS12 to jks:

```
keytool -importkeystore -deststorepass
<MY_DESTINATIONSTORE_PASSWORD> -destkeypass
<MY_DESTINATIONKEY_PASSWORD> -destkeystore <MY_KEYSTORE_FILENAME>
-srckeystore <MY_PKC12_KEYSTORE_NAME> -srcstoretype PKCS12
-srcstorepass <MY_PKC12_KEYSTORE_PASSWORD> -alias tomcat
```

Step 2. Modify the Server Configuration File in your Confluence Installation

1. Edit the server configuration file at this location: `{CONFLUENCE-INSTALLATION}>/conf/server.xml`.
2. Uncomment the following lines:

```
<Connector port="8443" maxHttpHeaderSize="8192"
                  maxThreads="150" minSpareThreads="25"
maxSpareThreads="75"
                  enableLookups="false" disableUploadTimeout="true"
                  acceptCount="100" scheme="https" secure="true"
                  clientAuth="false" sslProtocol="TLS" SSLEnabled="true"
                  URIEncoding="UTF-8"
keystorePass="<MY_CERTIFICATE_PASSWORD>"/>
```

3. Replace the text `<MY_CERTIFICATE_PASSWORD>` with the password you specified for your certificate.
4. Make sure that the attribute-value pair `SSLEnabled="true"` is part of the `Connector` element, as shown above. If this attribute is not present, attempts to access Confluence will time out.
5. Save the server configuration file.

Step 3. Specify the Location of your Certificate

By default, Tomcat expects the keystore file to be named `.keystore` and to be located in the user home directory under which Tomcat is running (which may or may not be the same as your own home directory). This means that, by default, Tomcat will look for your SSL certificates in the following location:

- On Windows: `C:\Documents and Settings\\#CURRENT_USER#\.keystore`
- On OS X and UNIX-based systems: `~/.keystore`

You may decide to move the certificate to a custom location. If your certificate is not in the default location, you will need to update your server configuration file as outlined below, so that Tomcat can find the certificate.

1. Edit the server configuration file at this location: `{CONFLUENCE-INSTALLATION}>/conf/server.xml`
2. Add the attribute `keystoreFile="<MY_CERTIFICATE_LOCATION>"` to the `Connector`element, so that the element looks like this:

```
<Connector port="8443" maxHttpHeaderSize="8192"
                  maxThreads="150" minSpareThreads="25"
maxSpareThreads="75"
                  enableLookups="false" disableUploadTimeout="true"
                  acceptCount="100" scheme="https" secure="true"
                  clientAuth="false" sslProtocol="TLS" SSLEnabled="true"
                  URIEncoding="UTF-8"
keystorePass="<MY_CERTIFICATE_PASSWORD>"
                  keystoreFile="<MY_CERTIFICATE_LOCATION>"/>
```

3. Replace the text `<MY_CERTIFICATE_LOCATION>` with the path to your certificate, including the path and the name of the `.keystore` file.
4. Save the server configuration file.

Step 4. Change your Confluence Base URL to HTTPS

1. In your browser, go to the Confluence Administration Console.
2. Change the Server Base URL to HTTPS. See the documentation on configuring the server base URL.
3. Restart Tomcat and access Confluence on `https://<MY_BASE_URL>:8443/`.

Step 5. Add a Security Constraint to Redirect All URLs to HTTPS

Although HTTPS is now activated and available, the old HTTP URLs (http://localhost:8090) are still available. Now you need to redirect the URLs to their HTTPS equivalent. You will do this by adding a security constraint in `web.xml`. This will cause Tomcat to redirect requests that come in on a non-SSL port.

1. Check whether your Confluence site uses the **RSS macro**. If your site has the RSS macro enabled, you may need to configure the URL redirection with a firewall rule, rather than by editing the `web.xml` file. Skip the steps below and follow the steps on the RSS Feed Macro page instead.
2. Otherwise, Edit the file at `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`.
3. Add the following declaration to the end of the file, **before** the `</web-app>`tag:

```
<security-constraint>
  <web-resource-collection>
    <web-resource-name>Restricted URLs</web-resource-name>
    <url-pattern>/</url-pattern>
  </web-resource-collection>
  <user-data-constraint>
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
  </user-data-constraint>
</security-constraint>
```

4. Restart Confluence and access http://localhost:8090. You should be redirected to https://localhost:8443/login.action.

ℹ️ Confluence has two web.xml files. The other one is at `<CONFLUENCE_INSTALLATION>/conf/web.xml`. Please only add the security constraints to `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`, as described above.

Notes

- **Background information on generating a certificate:** The `'keytool -genkeypair'` command generates a key pair consisting of a public key and the associated private key, and stores them in a keystore. The command packages the public key into an X.509 v3 self-signed certificate, which is stored as a single-element certificate chain. This certificate chain and the private key are stored in a new keystore entry, identified by the `alias` that you specify in the command. The Java SE documentation has a good overview of the utility.

- **Custom SSL port:** If you have changed the port that the SSL connector is running on from the default value of 8443, you must update the `redirectPort` attribute of the standard HTTP connector to reflect the new SSL port. Tomcat needs this information to know which port to redirect to when an incoming request needs to be secure.
- **Multiple instances on the same host:** When running more than one instance on the same host, it is important to specify the *address* attribute in the `<CONFLUENCE_INSTALLATION>/conf/server.xml` file because by default the connector will listen on **all available network interfaces**, so specifying the address will prevent conflicts with connectors running on the same default port. See the Tomcat Connector documentation for more about setting the address attribute: https://tomcat.apache.org/tomcat-8.0-doc/config/http.html

```
<Connector port="8443" address="your.confluence.url.com"
maxHttpHeaderSize="8192"
                maxThreads="150" minSpareThreads="25"
maxSpareThreads="75"
                enableLookups="false" disableUploadTimeout="true"
                acceptCount="100" scheme="https" secure="true"
                clientAuth="false" sslProtocol="TLS" SSLEnabled="true"
                URIEncoding="UTF-8"
keystorePass="<MY_CERTIFICATE_PASSWORD>"
                keystoreFile="<MY_CERTIFICATE_LOCATION>"/>
```

- **Protection for logins only or for individual spaces:** As of Confluence 3.0, Atlassian does not support HTTPS for logins only or for specific pages. We support only site-wide HTTPS. To see the reasoning behind this decision, please see CONF-18120 and CONF-4116.

Troubleshooting

- Check the Confluence knowledge base articles on troubleshooting SSL

- Problems with **Internet Explorer being unable to download attachments**: Applying SSL site wide can prevent IE from downloading attachments correctly. To fix this problem, edit `<CONFLUENCE_INST ALLATION>/conf/server.xml` and add the following line within the `<Context ... />`element:

```
<Valve className="org.apache.catalina.authenticator.NonLoginAuthenticator"
        disableProxyCaching="true" securePagesWithPragma="false" />
```

Connecting to LDAP or JIRA applications or Other Services via SSL

> This page documents configuration of SSL, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with SSL. Please be aware that this material is provided for your information only, and that you use it at your own risk.

Related pages:

- Configuring an SSL Connection to Active Directory
- Configuring Web Proxy Support for Confluence
- Running Confluence Over SSL or HTTPS

This page describes how to get Confluence connecting to external servers over SSL, via the various SSL-wrapped protocols.

Here are some examples of when you may need to connect to an external server over SSL/HTTPS:

- You need to connect to an LDAP server, such as Active Directory, if the LDAP server is running over SSL.
  For specific instructions for Active Directory, see Configuring an SSL Connection to Active Directory.
- You want to set up your JIRA application as a trusted application in Confluence, when JIRA is running over SSL.
- You want to refer to an https://... URL in a Confluence macro.

If you want to run Confluence itself over SSL, see Running Confluence Over SSL or HTTPS.

> There's a Confluence SSL plugin that facilitates this process.

Importing SSL Certificates

The following commands apply to JDK 1.5. For commands/syntax relevant to JDK 1.6, please refer to this document from Oracle.

1. Add the root certificate to your default Java keystore with the following command. This is the certificate that was used to authorise the LDAP server's certificate. It will be either the one that was used for signing it, or will come from further up in the trust chain, possibly the root certificate. This is often a self-signed certificate, when both ends of the SSL connection are within the same network. Again, the exact alias is not important.

```
keytool -import -alias serverCert -file RootCert.crt -keystore
%JAVA_HOME%/jre/lib/security/cacerts (Windows)
keytool -import -alias serverCert -file RootCert.crt -keystore
$JAVA_HOME/jre/lib/security/cacerts (Linux/Unix/Mac)
```

2. Import your LDAP or JIRA server's public certificate into the JVM Keystore. This is the certificate that the LDAP server will use to set up the SSL encryption. You can use any alias of your choosing in place of "JIRAorLDAPServer.crt".

```
keytool -import -alias ldapCert -file JIRAorLDAPServer.crt
-keystore %JAVA_HOME%/jre/lib/security/cacerts (Windows)
keytool -import -alias ldapCert -file JIRAorLDAPServer.crt
-keystore $JAVA_HOME/jre/lib/security/cacerts (Linux/Unix/Mac)
```

3. Verify that the certificate has been added successfully by entering the following command:

```
keytool -list -keystore %JAVA_HOME%/jre/lib/security/cacerts
(Windows)
keytool -list -keystore $JAVA_HOME/jre/lib/security/cacerts
(Unix/Linux)
keytool -list -keystore /Library/Java/Home/lib/security/cacerts
(Mac)
```

4. Ensure that you have updated JAVA_OPTS to specify the path to the keystore, as specified in Connecting to SSL services before restarting Tomcat/Confluence.
There is no need to specify an alias for Confluence to use. On connecting to the LDAP server, it will search through the keystore to find a certificate to match the key being presented by the server.

Troubleshooting

Check the following knowledge base articles:

- Unable to Connect to SSL Services due to PKIX Path Building Failed
- SSL troubleshooting articles

Configuring RSS Feeds

A Confluence System Administrator can configure the following aspects of RSS feeds:

- The maximum number of items that Confluence returns to an RSS feed request.
- The maximum time period that Confluence allows to respond to an RSS feed request.

Both of these are set in the 'Edit Security Configuration' screen.

On this page:
- Notes

Related pages:
- The RSS Feed Builder

To configure RSS feeds:

1. Choose the **cog icon** ⚙, then choose **General Configuration** under Confluence Administration
2. Choose **Security Configuration**.
3. Choose **Edit**.
4. Enter a value for **Maximum RSS Items**. The default value is 200.
5. Enter a value for **RSS timeout**.
6. Choose **Save**.

Screenshot: Configuring RSS feeds

Maximum RSS Items	200

Limit the maximum number of items an RSS Feed can request.

RSS timeout	60	seconds

The time in seconds allowed to create each RSS Feed. Any items rendered within the timeout will still be returned.

Notes

- When using the RSS Feed Builder, a user could potentially enter such a large value for the number of feed items returned that Confluence would eventually run out of memory.
- When using the Feed Builder, if a users a value greater than this setting (or less than 0) they will get a validation error.
- If any pre-existing feeds are set to request more than the configured maximum, they will be supplied with only the configured maximum number of items. This is done silently - there is no logging and no message is returned to the RSS reader.
- If Confluence times out when responding to an RSS feed request, any items already rendered are returned.

Preventing and Cleaning Up Spam

If your Confluence site is public-facing you may be affected by spammers.

Stopping Spammers

To prevent spammers:

1. Enable Captcha. See Configuring Captcha for Spam Prevention.
2. Run Confluence behind an Apache webserver and create rules to block the spammer's IP address.

Blocking Spam at Apache or System Level

If a spam bot is attacking your Confluence site, they are probably coming from one IP address or a small range of IP addresses. To find the attacker's IP address, follow the Apache access logs in real time and filter for a page that they are attacking.

For example, if the spammers are creating users, you can look for `signup.action`:

```
$ tail -f confluence.atlassian.com.log | grep signup.action
1.2.3.4 - - [13/Jan/2010:00:14:51 -0600] "GET /signup.action HTTP/1.1"
200 9956 "-" "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1)"
37750
```

Compare the actual spam users being created with the log entries to make sure you do not block legitimate users. By default, Apache logs the client's IP address in the first field of the log line.

Once you have the offender's IP address or IP range, you can add it to your firewall's blacklist. For example, using the popular Shorewall firewall for Linux you can simply do this:

```
# echo "1.2.3.4" >> /etc/shorewall/blacklist
# /etc/init.d/shorewall reload
```

To block an IP address at the Apache level, add this line to your Apache vhost config:

```
Deny from 1.2.3.4
```

You can restart Apache with a "graceful" command which will apply the changes without dropping any current sessions.

If this still does not stop the spam, then consider turning off public signup.

Deleting Spam

Profile Spam

By 'profile spam', we mean spammers who create accounts on Confluence and post links to their profile page.

If you have had many such spam profiles created, the easiest way to delete them is via SQL.

To delete a spam profile:

1. Shut down Confluence and back up your database.
   Note: This step is essential before you run any SQL commands on your database.
2. Find the last real profile:

```
SELECT bodycontentid,body FROM bodycontent WHERE contentid IN
 (SELECT contentid FROM content WHERE contenttype='USERINFO')
 ORDER BY bodycontentid DESC;
```

3. Look through the bodies of the profile pages until you find where the spammer starts. You may have to identify an number of ranges.
4. Find the killset:

```
CREATE TEMP TABLE killset AS SELECT bc.bodycontentid,c.contentid,c.username FROM
  bodycontent bc JOIN content c ON bc.contentid=c.contentid WHERE
  bodycontentid >= BOTTOM_OF_SPAM_RANGE AND bodycontentID <= TOP_OF_SPAM_RANGE

  AND  c.contenttype='USERINFO';

DELETE FROM bodycontent WHERE bodycontentid IN (SELECT bodycontentid FROM killset);

DELETE FROM links WHERE contentid IN (SELECT contentid FROM killset);

DELETE FROM content WHERE prevver IN (SELECT contentid FROM killset);

DELETE FROM content WHERE pageid IN (SELECT contentid FROM killset);

DELETE FROM content WHERE contentid IN (SELECT contentid FROM killset);

DELETE FROM os_user_group WHERE user_id IN (SELECT id FROM killset k JOIN os_user o ON o.username=k.username);

DELETE FROM os_user WHERE username IN (SELECT username FROM killset);
```

If you're using Confluence 5.6 or earlier use the SQL commands below:
⌄ For Confluence 5.6 and earlier...

```
 CREATE TEMP TABLE killset AS SELECT
bc.bodycontentid,c.contentid,c.username FROM
  bodycontent bc JOIN content c ON bc.contentid=c.contentid WHERE
  bodycontentid >= BOTTOM_OF_SPAM_RANGE AND bodycontentID <=
TOP_OF_SPAM_RANGE
  AND  c.contenttype='USERINFO';

DELETE FROM bodycontent WHERE bodycontentid IN (SELECT bodycontentid FROM
killset);

DELETE FROM links WHERE contentid IN (SELECT contentid FROM killset);

DELETE FROM content WHERE prevver IN (SELECT contentid FROM killset);

DELETE FROM attachments WHERE pageid IN (SELECT contentid FROM killset);

DELETE FROM content WHERE contentid IN (SELECT contentid FROM killset);

DELETE FROM os_user_group WHERE user_id IN (SELECT id FROM killset k JOIN
os_user o ON o.username=k.username);

DELETE FROM os_user WHERE username IN (SELECT username FROM killset);
```

5. Once the spam has been deleted, restart Confluence and rebuild the index. This will remove any references to the spam from the search index.

Scheduled Jobs

The administration console allows you to schedule various administrative jobs in Confluence, so that they are executed at regular time intervals. The types of jobs which can be scheduled cover:

- Confluence site backups
- Storage optimisation jobs to clear Confluence's temporary files and caches
- Index optimisation jobs to ensure Confluence's search index is up to date
- Mail queue optimisation jobs to ensure Confluence's mail queue is maintained and notifications have been sent.

ⓘ You'll need System Administrator permissions in order to edit and manually run jobs.

On this page:

- Accessing Confluence's Scheduled Jobs Configuration
- Running a job manually
- Changing a job's schedule
- Disabling/Re-enabling a Job
- Viewing a Job's Execution History
- Types of Jobs
- Cron Expressions

Related pages:

- Trigger Module
- Configuring Backups

Accessing Confluence's Scheduled Jobs Configuration

To access Confluence's Scheduled Jobs configuration page:

⚙ ▾

1. > **General Configuration** > **Scheduled Jobs**
2. All scheduled jobs are listed with:
   - **Status** - the job's status, which is either 'Scheduled' (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

Job	Status	Last Execution	Next Execution	Avg. Duration	Actions
Back Up Confluence	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 02:00	17450	History · Run · Edit · Disable
Check Cluster Safety	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:39	16	History · Run
Clean Journal Entries	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 02:00	882	History · Run · Edit · Disable
Clean Temporary Directory	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 04:00	1107	History · Run · Edit
Clear Expired Mail Errors	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 03:00	653	History · Run · Edit
Clear Expired Remember Me Tokens	Scheduled	Feb 20, 2015 08:17	Mar 20, 2015 00:00	300	History · Run · Edit
Cluster Cache Compaction	Scheduled	Mar 11, 2015 08:10	Mar 11, 2015 12:00	0	History · Run · Disable
EhCache Cache Compaction	Scheduled	Mar 11, 2015 11:34	Mar 11, 2015 11:39	10	History · Run · Disable
Email Daily Reports	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 00:00	6952	History · Run · Edit · Disable
Flush Edge Index Queue	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:38	37	History · Run · Edit · Disable
Flush Index Queue	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:38	20	History · Run · Edit
Flush Local Task Queue	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:39	65	History
Flush Mail Queue	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:39	89	History · Run · Edit · Disable
Flush Task Queue	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:39	47	History · Run · Disable
JIRA Metadata Cache Configuration	Scheduled	Mar 11, 2015 11:19	Mar 11, 2015 12:19	37	History · Run
scheduledjob.desc.conversionQueueMonitor	Scheduled	Mar 11, 2015 11:38	Mar 11, 2015 11:38	1	History · Run · Edit · Disable
scheduledjob.desc.purgeHistoryJob	Scheduled	Mar 11, 2015 08:10	Mar 12, 2015 00:00	1444	History · Edit
Send Recommended Updates Email	Scheduled	Mar 11, 2015 08:10	Mar 11, 2015 12:00	857	History · Run · Disable

Running a job manually

To run a job manually head to the Scheduled Jobs list and choose **Run** next to the job. It will run immediately.

Not all jobs can be run manually.

Changing a job's schedule

To change a job's schedule:

1. Choose **Edit** next to the job you want to change.
2. Enter the new day or time to run the job as a cron expression - there's more info about cron expressions below.
3. **Save** your changes to the job's schedule, or **Revert** back to the default setting.

Not all jobs' schedules are configurable.

Screenshot: Configuring a Job Schedule

Edit Schedule for Clean Index Queue

Advanced Configuration

Cron Expression   `0 0 2 * * ?`

Define schedule using a cron expression.

Future Schedules:

20-Feb-2013	02:00:00
21-Feb-2013	02:00:00
22-Feb-2013	02:00:00
23-Feb-2013	02:00:00
24-Feb-2013	02:00:00
25-Feb-2013	02:00:00
26-Feb-2013	02:00:00
27-Feb-2013	02:00:00
28-Feb-2013	02:00:00
01-Mar-2013	02:00:00

Default   Save   Cancel

Disabling/Re-enabling a Job

By default, all jobs in Confluence are enabled.

Use the **Disable** / **Enable** links in the action column to disable and re-enable each job.

Not all jobs in Confluence can be disabled.

Viewing a Job's Execution History

To see when a job was last run, and how long the job took to run, click the **History** link beside the job.

If a job has not run at least once the History link won't appear.

Screenshot: Job Execution History

History for Clean Temporary Directory

Started	Ended	Duration (ms)
19-Feb-2013 04:00:00	19-Feb-2013 04:00:01	1901
18-Feb-2013 04:00:00	18-Feb-2013 04:00:00	42
17-Feb-2013 04:00:00	17-Feb-2013 04:00:00	334

Types of Jobs

Job Name	Description	Execution Behaviour	Default Schedule
Back Up Confluence	Performs a backup of your entire Confluence site.	Per cluster	At 2am every day
Check Cluster Safety	For clustered Confluence installations, this job ensures that only one Confluence instance in the cluster writes to the database at a time. For standard (non-clustered) editions of Confluence, this job is useful for alerting customers who have accidentally connected a second Confluence instance to a Confluence database which is already in use.	Per cluster	Every 30 seconds
Clean Journal Entries	Periodically clears journal entries that have already been processed to ensure that its size does not grow indefinitely.	Per node	At 2am every day
Clean Temporary Directory	Cleans up temporary files generated in the `<confluence-home>/temp` directory. This temp directory is created by exports etc. This doesn't include the temp directory located in the confluence install directory.	Per node	At 4am every day
Clear Expired Mail Errors	Clears notification errors in the mail error queue. A notification error is sent to the mail error queue whenever the notification fails to be sent due to an error.	Per cluster	At 3am every day
Clear Expired Remember Me Tokens	Clears all expired 'Remember Me' tokens from the Confluence site. Remember Me tokens expire after two weeks.	Per cluster	On the 20th of each month
Email Daily Reports	Emails a daily summary report of all Confluence changes to all subscribers. Since each email report only records changes from the last 24-hour period, it is recommended that you only change the time of this job whilst keeping the job's frequency to 24 hours.	Per cluster	At 12am every day
Flush Edge Index Queue	Flushes the Edge Index Queue so Confluence's search results stay up to date.	Per node	Every 30 seconds
Flush Index Queue	Flushes changes to Confluence's index so that Confluence's search results are up to date. Confluence records each content update in its search index.	Per node	Every 5 seconds
Flush Local Task Queue	Flushes the local task queue. (These are internal Confluence tasks that are typically flushed at a high frequency.)	Per node	Every minute
Flush Mail Queue	Sends notifications that have been queued up in the mail queue.	Per cluster	Every minute
Flush Task Queue	Flushes the task queue. (These are internal Confluence tasks that are typically flushed at a high frequency.)	Per node	Every minute
Send Recommended Updates Email	Triggers sending recommended update emails to users. The job runs hourly, but users will receive the notification weekly or daily, depending on the setting in their profile, at a time that matches their timezone.	Per cluster	Hourly

Cron Expressions

A cron expression is a string of 6-7 'time interval' fields that defines the frequency with which a job is executed. Each of these fields can be expressed as either a numerical value or a special character and each field is separated by at least one space or tab character.

The table below 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 Whitelist

Confluence administrators can choose to allow incoming and outgoing connections and content from specified sources for use in the RSS macro, HTML Include macro and gadgets, by adding URLs to the whitelist.

Confluence will display an error if content has been added that is not from an allowed source, and prompt the user to add the URL to the whitelist.

Application links are automatically added to the whitelist. You don't need to manually add them.

Note: The HTML Include macro is disabled by default.

Add allowed URLs to the whitelist

To add a URL to the whitelist:

1. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
2. Choose **Whitelist**.
3. Enter the URL or expression you want to allow.
4. Choose the **Type** of expression (see below for examples of the types available).
5. Choose **Allow Incoming** if you need to allow CORS requests (see below).
6. Choose **Add**.

Your URL or expression appears in the whitelist.

To test that your whitelisted URL is working as expected you can enter a URL in the **Test a URL** field. Icons will indicate whether incoming and / or outgoing traffic is allowed for that URL.

Confluence 5.9 Documentation

Expression Types

When adding a URL to the whitelist, you can choose from a number of expression types.

Type	Description	Example
Domain name	Allows all URLs from the specified domain.	`http://www.example.com`
Exact match	Allows only the specified URL.	`http://www.example.com/thispage`
Wildcard Expression	Allows all matching URLs. Use the wildcard * character to replace one or more characters.	`http://*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 format `scheme://host[:port]`, with no trailing slashes (`:port` is optional). So `http://example.com/` would not allow CORS requests from the domain `example.com`.

Disabling the whitelist

The whitelist is enabled by default. You can choose to disable the whitelist however this will allow all URLs, including malicious content, and is not recommended.

To disable the whitelist:

1. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
2. Choose **Whitelist**.
3. Choose **Turn off whitelist**.
4. Choose **Confirm**.

All URLs will now be allowed. We strongly recommend **not** disabling the whitelist.

Configuring the Time Interval at which Drafts are Saved

By default, Confluence saves a draft of your page once every thirty seconds. Confluence administrators can configure how often drafts are saved.

As a Confluence administrator, you can set the time interval at which drafts are saved as follows:

1. Choose the **cog icon**
   , then choose **General Configuration** under Confluence Administration
2. Click **Further Configuration** in the left-hand panel.
3. Edit the setting for **Draft Save Interval**.

Related pages:

- Drafts

Performance Tuning

This document describes tuning your application for improved performance. It is not a guide to troubleshooting Confluence outages. Check Troubleshooting Confluence hanging or crashing for help if Confluence is crashing.

Like any server application, Confluence may require

some tuning as it is put under heavier use. We do our best to make sure Confluence performs well under a wide variety of circumstances, but there's no single configuration that is best for everyone's environment and usage patterns.

If you are having problems with the performance of Confluence and need our help resolving them, you should read Requesting Performance Support.

Use the latest version of your tools

Use the latest versions of your application servers and Java runtime environments. Newer versions are usually better optimized for performance.

Avoid swapping due to not enough RAM

Always watch the swapping activity of your server. If there is not enough RAM available, your server may start swapping out some of Confluence's heap data to your hard disk. This will slow down the JVM's garbage collection considerably and affect Confluence's performance. In clustered installations, swapping can lead to a Cluster Panic due to Performance Problems. This is because swapping causes the JVM to pause during Garbage Collection, which in turn can break the inter-node communication required to keep the clustered nodes in sync.

Being aware of other systems using the same infrastructure

It may sound tempting: Just have one powerful server hosting your database and/or application server, and run **all** your crucial programs on that server. **If** the system is set up perfectly, then you might be fine. Chances are however that you are missing something, and then one application's bug might start affecting other applications. So if Confluence is slow every day around noon, then maybe this is because another application is using the shared database to generate complicated reports at that time? Either make sure applications can't harm each other despite sharing the same infrastructure, or get these systems untangled, for example by moving them to separate instances that can be controlled better.

Choice of database

The **embedded H2 database** is provided for evaluating Confluence, not for production Confluence sites. After the evaluation finishes, you must switch to a supported external database. We recommend using what you are familiar with, because your ability to maintain the database will probably make far more difference to what you get out of it than the choice of database itself.

Database connection pool

If load on Confluence is high, you may need more simultaneous connections to the database.

- If you are using JNDI data-sources, you will do this in your application server's configuration files.
- If you have configured Confluence to access the database directly, you will need to manually edit the hibernate.c3p0.max_size property in the confluence.cfg.xml file in your confluence.home directory.

After you have changed the URL in this file, restart Confluence.

To assess whether you need to tune your database connection pool, take thread dumps during different times (including peak usage). Inspect how many threads have concurrent database connections.

Database in general

If Confluence is running slowly, one of the most likely cause is that there is some kind of bottleneck in (or around) the database.

The first item you should check is the **"Database Latency"** field in the System Information tab in the admin console.

Database Connection Transaction Isolation	Read committed
Database Latency	0
	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 analysers

Modern databases have query optimisers based on collecting statistics on the current data. Using the SQL EXPLAIN statement will provide you information on how well the query optimiser is performing. If the cost estimate is wildly inaccurate then you will need to run statistics collection on the database. The exact command will depend on your database and version. In most cases you can run statistics collection while Confluence is running, but due to the increased load on the database it's best to do this after normal hours or on a week-end.

Cache tuning in Confluence and Apache

To reduce the load on the database, and speed up many operations, Confluence keeps its own cache of data. Tuning the size of this cache may speed up Confluence (if the caches are too small), or reduce memory (if the caches are too big).

Please have a look at our documentation on Cache Performance Tuning for information on how to tune Confluence caches.

To improve performance of a large Confluence site, we recommend that you move the caching of static content from the JVM into Apache. This will prevent the JVM from having a number of long running threads serving up static content. See Configuring Apache to Cache Static Content via mod_disk_cache.

Antivirus software

Antivirus software greatly decreases the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental, and may even cause errors with Confluence. You should configure your antivirus software to ignore the Confluence home directory, its index directory and any database-related directories.

Enabling HTTP compression

If bandwidth is responsible for bottlenecking in your Confluence installation, you should consider enabling HTTP compression. This may also be useful when running an external facing instance to reduce your bandwidth costs.

Take note of the known issues with HTTP compression in versions of Confluence prior to 2.8, which may result in high memory consumption.

Performance testing

You should try out all configuration changes on a demo system. Ideally, you should run and customize loadtests that simulate user behaviour.

Access logs

You can find out which pages are slow and which users are accessing them by enabling Confluence's built-in access logging.

Built-in profiler

You can identify the cause of page delays using Confluence's built-in profiler according to Troubleshooting Slow Performance Using Page Request Profiling.

Application server memory settings

See How to fix out of memory errors by increasing available memory.

Web server configuration

For high-load environments, performance can be improved by using a web server such as Apache in front of the application server. There is a configuration guide to Running Confluence behind Apache.

When configuring your new web server, make sure you configure sufficient threads/processes to handle the load. This applies to both the web server and the application server connector, which are typically configured separately. If possible, you should enable connection pooling in your web server connections to the application server.

Troubleshooting possible memory leaks

Some external plugins, usually ones that have been written a long time ago and that are not actively maintained anymore, have been reported to consume memory and never return it. Ultimately this can lead to a crash, but first this manifests as reduced performance. The Troubleshooting Confluence hanging or crashing guide is a good place to start. Some of the known causes listed there could result in performance issues short of a crash or hang.

Cache Performance Tuning

Confluence performance can be significantly affected by the performance of its caches. It is essential for the administrator of a large production installation of Confluence to tune the caches to suit its environment.

There are several configurable parameters for each of the cache regions, most notably cache size, cache expiry delay and eviction policy. In most cases, cache size is the parameter you would want to change.

To change the size of a cache:

1. Go to

On this page:

- Cache tuning example
- Finding the configuration file
- Cache key mappings
- Caching in Confluence Data Center
- Important caches
- Monitoring the contents of a cache
- Notes

 > **General Configuration** > **Cache Management**.

2. Choose **Show Advanced View**.
3. Choose **Adjust Size** next to the cache you
   want to change.

To modify other parameters you can modify the cache configuration files manually.

Cache tuning example

As an example of how to tune Confluence's caches, let's have a look at the following table:

Caches	% Used	% Effectiveness	Objects/Size	Hit/Miss/Expiry
Attachments	87%	29%	874/1000	78226/189715/187530
Content Attachments	29%	9%	292/1000	4289/41012/20569
Content Bodies	98%	81%	987/1000	28717/6671/5522
Content Label Mappings	29%	20%	294/1000	4693/18185/9150
Database Queries	96%	54%	968/1000	105949/86889/83334
Object Properties	27%	18%	279/1000	5746/25386/8102
Page Comments	26%	11%	261/1000	2304/17178/8606
Users	98%	5%	982/1000	6561/115330/114279

The maximum size of the caches above is 1000 (meaning that it can contain up to 1000 objects). You can tell when a cache size needs to be increased because the cache has both:

- a high usage percentage (above 75%)
- a low effectiveness percentage.

Check the 'effectiveness' versus the 'percent used'. A cache with a low percent used need not have its size lowered; it does not use more memory until the cache is filled.

Based on this, the sizes of the "Attachments", "Database Queries", and "Users" caches should be increased to improve their effectiveness.

As the stored information gets older or unused it will expire and be eliminated from the cache. Cache expiry may be based on time or on frequency of use.

ℹ There is not much that you can do with a cache that has both a low percentage of usage and effectiveness. Over time, as the cache is populated with more objects and repeat requests for them are made, the cache's effectiveness will increase.

Finding the configuration file

Cache configurations are stored in `<confluence-home>/shared-home/config/cache-settings-ov errides.properties`

For Confluence Data Center (clustered) it can be found in `<confluence-shared-home>/config/cache -settings-overrides.properties` (in the shared home directory for the cluster).

Cache key mappings

The cache configuration file configures caches by their keys. To find out a cache key hover your mouse over the cache name in the **Cache Management** screen.

Caching in Confluence Data Center

In Confluence Data Center (clustered) you have a distributed cache and a cluster node-local cache. The Cluster Management page will indicate cluster distributed cache and cluster node-local cache.

The cache configuration file is stored in the shared home directory for the cluster.

Important caches

> The following suggestions are general guidelines. In cases of large databases, 20-30% of the size of the table may be unnecessarily large. Check the effectiveness and percent used categories in the cache for more specific assessments.

- **Content Objects cache** (`com.atlassian.confluence.core.ContentEntityObject`) should be set to at least 20-30% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the query `select count(*) from CONTENT where prevver is null`.
- **Content Body Mappings cache** (`com.atlassian.confluence.core.ContentEntityObject.bodyContents`) should be set to at least 20% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the query `select count(*) from CONTENT where prevver is null`.
- **Embedded Crowd Internal User cache** (`com.atlassian.crowd.model.user.InternalUser`) should be set to the number of users you have in the internal directory. You can discover this number by using the following SQL:

```
SELECT
    COUNT(*)
FROM
    cwd_user u
JOIN
    cwd_directory d
ON
    u.directory_id = d.id
AND d.directory_name = 'Confluence Internal Directory';
```

- **Embedded Crowd Users cache** `com.atlassian.confluence.user.crowd.CachedCrowdUserDao.USER_CACHE` should be set to the number of rows in the cwd_user table.

```
SELECT
  COUNT(*)
FROM
  cwd_user u;
```

- **Space permissions by ID cache** (`com.atlassian.confluence.security.SpacePermission`) should be set to the number of space permissions in your deployment (a good rule of thumb is 20 times the number of spaces). You can find the number of space permissions using the query `select count(*) from SPACEPERMISSIONS`.

Monitoring the contents of a cache

To monitor what is in a cache:

1. In your browser go to `<confluence-URL>/admin/cachecontents.jsp`
   All caches that contain items will appear.
2. Select a cache from the list.

Notes

- To improve performance of a large Confluence site, we recommend that you move the caching of static content from the JVM into Apache. This will prevent the JVM from having a number of long running threads serving up static content. See Configuring Apache to Cache Static Content via mod_disk_cache.

Cache Statistics

Confluence provides statistics about its internal caches that allow you to track the size and hit ratio of each cache and tune it for better performance (if necessary). See Performance Tuning for more information.

Configurable Caches

System administrators can change the sizes of Confluence's internal caches through the Administration Console and these changes will take effect without the need to first shut down and then restart Confluence. The maximum number of units for any of the defined cache regions can be adjusted individually.

Note that larger cache sizes will require more memory at runtime, so you should review the memory allocation of the Confluence Java process and the physical memory available on your server.

Viewing Cache Statistics and Modifying Cache Sizes

To view the cache statistics:

1. Go to

   ⚙ ▾

   > **General Configuration** > **Cache Management**.
2. Choose **Show Advanced View**.

Here is an example for one of the most frequently used caches, the 'Content Object' cache.

Cache Statistics

This page displays statistics about Confluence's internal caches. You may wish to tune cache sizes and expiration policies to improve memory usage or performance.

Cache Name	Capacity Utilisation	Effectiveness	Current / Max Entries	Current Heap Size (MB)	Hit / Miss / Evicted	Adjust Max Entries	Flush Entries
Content Objects	36%	73%	3653 / 10000	Unknown	307748 / 112538 / 37051	Adjust Size	Flush

About the generated numbers:

Capacity Utilisation	=(Objects)/(Size) 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.

Changes to cache size configurations are saved in the `config/cache-settings-overrides.properties` file in your home directory (or shared home directory in a Confluence Data Center cluster).

Memory Usage and Requirements

Managing Confluence's performance and memory usage really depends on what resources are available. Confluence will run faster if you give it lots of memory for its caches, but it should still be able to run quite well in low-memory environments, with the right tuning. Below are some tips on getting the most out of your Confluence site.

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

Increasing the amount of memory available to Confluence

See Increasing JIRA Memory for details on how to increase the memory available to web application servers typically used to run Confluence.

Embedded database

The embedded HSQL database that comes with Confluence essentially holds all your data in memory while the Confluence server is running. If you are running out of memory, you should consider migrating Confluence to an external database.

Caching

By default, Confluence keeps large in-memory caches of data to improve its responsiveness and the user experience. The trade off is an increase in memory requirements to support the cache. Administrators of larger Confluence sites may need to configure the size of their caches to improve performance.

To customise Confluence's cache to meet your needs, see cache tuning.
To increase the amount of memory available to Confluence, see How to fix out of memory errors by increasing available memory.

Mail error queue

Confluence keeps a copy of all emails that it failed to send within an internal error queue. In the event of intermittent failures such as network connectivity issues, the emails in this queue can be manually resent when the problem is fixed. Under certain circumstances, the mail queue can fill up with large objects. The queue is regularly flushed, but if you get a *lot* of mail errors, you might get a spike in memory usage.

Attachments

The indexing of large attachments requires that the attachment be loaded into memory. In the case of large attachments, this can cause a temporary strain on the systems resources, and may result in indexing failing because the attachment could not be fully loaded into memory.

System backup and restore

The Confluence backup and restore process scales linearly with the size of data. This can have a significant impact on large Confluence instances where the amount of data exceeds the amount of available memory. If you are experiencing an `OutOfMemoryError` during either a backup or restore processes, then we strongly recommend that you choose and Production Backup Strategy.

If you encounter an `OutOfMemoryError` while restoring a backup and wish to overcome this issue by increasing memory, how much more will you need to make this process work? A good rule of thumb is to have a look at the size of the `entities.xml` file in your backup. This file contains all of the data Confluence will be loading, so at least that much is required. Add another 64-128Mb to ensure that Confluence has enough memory to load and function and that should be enough. To increase the amount of memory available to Confluence, see How to fix out of memory errors by increasing available memory.

Known issues that we do not have control over

There are also some memory issues we don't have any control over. For example,

- There's a memory leak in the Oracle 10g JDBC drivers. Not much we can do about that.
- One customer found a rather nasty memory leak that appeared to originate inside Tomcat 5, but only using the IBM JDK on PowerPC.

If you are having problems that appear to result from a memory leak, log an issue on http://support.atlassian.com. Our memory profiler of choice is YourKit. It would be helpful to us if you can provide us with a memory dump from that tool showing the leak.

Confluence is taking long periods of time to respond to some actions

A common cause of random pauses in Confluence is the JVM running garbage collection. To determine if this is what is happening, enable verbose garbage collection and look at how long Java is taking to free up memory. If the random pauses match when Java is running its garbage collection, garbage collection is the cause of the pause.

Verbose garbage collection will generate log statements that indicate when Java is collecting garbage, how long it takes, and how much memory has been freed.

To enable gc (garbage collection) logging, start Confluence with the option `-XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc -Xloggc:gc.log`. Replace `gc.log` with an absolute path to a `gc.log` file.

For example, with a Windows service, run:

```
tomcat5 //US//Confluence ++JvmOptions="-XX:+PrintGCDetails
-XX:+PrintGCTimeStamps -verbose:gc -Xloggc:c:\confluence\logs\gc.log"
```

or in `bin/setenv.sh`, set:

```
export CATALINA_OPTS="$CATALINA_OPTS -XX:+PrintGCDetails
-XX:+PrintGCTimeStamps -verbose:gc
-Xloggc:${CATALINA_BASE}/logs/gc.log"
```

If you modify `bin/setenv.sh`, you will need to restart Confluence for the changes to take effect.

What can you do to minimise the time taken to handle the garbage collection? See http://java.sun.com/docs/hotspot/gc1.4.2/ for details on tuning the JVM to minimise the impact that garbage collection has on the running application.

Requesting Performance Support

Basic performance troubleshooting steps

Begin with the following procedures:

Requesting basic performance support

If the above tips don't help or you're not sure where to start, open a support ticket starting with at least the basic information:

1. The atlassian-confluence.log
2. The `catalina.out` log (or your application server log), with a series of three thread dumps separated by 10 seconds
3. A description with as much detail as possible regarding:
   a. What changes have been made to the system?
   b. When did performance problems begin?
   c. When in the day do performance issues occur?
   d. What pages or operations experience performance issues?
   e. Is there a pattern?

Continue with as much of the advanced performance troubleshooting information as you can.

Advanced performance troubleshooting

Please gather **all** of the information listed below and include it in your support request, even if you think you have a good idea what's causing the problem. That way we don't have to ask for it later.

System information

Confluence server

- Take a screenshot of Confluence's `Administration  System Information` (or save the page as HTML)
- Take a screenshot of Confluence's `Administration  Cache Statistics` (or save the page as HTML)
- Find out the exact hardware Confluence is running on
  - How many CPUs? What make and model? What MHz?
  - How much memory is installed on the machine?
  - How much memory is assigned to Confluence's JVM? (i.e. what are the -Xmx and -Xms settings for the JVM?)
  - What other applications are being hosted on the same box?

Confluence content

- How many *users* are registered in Confluence?
- On average, to how many *groups* does each user belong?
- How many *spaces* (global and personal) are there in your Confluence server?
- How many of those spaces would be viewable by the average user?
- Approximately how many *pages*? (Connect to your database and perform `'select count(*) from content where prevver is null and contenttype = 'PAGE''`)
- How much data is being stored in Bandana (where plugins usually store data)? (Connect to your database and perform `'select count(*), sum(length(bandanavalue)) from bandana'`)

The database

- What is the exact version number of Confluence's database server?
- What is the exact version number of the JDBC drivers being used to access it? (For some databases,

the full filename of the driver JAR file will suffice)
- Is the database being hosted on the same server as Confluence?
- If it is on a different server, what is the network latency between Confluence and the database?
- What are the database connection details? How big is the connection pool? If you are using the standard configuration this information will be in your confluence_cfg.xml file. Collect this file. If you are using a Data source this information will be stored in your application server's configuration file, collect this data.

User management

- Are you using external user management or authentication? (i.e. JIRA or LDAP user delegation, or single sign-on)
- If you are using external JIRA user management, what is the latency between Confluence and JIRA's database server?
- If you are using LDAP user management:
  - What version of which LDAP server are you using?
  - What is the latency between Confluence and the LDAP server?

Diagnostics

Observed problems

- Which pages are slow to load?
  - If it is a specific wiki page, attach the wiki source-code for that page
- Are they always slow to load, or is the slowness intermittent?

Monitoring data

Before drilling down into individual problems, helps a lot to understand the nature of the performance problem. Do we deal with sudden spikes of load, or is it a slowly growing load, or maybe a load that follows a certain pattern (daily, weekly, maybe even monthly) that only on certain occasions exceeds critical thresholds? It helps a lot to have access to continuous monitoring data available to get a rough overview.

Here are sample graphs from the confluence.atlassian.com system, showing

Load
This graph shows the load for two consecutive days. The obvious pattern is that the machine is under decent load, which corresponds to the user activity, and there is no major problem.

Resin threads and database connections

Active number of Java Threads

These two charts show the active threads in the application server (first chart) and the size database connection pool (second chart). As you can see, there was a sudden spike of server threads and a corresponding spike of db-connections.

The database connection pool size

The database connection pool size peaked over 112, which happened to be more than the maximum number of connections the database was configured for (100). So it was no surprise that some requests to Confluence failed and many users thought it had crashed, since many requests could not obtain the crucial database connections.

We were able to identify this configuration problem quite easily just by looking at those charts. The next spikes were uncritical because more database connections were enabled.

The bottom line being: it helps a lot to monitor your Confluence systems continuously (we use Hyperic, for example), and it helps even more if you are able to send us graphs when you encounter problems.

Access logs

- How to Enable User Access Logging, including redirecting the logs to a separate file
  - You can run this file through a log file analyser such as AWStats, or manually look through for pages which are slow to load.

Profiling and logs

- Enable Confluence's built-in profiling for long enough to demonstrate the performance problem using 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, in particular with the Top threads plugin.

Site metrics and scripts

- It is essential to understand the user access and usage of your instance. Please use the access log scripts and sql scripts to generate Usage statistics for your instance.

Next step

Open a ticket on https://support.atlassian.com and attach all the data you have collected. This should give us the information we need to track down the source of your performance problems and suggest a solution. Please follow the progress of your enquiry on the support ticket you have created.

Access Log Scripts

The access log scripts are attached to this page. To use the scripts:

1. Unzip the 7z file.
2. Copy all the daily access logs to a folder called `logs`.
3. Run `Atlassian-processDailyLog.rb`. This will generate a `csv` file called `summary.csv` and several directories which contain the access logs of each defined user action.
4. Run the appropriate script `Atlassian-processDailyLog-hourly.rb` `<admin/comment/create/edit/search/rss>`. Each script will generate a different `csv` file. For example, `Atlassian-processDailyLog-hourly.rb admin` will process the admin logs extracted in step 3.
5. Import the `csv` files to `www-log-Analysis.xls` (`summary.csv` to 'raw stats - daily' sheet and `admin.csv` to 'admin -hours' sheet, etc) to generate the load profiles and graphs. You may need to modify the number of rows in each sheet depending on the number of logs.

> **Note**
> All scripts are written in Ruby and assume the log file name contains the string 'confluence.atlassian.com-access.log'. Scripts need to be changed if another name is used. Modify the line: `filenameRegexp = Regexp.new('confluence.atlassian.com-access.log')`

Troubleshooting Slow Performance Using Page Request Profiling

This page tells you how to enable page-request profiling. With profiling turned on, you will see a record of the time it takes (in milliseconds) to complete each action made on any Confluence page. If Confluence is responding slowly, an internal timing trace of the slow page request can help to identify the cause of the delay.

You will need access to the Confluence server to view a profile.

Enabling Page-Request Profiling

On this page:

- Enabling Page-Request Profiling
- Profiling an Activity
- Example of a Profile
- Start Confluence with Profiling Enabled

Related pages:

- Requesting Performance Support
- Working with Confluence Logs

> To see just the slow performing macros, see Identifying Slow Performing Macros.

ℹ You need to have System Administrator permissions in order to enable or disable profiling.

To enable page profiling:

1. Choose the **cog icon**

    ⚙ ▾

    , then choose **General Configuration** under Confluence Administration
2. Choose '**Logging and Profiling**' in the left-hand panel.
3. The '**Logging and Profiling**' screen appears. Choose '**Enable Profiling**'.
    ℹ If profiling is already enabled, the button will be labelled 'Disable Profiling'.

To disable page profiling:

1. Choose the **cog icon**

    ⚙ ▾

    , then choose **General Configuration** under Confluence Administration
2. Choose '**Logging and Profiling**' in the left-hand panel.
3. The '**Logging and Profiling**' screen appears. Choose '**Disable Profiling**'.
    ℹ If profiling is already disabled, the button will be labelled 'Enable Profiling'.

Screenshot: Changing Log Levels and Profiling

Logging and Profiling

Performance Profiling
Profiling is currently OFF.
[Enable Profiling]

SQL Logging
[Enable SQL Logging]

Log4j Logging

loglevel.profileDesc
[Production] [Diagnostic]

OR:

Customise specific logging settings

Add New Entry

Class/Package Name	New Level
	INFO ▾ [Add entry]

Existing Levels

Class/Package Name	Current Level	New Level	
atlassian.plugin	INFO	INFO ▾	Remove
com.atlassian.confluence.admin.actions	INFO	INFO ▾	Remove

Profiling an Activity

1. Enable profiling, using either of the methods described above.
    Profiles for every page hit, for all users, will now be logged to your application server's default logs until Confluence is restarted. Note that each time a user visits a link, a single profile is printed.

2. Confirm that profiles are being written to the Confluence log file — see Working with Confluence Logs for location of the log files and other details.
3. Perform the activity that is resulting in unusually slow response time.
4. Copy the profile for that action. When deciding which profiles to copy, look for the links that took a long time to respond. If a single page is slow, only that profile is necessary. If Confluence is generally or intermittently slow, copy all profiles logged during the slowdown until a reasonable sample has been collected.
5. If you were instructed to profile your instance by Atlassian technical support, attach all relevant profiles to your support ticket.
6. Turn profiling off again, using either of the methods described above.
7. Confirm that profiles are no longer being printed to the Confluence log file.

Example of a Profile

Below are the first few lines of a normal profile for accessing a page called Confluence Overview.

```
[344ms] - /display/ds/Confluence+Overview
  [313ms] - SiteMesh: parsePage:
http://localhost:8080/display/ds/Confluence+Overview
    [313ms] - XW Interceptor: Before defaultStack:
/pages/viewpage.action (ViewPageAction.execute())
      [0ms] - SpaceAwareInterceptor.intercept()
      [16ms] - PageAwareInterceptor.intercept()
        [0ms] - AOP: PageManager.getPage()
        [16ms] - AOP: PermissionManager.hasPermission()
          [0ms] - AOP: SpacePermissionManager.hasPermission()
          [16ms] - AOP: SpacePermissionManager.hasPermission()
        [0ms] - AOP: SpacePermissionManager.hasPermission()
      [0ms] - AOP: SpacePermissionManager.hasPermission()
      [281ms] - XW Interceptor: After defaultStack:
/pages/viewpage.action (ViewPageAction.execute())
        [281ms] - XW Interceptor: After validatingStack:
/pages/viewpage.action (ViewPageAction.execute())
            ...
```

ℹ Notice that each indented line is a recursive call that rolls up into the parent line. In the example above, the Confluence Overview page takes 344ms. Part of that, 313ms, is spent in sitemesh.

Start Confluence with Profiling Enabled

There may be some situations where you may wish to have Confluence profiling enabled during startup. This may be useful if you restart often and may forget to enable profiling for Support/Trouble-shooting purposes.

Edit the file *CONFLUENCE_HOME\confluence\WEB-INF\web.xml*. You should see a stanza similar to the one below. Set the parameter value for **autostart** to **true**:

```
<filter>
        <filter-name>profiling</filter-name>

<filter-class>com.atlassian.core.filters.ProfilingAndErrorFilter</fil
ter-class>
        <init-param>
            <!-- specify the which HTTP parameter to use to turn the
filter on or off -->
            <!-- if not specified - defaults to "profile.filter" -->
            <param-name>activate.param</param-name>
            <param-value>profile</param-value>
        </init-param>
        <init-param>
            <!-- specify the whether to start the filter
automatically -->
            <!-- if not specified - defaults to "true" -->
            <param-name>autostart</param-name>
            <param-value>true</param-value>
        </init-param>
    </filter>
```
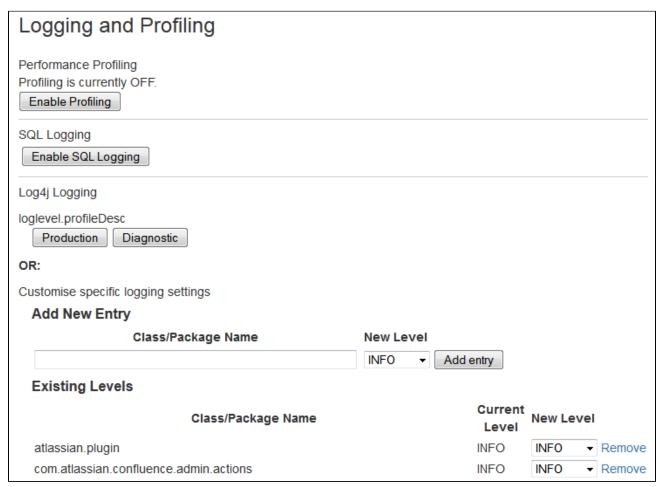
Remember to turn it back to **false** or your logs will grow very large.

Identifying Slow Performing Macros

Page Profiling gives good detail on what operations are slow in a page load. In addition, you can add debug level logging:
Version 3.1 and Later

Set the package name `com.atlassian.renderer.v2.components.MacroRendererComponent` to DEBUG in Administration >> Logging and Profiling.
Prior to version 3.1

Download WikiMarkupParser.class, available from the attachments to this page. This will result in logs like:

```
2009-04-23 10:27:54,789 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Entering macro rendering. Processed
text: {spaces}
2009-04-23 10:27:55,768 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Exiting macro text rendering. Total
time: 979ms
2009-04-23 10:27:55,785 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Entering macro rendering. Processed
text: {create-space-button}
2009-04-23 10:27:55,857 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Exiting macro text rendering. Total
time: 72ms
2009-04-23 10:27:55,862 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Entering macro rendering. Processed
text: {recently-updated-dashboard:dashboard|showProfilePic=true}
2009-04-23 10:27:56,704 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Exiting macro text rendering. Total
time: 842ms
2009-04-23 10:27:56,707 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Entering macro rendering. Processed
text: {favpages:maxResults=10}
2009-04-23 10:27:56,889 DEBUG [http-8080-1]
[atlassian.renderer.v2.WikiMarkupParser] parse Exiting macro text rendering. Total
time: 182ms
```

To add the class:

1. Add this line to the file <confluence-install>/confluence/WEB-INF/classes/log4j.properties:
   log4j.logger.com.atlassian.renderer=DEBUG
2. Add the appropriate WikiMarkupParser.class to /confluence/WEB-INF/classes/com/atlassian/renderer/v2.
   You'll have to make the renderer and v2 folders.

In combination with page profiling, this should give good specifics on the amount of time various plugins take. You can also use this utility to Search Confluence for Uses of a Macro.

Resolution

Experiment with the tips from the performance tuning page, or open an enhancement request about the specific macro. In some instances there is no resolution - you'll just be aware of the overhead of various macros.

Compressing an HTTP Response within Confluence

Confluence supports HTTP GZip transfer encoding. This means that Confluence will compress the data it sends to the user, which can speed up Confluence over slow or congested Internet links, and reduce the amount of bandwidth consumed by a Confluence server.

Turn on Confluence's GZip encoding if:

- Users are accessing Confluence over the Internet, or a WAN connection with limited bandwidth.
- You wish to reduce the amount of data transfer between the Confluence server and client.

If you are accessing Confluence over a Local Area Network or over a particularly fast WAN, you may wish to leave GZip encoding disabled. If the network is fast enough that transferring data from Confluence to the user isn't a limiting factor, the additional CPU load caused by compressing each HTTP response may slow Confluence down.

Enabling HTTP Compression

1. Choose the **cog icon**

, then choose **General Configuration** under Confluence Administration
2. Select '**General Configuration**' in the left-hand panel.
3. Enable '**Compress HTTP Responses**'.

It is possible to configure which types of content are compressed within Confluence. By default, the following mime types will be compressed:

- text/htmltext
- javascript
- text/css
- text/plain
- application/x-javascript
- application/javascript

If you wish to change the types of content to be compressed, add a replacement `urlrewrite-gzip-default .xml` file within the `WEB-INF/classes/com/atlassian/gzipfilter/` directory in your Confluence Installation Directory. A sample file is provided as an attachment. It is unlikely that you will need to alter this file.

Garbage Collector Performance Issues

> This document relates broadly to memory management with Oracle's Hotspot JVM. These recommendations are based on Support's successful experiences with customers with large Confluence instances.

> Please **do not** use the Concurrent Mark Sweep (CMS) Collector with Confluence, unless otherwise advised by Atlassian Support. It requires extensive manual tuning and testing, and is likely to result in degraded performance.

Use a small heap

Keep your heap as small as possible, without the instance experiencing OutOfMemory errors. If you experience OutOfMemory errors and need to increase this, we recommend you do it in 512mb or 1gb allotments, and monitor the instance. If you continue to receive OutOfMemory errors, increase the heap by another 512mb or 1gb, and continue this process until you are operating stably with no OutOfMemory errors. Do not increase the heap further than required, as this will result in longer garbage collections.

Remove any old tuning parameters

On every full GC, the JVM will resize the allocations of Eden, Survivor etc based on the throughput it is actually seeing. It will tune itself based on the real world data of the objects that are being created and collected. Most of the time simply allowing JVM to tune itself will give you better performance.

If you have added JVM parameters in the past and are experiencing difficulties with GC now, we'd recommend you remove all GC related parameters, unless you added them to solve a specific problem, and they did in fact solve that problem. You should also consider re-benchmarking now to ensure that they are still solving that problem, and are not causing you any other issues.

Check your VM resources

If you run Confluence on a VM, check that 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, we recommend following the instructions in our Garbage Collection (GC) Tuning Guide. This document will take you through the process of choosing performance goals (throughput/footprint/latency), and how to tune for those goals.

Viewing your GC logs

How to Enable Garbage Collection (GC) Logging, and use a tool like Chewiebug's GCViewer to view the resulting logs.

Cluster Troubleshooting

This page covers troubleshooting for a clustered installation of Confluence.

<table>
<tr><td>On this page:</td></tr>
<tr><td>

• Symptoms
• Didn't find
  a solution?

</td></tr>
</table>

- For information about clustering in general, refer to Confluence Enterprise Resources.
- If you're experiencing Cluster Panic messages in non-clustered installation of Confluence, visit the Knowledge Base article 'Database is being updated by an instance which is not part of the current cluster' Error Message.

⚠ You must ensure the clocks on your cluster nodes don't diverge, as it can result in a range of problems with your cluster.

Symptoms

Below is a list of potential problems with a Confluence cluster, and their likely solutions. The solutions are listed below.

Problem	Likely solutions
`Database is being updated by an instance which is not part of the current cluster` errors on a stand-alone	'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
Add-on is unlicensed on some nodes after updating the license on one node.	Disable and re-enable the add-on in the Universal Plugin Manager.
After a add-on update, strings appear in the UI instead of buttons and icons on some nodes.	Restart the affected node.
Any issue not covered here	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.xx.xxx.xxx</property>
```

A warning message is displayed when an user changes the address from the one that Confluence has generated by the hashing of the name. There is no way of eliminating the message any other way other than by returning the address to the one that matches the cluster name. Purpose of the warning message is to remind the user that the address has been changed - as it is not the hashed version any longer - consequently the node can not join the cluster just by using the name. It is also necessary to provide the correct address as well.

Mapping interface to IP address.

To ensure that the interface name is mapped correctly, the following tool can be used. It shows the mapping of the interface name to the IP address.

```
C:\>java -jar list-interfaces.jar
interfaces.size() = 4
networkInterface[0] = name:lo (MS TCP Loopback interface) index: 1
addresses:
/127.0.0.1;

networkInterface[1] = name:eth0 (VMware Virtual Ethernet Adapter for
VMnet8) index: 2 addresses:
/192.168.133.1;

networkInterface[2] = name:eth1 (VMware Virtual Ethernet Adapter for
VMnet1) index: 3 addresses:
/192.168.68.1;

networkInterface[3] = name:eth2 (Broadcom NetXtreme 57xx Gigabit
Controller - Packet Scheduler Miniport) index: 4 addresses:
/192.168.0.101;
```

Debugging tools

Listed below are some debugging tools that help determine what the status of the multicast traffic is:

Tool	Information provided
`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 `confl uence.cluster.interface` property in `<local-home>/confluence.cfg.xml` and specify the network interface. For example to tell Confluence to use `eth1`:

```
<property name="confluence.cluster.interface">eth1</property>
```

Overriding Hazelcast Configuration

If the solution to your problem involves changes to the Hazelcast configuration, these changes should **not** be made to the Confluence configuration files. Instead, to ensure your configuration survives upgrades, make your changes by creating a Hazelcast override file.

Increase multicast TTL

The multicast time-to-live (TTL) specifies how many *hops* a multicast packet should be allowed to travel before it is discarded by a router. It should be set to the number of routers in between your clustered nodes: 0 if both are on the same machine, 1 if on two different machines linked by a switch or cable, 2 if on two different machines with one intermediate router, and so on.

To increase the multicast TTL by edit the `confluence.cluster.ttl` property in the `<local home>/confluence.cfg.xml` file on each node. For example to set the TTL to 3:

```
<property name="confluence.cluster.ttl">3</property>
```

Check intermediate routers

Advanced switches and routers have the ability to understand multicast traffic, and route it appropriately. Unfortunately sometimes this functionality doesn't work correctly with the multicast management information (IGMP) published by the operating system running Confluence.

If multicast traffic is problematic, try disabling advanced multicast features on switches and routers in between the clustered nodes. These features can prevent multicast traffic being transmitted by certain operating systems.

Didn't find a solution?

Check Related Articles from the Confluence Knowledge Base

- Recovering from a Data Center cluster split-brain
- Port [5801] is already in use and auto-increment is disabled. Hazelcast cannot start.
- "Exception bootstrapping cluster:Shared home directory is not configured correctly" Error during Confluence Data Center startup
- Multicast communication works only one-way
- Cannot find "external_id" column when trying to upgrade to a Confluence CDC license after upgrading from a pre-5.5 Confluence Clustered installation
- Hazelcast CANNOT start on this node. No matching network interface found.
- Cluster Panic due to Multicast Traffic Communication Problem
- Configuration of Confluence Cluster Fails with 'Cannot assign requested address'
- How to suppress cluster warning messages in the Confluence log files

Contact Atlassian support

We have dedicated staff on hand to support your installation of Confluence. Please follow the instructions for raising a support request and mention that you're having trouble setting up your Confluence cluster.

Data Collection Policy

Why does Confluence collect usage data?

We're proud that Confluence is one of the most versatile collaboration tools on the planet, and we will continue to deliver innovative new features as quickly as we can. In order to prioritize the features we deliver, we need to understand how our customers use Confluence, what's important, what's not, and what doesn't work well. The collection of usage data allows us to measure the user experience across many thousands of users and deliver features that matter.

What data is collected?

The type of data we collect is covered in our Privacy Policy. Please read it - we've tried to avoid legal jargon and made it as straightforward as possible.

To view a sample of data that might be collected from your specific installation, go to

> **General Configuration** > **Analytics**.

Data is always collected in Confluence Cloud.

How is data collected from Confluence?

Older versions of Confluence (prior to Confluence 5.6 or Confluence Questions 1.0.618) didn't collect usage data. Analytics are collected using the Atlassian Analytics add-on. The add-on collects analytics events in a log file which is located in `<confluence-home>/analytics-logs`. The logs are periodically uploaded using an encrypted session and then deleted. If Confluence is unable to connect to the Internet, no logs are ever uploaded.

Enabling/disabling data collection in Confluence

You can turn off analytics collection at any time. Go to

> **General Configuration** > **Analytics**.

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.

- System Requirements
  - Server Hardware Requirements Guide
  - Example Size and Hardware Specifications From Customer Survey
  - Running Confluence in a Virtualised Environment
- Confluence Installation Guide
  - Installing Confluence
  - Installing 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
- Supported Platforms
  - End of Support Announcements for Confluence
  - Supported Platforms FAQ
- Migrating Confluence Between Servers
  - From Confluence Evaluation through to Production Installation
  - Migrate from Confluence Cloud to Server

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 works with a broad range of operating systems, database systems and application servers. Provided you have the technical knowledge, it is very likely that you will be able to run Confluence with an 8-year-old database or even on some 8-year-old hardware. Realistically, it is not technically feasible for us to provide our legendary support service on all environments available. There can only be a finite number of platforms and release versions of those that we support.

Our rule of thumb when releasing a new version of Confluence is that we will officially support platforms that have been released within the last one to two years (or the latest version of that platform if no new

On this page:
- Confluence Software Requirements
  - Operating Systems
  - Application Servers
  - Databases
  - Java
  - Please Note: Impact of Antivirus Software
- Confluence Hardware Requirements
- Hosted Solutions – Confluence Cloud

version of it was released in that period). This does not necessarily mean that you will need to upgrade your database or application server every time you upgrade Confluence. However, if you do run into problems with an unsupported version of a database or application server, we may have to ask you to upgrade to something newer.

Please refer to our Supported Platforms topic for details on platforms that we currently support in this version of Confluence and our Supported Platforms FAQ topic for details on our support handling procedures.

Confluence Software Requirements

Please read the Supported Platforms page for Confluence. That page contains a list of specific software that Confluence will work with.

Operating Systems

Atlassian supports the operating systems listed on the Supported Platforms page.

If you would like to run Confluence on virtualised hardware, please read our Running Confluence in a Virtualised Environment document first.

Application Servers

An application server is required to run Confluence. Apache Tomcat is bundled with the distribution.

Atlassian only supports the application servers listed on the Supported Platforms page, provided they are running on Windows, Linux, or Solaris. If you are using a different application server or earlier version, we may ask you to migrate to one of the supported application servers before we can provide you with further support.

Databases

A database is required to run Confluence. Atlassian supports the databases listed on the Supported Platforms page.

When evaluating Confluence, you can use the embedded database included in the Confluence installation.

When moving to a production installation, you must set up an external database server. If you have no preference for a particular database, we highly recommend using PostgreSQL. This is a scalable, robust and free database server that is also easy to set up. For database setup information, please refer to Database Configuration.

Java

Confluence requires the Java Runtime Environment (JRE) installed.

If using the zip or archive distribution of Confluence, you will need to install a supported JRE. The automated installer bundles Java and will install this for you.
For instructions on installing the JRE for Windows and Linux/Solaris, please refer to Installing Java for Confluence.

Please Note: Impact of Antivirus Software

The presence of antivirus software on your operating system running Confluence greatly decreases the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental and may even cause errors in Confluence.

You should configure your antivirus software to ignore the following directories:

- Confluence home directory
- Confluence's index directory
- All database-related directories

⚠️ This recommendation above is particularly important if you are running Confluence on Windows. No matter how fast your hardware is, antivirus software will almost always have a negative impact on Confluence's performance and may render Confluence impossible to use.

Confluence Hardware Requirements

Please be aware that while some of our customers run Confluence on SPARC-based hardware, Atlassian only officially supports Confluence running on x86 hardware and 64-bit derivatives of x86 hardware.

See Server Hardware Requirements Guide for details.

Refer also to the tips on reducing out of memory errors, in particular the section on Permanent Generation Size.

Hosted Solutions – Confluence Cloud

If you do not have the resources to set up and maintain a Confluence installation locally, consider 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.

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 512mb. 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.

ℹ️ 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 your configuration will likely require better hardware.

5 Concurrent Users

- **CPU:** 2 x Intel Core 2 (2.66 Ghz, 128K cache)
- **RAM:** 2GB
- **Minimum database space:** 10GB

25 Concurrent Users

- **CPU:** Quad 2GHz+ CPU
- **RAM:** 4GB
- **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
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. See Accenture's slides and video for full details (That link isn't working, but the slides can be found here)

Server load and scalability

When planning server hardware requirements for your Confluence deployment, you will need to estimate the server scalability based on peak visitors, the editor to viewer ratio and total content.

- The editor to viewer ratio is how many visitors are performing updates versus those only viewing content
- Total content is best estimated by a count of total spaces

Confluence scales best with a steady flow of visitors rather than defined peak visitor times, few editors and few spaces. Users should also take into account:

- Total pages is not a major consideration for performance. For example, instances hosting 80K of pages can consume under 512MB of memory
- Always use an external database, and check out the performance 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 use either the database or file system. For example, the wiki instance you are reading now uses approximately 2.8 GB of database space and 116 GB of disk space. The more attachments you have, the more disk space you will require.

Private and public comparison

Private instances manage their users either internally or through a user repository such as LDAP, while online instances have public signup enabled and must handle the additional load of anonymous internet visitors. Please keep in mind that these are examples only, not recommendations:

Use Case	Spaces	User Accounts	Editors	Editor To Viewer Ratio	Pages	Page Revisions	Attachments	Comment

Online Documentation	140	11,500	1,000	9%	8,800	65,000	7,300	11,500
Private Intranet	130	180	140	78%	8,000	84,000	3,800	500
Company-Wide Collaboration	100	85,000	1,000+	1%+	12,500	120,000	15,000	

Professional assistance

For large instances, it may be worthwhile contacting an Atlassian Expert for expertise on hardware sizing, testing and performance tuning. Simply contact a local Expert directly or email our Experts team for a recommendation.

Example - https://confluence.atlassian.com/

Here is a breakdown of the disk usage and memory requirements for this wiki, as at April 2013:

Database size	2827 MB
Home directory size	116 GB
Average memory in use	1.9 GB

Size of selected database tables

Data	Relevant Table	Rows	Size
Attachment metadata	`attachments`	193903	60 MB
Content and user properties	`os_propertyentry (?)`	639737	255 MB
Content bodies (incl. all versions of blogs, pages and comments)	`bodycontent`	517520	1354 MB
Content metadata (incl. title, author)	`content`	623155	459 MB
Labels	`label (5982, 1264 kB),` `content_label (134151, 46 MB)`	140133	47.2 MB
Users	`users`	38766	6200 kB

Note: not all database tables or indexes are shown, and average row size may vary between instances.

Size of selected home directory components

Data	Files	Size
Attachments (incl. all versions)	207659	105 GB
Did-you-mean search index	10	14 MB
Office Connector cache	3506	456 MB
Plugin files	1851	669 MB
Search index	448	3.9 GB
Temporary files	14232	5 GB
Thumbnails	86516	1.7 GB
Usage index (now disabled)	239	2.6 GB

Note: not all files are shown, and average file size may vary between instances.

Example Size and Hardware Specifications From Customer Survey

Below are the results of a survey conducted by Atlassian in July 2007, showing some capacity statistics for Confluence users. The figures are broken down by industry and number of users.

Num Users	Length of time in production	Database	Application Server	Num CPUs/Cores	Physical Memory/RAM
Banking/Finance					
26 - 50	3-6 Months Ago	Microsoft SQL Server	Confluence distribution/Apache Tomcat	2	2G
26 - 50	2 Years Ago	Sybase ASE	Weblogic	>8	>16G
51 - 250	3-6 Months Ago	Oracle	Confluence distribution/Apache Tomcat	2	4G
501 - 1,000	3-6 Months Ago	Microsoft SQL Server	Websphere	2	2G
1,001 - 5,000	3-6 Months Ago	Oracle	Confluence distribution/Apache Tomcat	2	4G
1,001 - 5,000	2 Years Ago	Oracle	Websphere	4	>16G
5,001 - 10,000	10-12 Months Ago	Microsoft SQL Server	Confluence distribution/Apache Tomcat	4	16G
Education					
1-25	2 Years Ago	DB2	Confluence distribution/Apache Tomcat	2	2G

26 - 50	10-12 Months Ago	MySQL	Confluence distribution/Apache Tomcat	2	2G
51 - 250	<3 Months Ago	Oracle	Confluence distribution/Apache Tomcat		1G
51 - 250	10-12 Months Ago	Oracle	Confluence distribution/Apache Tomcat	1	2G
Engineering/Aerospace					
251 - 500	7-9 Months Ago	Oracle	Confluence distribution/Apache Tomcat	1	1G
1,001 - 5,000	7-9 Months Ago	Microsoft SQL Server	JBoss	2	4G
Entertainment					
1,001 - 5,000	10-12 Months Ago	PostgreSQL	Confluence distribution/Apache Tomcat	2	8G
Government					
51 - 250	2 Years Ago	MySQL	Confluence distribution/Apache Tomcat	2	2G
Technology					
501 - 1,000	7-9 Months Ago	MySQL	Confluence distribution/Apache Tomcat	1	2G
Telecommunications & Media					
1-25	3-6 Months Ago	Confluence distribution/HSQL	Confluence distribution/Apache Tomcat	1	
1-25	7-9 Months Ago	MySQL	Confluence distribution/Apache Tomcat	1	2G
26 - 50	10-12 Months Ago	MySQL	Confluence distribution/Apache Tomcat	2	2G

Running Confluence in a Virtualised Environment

This page provides pointers for things to look at when running Confluence on virtualised hardware.

Summary

Running Confluence in a virtual machine (VM) requires specialised skills to set up and manage the virtualised environment. In particular, the performance of Confluence can be affected by the activity of other VMs running on the same infrastructure, as well as how you configure the Confluence VM itself.

Atlassian supports running Confluence and Confluence Data Center in a virtualised 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 Enterprise Resources
- Confluence Installation Guide

Recommendations

The following recommendations come from our experience in running and testing Confluence in virtualised environments like VMWare and KVM, and our experience in working with customers running on these platforms.

- **Know your platform.** Consult the documentation for your operating system and your chosen virtualisation technology, for details on setting up a reliable VM (virtual machine) image.
- **Allocate enough memory.** As a Java web application, Confluence requires a relatively large memory allocation, compared to some other web technologies. Ensure that your VM images have enough physical memory allocated to run Confluence without swapping.
- **Handle high I/O.** Under normal usage, Confluence requires a significant number of input/output (I/O) operations to the database and home directory for each web request. Ensure that you use the correct drivers and consider how you make storage available to your VMs to optimise this access.
- **Handle peak CPU and memory usage.** For certain operations (including PDF export, Office document processing, and displaying large pages) Confluence requires a significant amount of CPU and memory. Ensure that your virtualisation infrastructure has the flexibility and capacity to deal with peak load, not just idle load.
- **Synchronise time correctly.** Some customers have had problems with time synchronisation between the VM and the host system. This causes problems in Confluence due to irregularities in the execution of scheduled tasks. We strongly recommend checking your VM time sync if you have issues with scheduled tasks in a virtualised environment.

Further help

For further assistance in setting up a virtualised environment for running Confluence, you may want to consult an Atlassian Expert. Several experts have experience with installation and performance tuning, and can help you with your Confluence configuration.

Confluence Installation Guide

Before you start

Before installing Confluence, please check that you meet the minimum system requirements.

If you've chosen a package that includes add-ons such as Confluence Questions, Team Calendars or the SharePoint Connector you'll need to install these from within Confluence after setup is complete. See Finding new add-ons for how to find and install add-ons.

Choose the Confluence Installation Type

There are two methods of installing Confluence, using an installer, or manually from a standalone zip or archive file.

Install option	Is this right for you?
Install Confluence using an installer - Windows - Linux	This is the easiest method for installing Confluence. If you're evaluating Confluence, use this option or try Confluence Cloud free.

Install Confluence from a zip or archive file • Windows • Linux	This option requires you to manually install files and configure some system properties. Use this option if there isn't an installer for your operating system. The EAR/WAR distribution is no longer available, you should use this option if you previously deployed Confluence into an existing application server.
Install Confluence in a cluster Windows or Linux	Confluence Data Center is a clustered solution for large enterprises. Read the Confluence Data Center Technical Overview to find out if Confluence Data Center is right for your organisation.

Please read Running Confluence in a Virtualised Environment if you are interested in running Confluence in a virtual machine.
If you wish to upgrade Confluence, see Upgrading Confluence.

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 the supported browsers.

Installing Confluence

Choose the type of Confluence installion you'd like from the table below and follow the link to the installation instructions. When you have finished the installation phase, you will be prompted to start the setup phase.

Install option	Is this right for you?
Install Confluence using an installer • Windows • Linux	This is the easiest method for installing Confluence. If you're evaluating Confluence, use this option or try Confluence Cloud free.
Install Confluence from a zip or archive file • Windows • Linux	This option requires you to manually install files and configure some system properties. Use this option if there isn't an installer for your operating system. The EAR/WAR distribution is no longer available, you should use this option if you previously deployed Confluence into an existing application server.
Install Confluence in a cluster Windows or Linux	Confluence Data Center is a clustered solution for large enterprises. Read the Confluence Data Center Technical Overview to find out if Confluence Data Center is right for your organisation.

ⓘ If you have not already done so, please verify that this version of the Confluence documentation matches that of the Confluence version you are installing. The Confluence documentation version you are currently viewing is indicated toward the top of the page tree on the left or in the 'breadcrumb trail' in the top banner of this page. If you need to access a different version of the Confluence documentation, use the control at the top of the page tree on the left or you can access it from the documentation home page.

Take me back to the Confluence Installation Guide.

Installing Confluence on Windows

This guide describes how to install a new Confluence installation on Windows using the installer. You can also install Confluence manually from a 'zip' archive - see Installing Confluence on Windows from Zip File for details.

If you are upgrading Confluence, please refer to the Upgrading Confluence guide.

⚠️ **Note:** Some anti-virus or other Internet security tools may interfere with the Confluence installation process and prevent the process from completing successfully. If you experience or anticipate experiencing such an issue with your anti-virus/Internet security tool, disable this tool first before proceeding with the Confluence installation.

Before you Start

1. Ensure that your system meets the minimum requirements to run Confluence. See System Requirements, Supported Platforms and be sure to check the Database Configuration documentation for any specific requirements related to your chosen database.
2. Have your Confluence license key ready. You can obtain a trial, free or commercial license, or retrieve your existing license key at my.atlassian.com.

Using the Installation Wizard

Use the installation wizard if you are installing Confluence on your server for the first time or you wish to specify your installation options.

If you have previously installed Confluence using the installation wizard and wish to re-install Confluence again with the same installation options, you can re-install Confluence in 'unattended mode' without any user input required (see below for details).

1. Download and Run the Confluence 'Windows Installer'

✅ To install Confluence as a service, the Windows Installer must be run using a Windows administrator account. While you can run the Windows Installer with a non-administrator account, your installation options will be much more limited.

1. Download the **Confluence 'Windows Installer'** (.exe) file from the Confluence Download page.
2. Run the installer file to start the installation wizard.
   ℹ️ If a Windows 7 (or Vista) 'User Account Control' dialog box requests if you want to allow the installation wizard to make changes to your computer, click '**Yes**'. If you do not, the installation wizard will have restricted access to your operating system and any subsequent installation options will be limited.
3. Choose between the '**Express Install**' or '**Custom Install**' options:
   a. **Express Install** — If you choose this option, Confluence will be installed with default settings which are shown in the next step of the installation wizard. If you want to customise any of these options, click the '**Back**' button and choose the '**Custom Install**' option instead.
   b. **Custom Install** — If you choose this option, Confluence will prompt you to specify the following options (which are presented during subsequent steps of the installation wizard and pre-populated with default values):
      - The 'Destination Directory' in which to install Confluence.
      - The Confluence Home and other important directories (which must be unique for each Confluence installation).
      - The Windows 'Start' menu folder options.
      - The *TCP ports* (i.e. an HTTP connector port and a control port) that Confluence will operate on.
      - If you are running the installer using an administrator account, you will be prompted to 'Install Confluence as a service' (recommended). You can also do this manually later, as described in Start Confluence Automatically on Windows as a Service.

Confluence 5.9 Documentation

> ✅ If you installed Confluence as a service, you must start Confluence through the Windows 'Start' menu, since Confluence will not start if you run `start-confluence.bat` at the Windows Command Prompt.

4. The installation wizard will install Confluence onto your operating system and will start Confluence automatically when the wizard finishes. Confluence will also be launched automatically in your browser window if you chose this option.

ℹ️ **Please Note:**

- If you chose to install Confluence as a service, the Confluence service will be run as the Windows 'SYSTEM' user account. To change this user account, see Changing the Windows user that the Confluence service uses.
- If you do not install Confluence as a service, then once started, Confluence will be run as the Windows user account under which Confluence was installed.
- If you use Confluence running on a Windows Server in production, we strongly recommend creating a dedicated user account (e.g. with username 'confluence') for running Confluence.
  - For more information about creating a dedicated user account and defining which directories this account should have write access to, refer to our guidelines.
    ✅ If your Windows Server is operating under Microsoft Active Directory, ask your Active Directory administrator to create a dedicated user account that you can use to run Confluence (with no prior privileges).
  - If Confluence is installed as a service, do not forget to change the user account that runs the Confluence service to your dedicated user account for running Confluence.

2. Starting Confluence

If Confluence is not already started, you can start Confluence using the appropriate Windows 'Start' menu shortcut or command prompt option.

Once Confluence is started, you can access Confluence from the appropriate Windows 'Start' menu shortcut or a browser on any computer with network access to your Confluence server.

2.1 Windows 'Start' Menu Shortcuts

The Installer will have created the following Windows 'Start' menu shortcuts:

- **Access Confluence** — opens a web browser window to access your Confluence application.
  ℹ️ Your Confluence server must have been started for this shortcut to work.
- **Start Confluence Service** — starts up the Apache Tomcat application server which runs your Confluence installation, so that you can access Confluence through your web browser.
- **Stop Confluence Service** — stops the Apache Tomcat application server which runs your Confluence installation. You will not be able to access Confluence through your web browser after choosing this shortcut.
- **Uninstall Confluence** — uninstalls Confluence from your Windows operating system.

2.2 Starting and Stopping Confluence from a Command Prompt

Enter the `bin` subdirectory of your Confluence installation directory and run the appropriate file:

- `start-confluence.bat` (to start Confluence)
- `stop-confluence.bat` (to stop Confluence)

ℹ️ If you followed our guidelines for running Confluence with a dedicated user account, then to run Confluence as this user account (e.g. 'confluence'), use the `runas` command to execute `start-confluence.bat`. For example:

- `> runas /env /user:<DOMAIN>\confluence start-confluence.bat`
  (where `<DOMAIN>` is your Windows domain or computer name.)

2.3 Accessing Confluence from a Browser

You can access Confluence from any computer with network access to your Confluence server by opening a supported web browser on the computer and visiting this URL:

- `http://<computer_name_or_IP_address>:<HTTP_port_number>`

where:

- `<computer_name_or_IP_address>` is the name or IP address of the computer on which Confluence is installed and
- `<HTTP_port_number>` is the HTTP port number specified when you installed Confluence (above).

ⓘ If Confluence does not appear in your web browser, you may need to change the port that Confluence runs on.

3. Run the Setup Wizard

See the Confluence Setup Guide.

4. Next Steps

- See Confluence Getting Started.
- If you did not install Confluence as a service, you will need to start Confluence manually every time you restart your computer. To change your Confluence installation to run as a service, please see Start Confluence Automatically on Windows as a Service.
- To get the most out of Confluence, please see Performance Tuning.

Performing an Unattended Installation

If you have previously installed Confluence using the installation wizard (above), you can use a configuration file from this Confluence installation (called `response.varfile`) to re-install 'unattended mode' without any user input required.

Installing Confluence in unattended mode saves you time if your previous Confluence installation was used for testing purposes and you need to install Confluence on multiple server machines based on the same configuration.

⚠ **Please Note:**

- The `response.varfile` file contains the options specified during the installation wizard steps of your previous Confluence installation. Hence, do not uninstall your previous Confluence installation just yet.
- If you intend to modify the `response.varfile` file, please ensure all directory paths specified are absolute, for example, `sys.installationDir=C\:\\Program Files\\Atlassian\\Confluence`
  Unattended installations will fail if any relative directory paths have been specified in this file.

Download and Run the Confluence 'Windows Installer' in Unattended Mode

1. Download the **Confluence 'Windows Installer'** (.exe) file from the Confluence Download Center to a suitable location.
2. Open the Windows command prompt and perform the remaining steps in the command prompt.
3. `copy` the `response.varfile` file located in the `.install4j` subdirectory of your previous Confluence installation directory, to the same location as the downloaded 'Windows Installer' file.
   ⓘ You can uninstall your previous Confluence installation after this step. Save your `response.varfile` if you need to install Confluence on multiple machines.
4. Change directory (`cd`) to the location of the 'Windows Installer' file and run the following command:

   ```
   atlassian-confluence-X.Y.exe -q -varfile response.varfile
   ```

   Where:
   - `X.Y` — refers to the version of Confluence you are about to install.
   - `-q` — instructs the installer to operate in unattended mode (i.e. 'quietly').
   - `-varfile response.varfile` — specifies the configuration file containing the configuration options used by the installer. The location and name of the configuration file should be specified after the `-varfile`option.
5. Confluence will start automatically when the silent installation finishes. Continue from step 2 Starting Confluence (above).

Installing Confluence on Windows from Zip File

> **These instructions apply to:**
> - Confluence distributed as an archive (zip) file. This distribution includes Apache Tomcat as the application server.
> - **Windows** systems. For other operating systems please refer to the Confluence Installation Guide .
> - Manual installation and configuration using a **zipped download file**. For a simpler installation process, please use the Confluence Installer instead.
>
> Also, please check that the version of Confluence which you are installing coincides with the version that this documentation is written for.

On this page:

1. Before you Start

Please check the following points:

1. Ensure that your system meets the minimum requirements to run Confluence. For more information, please refer to our Supported Platforms topic and for further details, our System Requirements topic.
2. Have your Confluence license key ready. You can obtain a trial, free or commercial license, or retrieve your existing license key at my.atlassian.com.

2. Install Java

Please refer to Installing Java for Confluence. If you are certain that this has already been installed and that the `JAVA_HOME` environment variable has been correctly configured, then proceed to the next step.

3. Download the Confluence Installation File

1. If you have not downloaded Confluence already, download the zip file.
2. Please check your unzip program before extracting the downloaded zip file. You should use a third-party unzip program like 7Zip or Winzip. If you do not have one, please download and install one of these before continuing:
   - 7Zip (recommended). If in doubt, download the '32-bit.exe' version.
   - Winzip.
3. Use your unzip program to unzip the installation file to a directory such as `c:\confluence`.
   - Do not use spaces in your directory path.

The directory into which you unzipped the Confluence installation is called the Confluence Installation directory. Next, you will define the Confluence Home directory.

4. Define your Confluence Home Directory

Now you need to define the Confluence Home directory. This is where Confluence will store its configuration information, indexes and attachments.
Tip: Another term for 'home directory' would be 'data directory'.

We suggest using different paths for your installation and home directories. This will facilitate easier upgrades.

> **Examples of Installation and Home Directories**
> Installation directory: `c:\confluence\confluence-vX.X`
> Home directory: `c:\confluence\data`

1. Open your Confluence Installation directory (created when you unzipped Confluence — see above).
2. Under the Installation directory, open this file: `confluence\WEB-INF\classes\confluence-init.properties` in a text editor such as Notepad.
3. Scroll to the bottom of the text and find this line:

```
# confluence.home=c:/confluence/data
```

4. Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with `confluence.home`
5. If you decide to change the Confluence Home directory from the default, please note the following:
     * Avoid spaces in the directory path or file name.
     * Use forward slashes '/' to define the path.

For example:

```
confluence.home=c:/data/confluence-home
```

5. Check the Ports

If you have another application running on your machine which is using the same ports that Confluence uses by default, you may need to change the port which Confluence will use. For example, if you have a JIRA application running on this machine, JIRA might be already using the port which Confluence requests by default.

By default, Confluence listens on port '8090'. If this port is already in use in your installation, follow these instructions to change the ports:

* To change the ports for Confluence, open the file `conf/server.xml` under your Confluence Installation directory. The first four lines of the file look like this:

  Default conf/server.xml
  ```
  <Server port="8000" shutdown="SHUTDOWN" debug="0">
      <Service name="Tomcat-Standalone">
          <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
  port="8090" minProcessors="5" maxProcessors="75"
              enableLookups="true" redirectPort="8443" acceptCount="10"
  debug="0" connectionTimeout="20000" useURIValidationHack="false"/>
          ...
  ```

  You need to modify both the **server** port (default is 8000) and the **connector** port (default is 8090) to ports that are free on your machine. The server port is required by Tomcat but is not user facing in any way. The connector port is what your users will use to access Confluence, eg in the snippet above, the URL would be `http://example.com:8090`.

  ✅ Hint: You can use netstat to identify free ports on your machine. See more information on using netstat on Windows or on Linux.

  For example, here are the first four lines of a modified `server.xml` file, using ports '8020' and '8099':

> **Modified conf/server.xml using ports 8020 and 8099**
>
> ```
> <Server debug="0" shutdown="SHUTDOWN" port="8020">
>     <Service name="Tomcat-Standalone">
>         <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
> port="8099" minProcessors="5" maxProcessors="75"
>             enableLookups="true" redirectPort="8443" acceptCount="10"
> debug="0" connectionTimeout="20000" useURIValidationHack="false"/>
>         ...
> ```

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.

You will find more information on this page.

6. Select an External Database

This step is optional for users evaluating Confluence. However, if you are installing Confluence for production purposes, this step is mandatory. Please refer to the database requirements listed on our System Requirements topic for help in choosing an external database.

ℹ️ External databases are those listed on our Supported Platforms topic, excluding HSQLDB, which is bundled with Confluence and should not be used in production.

When you have chosen your external database, follow the the appropriate database setup guide to set up your database to work with Confluence.

You can learn more about migration from an existing installation or use of the evaluation database here. You will continue to use the database setup guide during the Confluence Setup Wizard. (See step 8 below.)

7. Start Confluence

1. Go to your Confluence Installation directory (created when you unzipped Confluence — see above).
2. Under your Confluence Installation directory, open the `bin` directory and run the startup script: `startup.bat`. A command prompt window should appear.
   ⚠️ Please do not close this command prompt window. If you do so, Confluence will stop running.

> **Troubleshooting**
> If the window closes immediately when started, this means that an error is preventing Confluence from starting. To view this error:
>
> a. Open a command prompt: Click on your 'Start' menu, then click 'Run'. In the Run box, type `cmd` and click 'OK'.
> b. From the command prompt, go to your Confluence Installation directory.
> c. Go into the `bin` subdirectory.
> d. Run `catalina.bat run`.
>   ⚠️ You should not run `startup.bat` at this point, because that would still produce a popup window that would close straight away.
> e. Read the error message.
> f. Find the solution to that error in the Installation Troubleshooting section of the Confluence Knowledge Base.

3. Once Confluence is running, open a web browser and visit `http://localhost:8090/`.
   ✅ If you changed the port earlier, use the port you specified in step 5 above.
   ℹ️ If your web browser window shows an error, try waiting for 30 seconds or so and then refresh the

browser page.

8. Next Step is the Confluence Setup Wizard

The Confluence Setup Wizard should appear in your web browser, prompting you to enter your license key. Follow the instructions on the screens, and read more guidelines on the Confluence Setup Wizard.

9. Start Confluence automatically on Windows as a Service

Confluence should be run as a service.

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 **Progr ams and Features** on (Windows 7, Vista) and then select **Confluence X.Y** from the list of applications and click **Uninstall/Change**.
      OR
    - Open the Windows command prompt and do the following:
        a. Change directory to your Confluence installation directory
        b. Run the `uninstall.exe` file
3. Follow the prompts to uninstall Confluence from your computer.

ℹ️ 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` f older located in the Confluence Installation Directory).
- The uninstaller can be made to operate in unattended mode by specifying the `-q` option at the Windows command prompt — i.e. `uninstall -q`
- If you wish to re-install Confluence in 'unattended mode', do not uninstall your previous installation of Confluence just yet. See Using the Silent Installation Feature for more information.

Installing Confluence on Linux

This guide describes how to install a new Confluence installation on Linux using the installer. You can also install manually from a zip archive - see Installing Confluence on Linux from Archive File for details.

If you are upgrading Confluence, please see Upgrading Confluence.

Before you Start

Ensure that your system meets the minimum requirements to run Confluence. See System Requirements, Supported Platforms and be sure to check the Database Configuration documentation for any specific requirements related to your chosen database.

On this page:

- Before you Start
- Using the Console Wizard
    - 1. Download and Install the Confluence 'Linux Installer'
    - 2. Start Confluence
    - 3. Run the Setup Wizard
    - 4. Next Steps
- Performing an Unattended Installation
    - Download and Run the Confluence 'Linux Installer' in Unattended Mode

⚠️ **Note:** It is possible that any anti-virus or other Internet security tools installed on your Linux operating system may interfere with the Confluence installation process and prevent the process from completing successfully. If you experience or anticipate experiencing such an issue with your anti-virus/Internet security tool, disable this tool first before proceeding with the Confluence installation.

You may also wish to consider disabling the Linux OutOfMemory Killer (OOM Killer). This is a Linux function that can kill processes when memory on the server becomes low, and sometimes targets Confluence.

Using the Console Wizard

Use the console wizard if you are installing Confluence on your server for the first time or you wish to specify your installation options.

If you have previously installed Confluence using the installation wizard and wish to re-install Confluence again with the same installation options, you can re-install Confluence in 'unattended mode' without any user input required (see below for details).

1. Download and Install the Confluence 'Linux Installer'

✅ If you execute the Linux Installer with 'root' user privileges, the installer will create and run Confluence using a dedicated user account. You can also execute the Linux Installer without 'root' user privileges, although your installation options will be much more limited and a dedicated user account (to run Confluence) will not be created. To run Confluence as a service, the Linux Installer must be executed with 'root' user privileges.

1. Download the appropriate **Confluence 'Linux 64-bit / 32-bit Installer'** (.bin) file from the Confluence Download page.
   ℹ️ **Please Note:**
   - To access the 32-bit installer, you may need to click the '**Show all**' link on the 'Confluence Download' page to access the other installation packages.
   - The difference between the 64-bit / 32-bit .bin installers relates to their bundled Java platforms that run Confluence. Bear in mind that a Confluence installation installed using the 64-bit installer may require additional memory (to run at a similar level of performance) to a Confluence installation installed using the 32-bit installer. This is because a 64-bit Java platform's object references are twice the size as those for a 32-bit Java platform.
2. Open a Linux console and change directory (`cd`) to the '.bin' file's directory.
   ⚠️ If the '.bin' file is not executable after downloading it, make it executable, for example:
   `chmod a+x atlassian-confluence-X.Y.bin`
   (where X.Y represents your version of Confluence)
3. Execute the '.bin' file to start the console wizard.
4. When prompted to choose between '**Express Install**', '**Custom Install**' or '**Upgrade an existing Confluence installation**', choose either the '**Express Install**' or '**Custom Install**' options:
   - **Express Install** — If you choose this option, Confluence will be installed with default settings which are shown in the next step of the console wizard.
     ℹ️ **Please Note:**
     - If you are running the installer with 'root' user privileges, Confluence will be installed as a service.
     - If you want to customise any of these options:
       i. Enter '**e**' to exit the console wizard.
       ii. Execute the console wizard again (step 3 above).
       iii. Choose the '**Custom Install**' option instead.
   - **Custom Install** — If you choose this option, Confluence will prompt you to specify the following options (which are presented during subsequent steps of the console wizard and pre-populated with default values):
     - The 'Destination Directory' in which to install Confluence.
     - The Confluence Home directory (which must be unique for each Confluence installation).
     - The *TCP ports* (i.e. an HTTP and a Control port) that Confluence will run through.
     - If you are running the installer with 'root' user privileges, you will be prompted to 'Run Confluence as a service' (recommended). You can also do this manually later, as described in Start Confluence Automatically on Linux.
5. The console wizard will install Confluence onto your operating system and will start Confluence automatically when the wizard finishes.

ℹ️ **Please Note:**

If you executed the Linux Installer with 'root' user privileges, the Linux Installer creates a dedicated Linux user account with username 'confluence', which is used to run Confluence. This account has only:

- Full write access to your Confluence Home and other important directories.
- Limited write access to your Confluence Installation Directory.

If you executed the Linux Installer without 'root' user privileges, be aware that Confluence can still be run with 'root' privileges. However, to protect the security of your operating system, *this is not recommended*.

2. Start Confluence

If Confluence is not already started, you can start Confluence using the appropriate command at the Linux console.

Once Confluence is started, you can access Confluence from a browser on any computer with network access to your Confluence server.

2.1 Starting and Stopping Confluence manually

In the Linux console, enter the `bin` subdirectory of your Confluence installation directory and execute the appropriate file:

- `start-confluence.sh` (to start Confluence)
- `stop-confluence.sh` (to stop Confluence)

Confluence will be ready to access (from a browser window) when the following message appears in the application's log file:

```
*********************************************************
 ... You can now access Confluence through your web browser.
*********************************************************
```

2.2 Accessing Confluence from a Browser

You can access Confluence from any computer with network access to your Confluence server by opening a supported web browser on the computer and visiting this URL:

- `http://<computer_name_or_IP_address>:<HTTP_port_number>`

where:

- `<computer_name_or_IP_address>` is the name or IP address of the computer on which Confluence is installed and
- `<HTTP_port_number>` is the HTTP port number specified when you installed Confluence (above).

If Confluence does not appear, you may need to change the port that Confluence runs on.

Note: Application server logs (i.e. for Apache Tomcat) will be written to `logs/catalina.out`.

3. Run the Setup Wizard

See the Confluence Setup Guide.

4. Next Steps

- See Confluence Getting Started.
- If you did not install Confluence to run as a service, you will need to start Confluence manually every time you restart your computer. To change your Confluence installation to run as a service, please see Start Confluence Automatically on Linux.
- To get the most out of Confluence, please see Performance Tuning.

Performing an Unattended Installation

If you have previously installed Confluence using the console wizard (above), you can use a configuration file from this Confluence installation (called `response.varfile`) to re-install Confluence in 'unattended mode' without any user input required.

Installing Confluence in unattended mode saves you time if your previous Confluence installation was used

for testing purposes and you need to install Confluence on multiple server machines based on the same configuration.

> ⚠️ **Please Note:**
>
> - The `response.varfile` file contains the options specified during the installation wizard steps of your previous Confluence installation. Hence, do not uninstall your previous Confluence installation just yet.
> - If you intend to modify the `response.varfile` file, please ensure all directory paths specified are absolute, for example, `sys.installationDir=/opt/atlassian/confluence`
>   Unattended installations will fail if any relative directory paths have been specified in this file.

Download and Run the Confluence 'Linux Installer' in Unattended Mode

1. Download the **Confluence 'Linux Installer'** (.bin) file from the Confluence Download Center to a suitable location.
2. Open a Linux console.
3. Copy (`cp`) the file `.install4j/response.varfile` located in your previous Confluence installation directory, to the same location as the downloaded 'Linux Installer' file.
   ℹ️ You can uninstall your previous Confluence installation after this step. Save your `response.varfile` if you need to install Confluence on multiple machines.
4. Change directory (`cd`) to the location of the 'Linux Installer' file and execute the following command:

```
atlassian-confluence-X.Y.bin -q -varfile response.varfile
```

   Where:
   - `X.Y` — refers to the version of Confluence you are about to install.
   - `-q` — instructs the installer to operate in unattended mode (i.e. 'quietly').
   - `-varfile response.varfile` — specifies the configuration file containing the configuration options used by the installer. The location and name of the configuration file should be specified after the `-varfile` option.
5. Confluence will start automatically when the silent installation finishes. Continue from the step above, Starting Confluence.

Installing Confluence on Linux from Archive File

> **These instructions apply to:**
> - Confluence distributed as an archive (tar.gz) file.
> - **Linux or Solaris** systems. If you are installing Confluence on a different system, please refer to Installing Confluence.

Hint: If you're evaluating Confluence on Solaris or OSX, this is the version you need to install.

On this page:

- 1. Before you start
- 2. Install Java
- 3. Download and extract the Confluence installation file
- 4. Define your Confluence Home Directory
- 5. Check the Ports
- 6. Select an External Database
- 7. Start Confluence
- 8. Confluence Setup Wizard

1. Before you start

Please check the following points:

1. Ensure that your system meets the minimum requirements to run Confluence. For more information, please read the detailed System Requirements.
2. Have your Confluence license key ready. You can obtain a trial, free or commercial license, or retrieve

your existing license key at my.atlassian.com.
3. You must be able to use a command prompt and install Java to continue.
4. Make sure that you use a Gnu version of zip application - Solaris and AIX are known to have problems wit h zip, because they use their own (old) versions instead of the Gnu version.

2. Install Java

Please refer to the Supported Platforms for the required version of Java. (OpenJDK is currently **not** supported)

If you're not sure whether you have Java installed correctly:

1. Open a shell console.
2. Type `echo $JAVA_HOME` in the shell console and then press Enter
   View the result:
   - If a line is displayed such as `/opt/jdk1.8.0_45` or `/usr/lib/jvm/java-8-sun`, then Java is installed and properly configured.
   - If nothing is displayed, then you either need to install Java or set the `$JAVA_HOME` environment variable. You can set this environment variable in your user account's 'profile' file. Alternatively, you can set this after installing Confluence (in step 4 below) by defining this path in your Confluence installation's `setenv.sh` file, usually located in the Confluence `bin` directory.

If you need to install Java, follow these instructions:

1. Go to the Java download page and download the latest JRE or JDK that is listed on the Confluence Supp orted Platforms page. (Confluence works with either the JDK or the JRE.)
2. When the download has finished, run the Java installer. Detailed installation instructions are provided on Oracle's website.
   Note: you will be asked to choose an installation directory. Make a note of this directory, you'l' need it later.

3. Download and extract the Confluence installation file

If you have not downloaded Confluence already, download the TAR.GZ file.

Use your unzip program to unzip the installation file to a directory such as `/home/jsmith/confluence-2.7.0-std/`.

Most Linux/Solaris users can use any unzip program (such as GNU Tar) to extract the Confluence installer. However, Solaris users should not use the Solaris Tar program due to a known issue associated with its use in extracting Confluence. Use another application such as GNU Tar instead.

For example, change directory to your home directory in Linux and enter the following commands in the shell console:

- `gunzip confluence-<version>-std.tar.gz`
- `tar -xf confluence-<version>-std.tar`
  (where `<version>` refers to the Confluence version you downloaded.)

ℹ️ As usual on Linux/Solaris-based operating systems, avoid using spaces in your directory path. The directory into which you unzipped the Confluence installation is called the Confluence Installation directory. Next you will define the Confluence Home directory.

4. Define your Confluence Home Directory

Now you need to define the Confluence Home directory. This is where Confluence will store its configuration information, indexes and attachments. Another term for 'Home directory' would be 'data directory'.

You should use different paths for your installation and home directories. This will facilitate upgrades. Examples of Installation and Home Directories:

- Installation directory: `/usr/local/confluence/`
  If you wish to install or maintain multiple versions of Confluence, you can add a version number to the Confluence installation directory name like `/usr/local/confluence-3.1-std/` and optionally, create the symbolic link `/usr/local/confluence/` that points to `/usr/local/confluence-3.1-std/`
- Home directory: `/usr/local/confluence-data/`

1. Open your Confluence Installation directory (created when you unzipped Confluence — see above).
2. Under the Installation directory, find this file: `confluence/WEB-INF/classes/confluence-init.properties`
3. Open the `confluence-init.properties` file in a text editor.
4. Scroll to the bottom and find this line:

```
# confluence.home=c:/confluence/data
```

5. Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with `confluence.home`
6. If you decide to change the Confluence Home directory from the default, use an absolute path rather than a symbolic link to specify the path and file name. For example:

```
confluence.home=/home/jsmith/confluence-data/
```

5. Check the Ports

If you have another application running on your machine which is using the same ports that Confluence uses by default, you may need to change the port which Confluence will use. For example, if you have a JIRA application running on this machine, JIRA might be already using the port which Confluence requests by default.

By default, Confluence listens on port '8090'. If this port is already in use in your installation, follow these instructions to change the ports:

- To change the ports for Confluence, open the file `conf/server.xml` under your Confluence Installation directory. The first four lines of the file look like this:

  Default conf/server.xml

  ```
  <Server port="8000" shutdown="SHUTDOWN" debug="0">
      <Service name="Tomcat-Standalone">
          <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
  port="8090" minProcessors="5" maxProcessors="75"
              enableLookups="true" redirectPort="8443" acceptCount="10"
  debug="0" connectionTimeout="20000" useURIValidationHack="false"/>
          ...
  ```

  You need to modify both the **server** port (default is 8000) and the **connector** port (default is 8090) to ports that are free on your machine. The server port is required by Tomcat but is not user facing in any way. The connector port is what your users will use to access Confluence, eg in the snippet above, the URL would be `http://example.com:8090`.

  ✅ Hint: You can use netstat to identify free ports on your machine. See more information on using netstat on Windows or on Linux.

  For example, here are the first four lines of a modified `server.xml` file, using ports '8020' and '8099':

Modified conf/server.xml using ports 8020 and 8099

```
<Server debug="0" shutdown="SHUTDOWN" port="8020">
    <Service name="Tomcat-Standalone">
        <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
port="8099" minProcessors="5" maxProcessors="75"
            enableLookups="true" redirectPort="8443" acceptCount="10"
debug="0" connectionTimeout="20000" useURIValidationHack="false"/>
        ...
```

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.

You will find more information on this page.

6. Select an External Database

This step is optional for users evaluating Confluence. However, if you are installing Confluence for production purposes, this step is mandatory. Please refer to the database requirements listed on our System Requirements topic for help in choosing an external database.

ⓘ External databases are those listed on our Supported Platforms topic, excluding HSQLDB, which is bundled with Confluence and should not be used in production.

When you have chosen your external database, follow the appropriate database setup guide to set up your database to work with Confluence.

You can learn more about migration from an existing installation or use of the evaluation database here. You will continue to use the Database Setup Guide during the Confluence Setup Wizard. (See step 8 below.)

7. Start Confluence

1. Go to your Confluence Installation directory (created when you unzipped Confluence — see above).
2. Under your Confluence Installation directory, open the `bin` directory and run the startup script: `start-co nfluence.sh`.
3. Once Confluence is running, open a web browser and visit http://localhost:8090/.
   ✅ Hint: If you changed the port earlier, use the port you specified in step 6 above.

8. Confluence Setup Wizard

The Confluence Setup Wizard should appear in your web browser, prompting you to enter your license key. Follow the instructions on the screens, and read more guidelines on the Confluence Setup Wizard.
Related Topics

Change listen port for Confluence
Running Confluence Over SSL or HTTPS
Confluence Setup Guide
Configuring Confluence

Uninstalling Confluence from Linux

This page describes the procedure for uninstalling Confluence, which had been installed using the Linux Installer .

To uninstall Confluence from Linux:

1. Open a Linux console.
2. Change directory (`cd`) to your Confluence installation directory.

3. Execute the command `uninstall`. This command must be executed as the same user account that was used to install Confluence with the [Linux Installer](#).
4. Follow the prompts to uninstall Confluence from your computer.

ⓘ Please note:

- The uninstaller will not delete the Confluence Home Directory.
- All log files that were generated while Confluence was running will not be deleted.
- All files within the Confluence Installation Directory will be deleted (with the exception of the Tomcat `log` f older located in the Confluence Installation Directory).
- The uninstaller can be made to operate in unattended mode by specifying the `-q` option — i.e. `uninstall -q`
- If you wish to re-install Confluence in 'unattended mode', do not uninstall your previous installation of Confluence just yet. See [Using the Silent Installation Feature](#) for more information.

Change listen port for Confluence

Problem

This page tells you what to do if you get errors like the following when starting Confluence, when you can't access Confluence on port **8090**.

If you see this error:

```
java.net.BindException: Address already in use: JVM_Bind:8090
```

This means you are running other software on Confluence's default port of **8090**. This may be another other process running on the same port. It may also be a previous instance of Confluence that hasn't been shut down cleanly.

To find out what process is listening on that port, load a command prompt and type: **netstat -an**

```
-a : Displays all active TCP connections and the TCP and UDP ports on
which the computer is listening.
 -n : Displays active TCP connections, however, addresses and port
numbers are expressed numerically and no attempt is made to determine
names.
```

There is also [Process Explorer tool](#) available to determine what is binding port **8090**.

Solution: Change the Ports which Confluence Listens On

To change the ports for Confluence, open the file `conf/server.xml` under your Confluence Installation directory. The first four lines of the file look like this:

Default conf/server.xml

```
<Server port="8000" shutdown="SHUTDOWN" debug="0">
    <Service name="Tomcat-Standalone">
        <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
port="8090" minProcessors="5" maxProcessors="75"
            enableLookups="true" redirectPort="8443" acceptCount="10" debug="0"
connectionTimeout="20000" useURIValidationHack="false"/>
        ...
```

You need to modify both the **server** port (default is 8000) and the **connector** port (default is 8090) to ports that are free on your machine. The server port is required by Tomcat but is not user facing in any way. The connector port is what your users will use to access Confluence, eg in the snippet above, the URL would be `http://example.com:8090`.

✅ Hint: You can use netstat to identify free ports on your machine. See more information on using netstat on Windows or on Linux.

For example, here are the first four lines of a modified `server.xml` file, using ports '8020' and '8099':

Modified conf/server.xml using ports 8020 and 8099

```
<Server debug="0" shutdown="SHUTDOWN" port="8020">
    <Service name="Tomcat-Standalone">
        <Connector className="org.apache.coyote.tomcat4.CoyoteConnector"
port="8099" minProcessors="5" maxProcessors="75"
            enableLookups="true" redirectPort="8443" acceptCount="10" debug="0"
connectionTimeout="20000" useURIValidationHack="false"/>
        ...
```

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, see Changing JIRA application TCP ports.

[3] You will need to restart Confluence after editing `server.xml` for the changes to take effect.

Installing Confluence Data Center

This guide provides step-by-step instructions for installing Confluence Data Center, which is a clustered solution, for the first time (with no existing data).

If you already have a Confluence Server instance, see Moving to Confluence Data Center.

1. Clustering requirements and terminology

To run Confluence in a cluster you must:

- Have a clustered license
- Use a supported external database, operating system and Java version
- Use a load balancer with session affinity in front of the cluster
- Have a shared directory accessible to all cluster nodes in the same path (this will be your shared home directory)
- Use OAuth authentication if you have application links to other Atlassian products (such as JIRA)

If you need a Data Center evaluation license please contact us.

In this guide we'll use the following terminology:

- Installation directory – The directory where you installed Confluence on a node.
- Local home directory – The home or data directory on each node (in non-clustered Confluence 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).

2. Install Confluence on the first node

1. Install Confluence on node 1
   See Installing Confluence on Windows from Zip File or Installing Confluence on Linux from Archive File fo r more information.
2. Start Confluence on Node 1
3. Complete the setup wizard by entering:

   - Your cluster license
   - A name for your cluster
   - The path to the shared home directory you created earlier
   - A multicast address (automatically generated or enter your own) or the IP addresses of each cluster node
   - The network interface Confluence will use to communicate between nodes

3. Load test the single node

Most Confluence installations don't need to be clustered. You might want to test your single node installation with the number of users and load you expect before going ahead with the additional complexity of clustering.

4. Copy Confluence to second node

To copy Confluence to the second node:

1. Shut down Confluence on node 1
2. Shut down your application server on node 2, or stop it automatically loading web applications
3. Copy the installation directory from node 1 to node 2
4. Copy the local home directory from node 1 to node 2
   If the file path of the local home directory is not the same on nodes 1 and 2 you'll need to update the `<installation directory>/confluence/WEB-INF/classes/confluence-init.properties` file on node 2 to point to the correct location.

Copying the local home directory ensures the Confluence search index, the database and cluster configuration, and any other settings are copied to node 2.

5. Start Confluence on the first node, wait, then start Confluence on second node

You must only start Confluence one server at a time.

1. Start Confluence on node 1
2. Wait for Confluence to become available on node 1
3. Start Confluence on node 2
4. Wait for Confluence to become available on node 2.

6. Test cluster connectivity
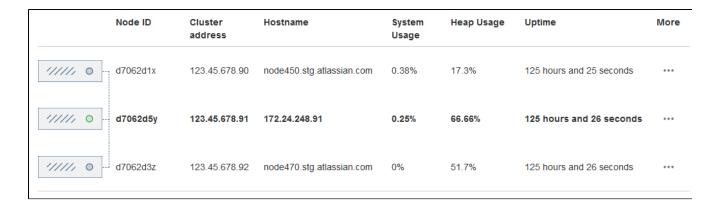
The Cluster monitoring console (

 > **General Configuration** > **Clustering**) includes 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.

Screenshot: *Cluster monitoring console*

	Node ID	Cluster address	Hostname	System Usage	Heap Usage	Uptime	More
///// ○	d7062d1x	123.45.678.90	node450.stg.atlassian.com	0.38%	17.3%	125 hours and 25 seconds	•••
///// ○	**d7062d5y**	**123.45.678.91**	**172.24.248.91**	**0.25%**	**66.66%**	**125 hours and 26 seconds**	•••
///// ○	d7062d3z	123.45.678.92	node470.stg.atlassian.com	0%	51.7%	125 hours and 26 seconds	•••

A simple process to ensure your cluster is working correctly is:

1. Create a new document on node 1
2. Ensure the new document is visible by accessing it directly on node 2
3. Search for the new document on node 1, and ensure it appears
4. Search for the new document on node 2, and ensure it appears

> If Confluence detects more than one instance accessing the database but not in a working cluster, it will shut itself down in a *cluster panic*. This can be fixed by troubleshooting the network connectivity of the cluster.

7. Configure load balancer

Install and configure your load balancer. You can use the load balancer of your choice, but it needs to support 'cookie-based session affinity'.

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.

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 yet running a clustered instance of Confluence, see Moving to Confluence Data Center.

In this guide we'll use the following terminology:

- installation directory - this is the directory where you installed Confluence on each node.
- local home directory - this is the home or data directory on each node (in non-clustered Confluence this is simply known as the home directory).
- shared home directory - this is a directory that is accessible to all nodes in the cluster via the

On this page:

- Step 1 Back up
- Step 2 Stop the cluster
- Step 3 Create a shared home directory
- Step 4 Upgrade the first node
- Step 5 Upgrade the other nodes
- Step 6 Start Confluence and check cluster connectivity

same path. If you're upgrading from Confluence 5.4 or earlier you'll create this directory as part of the upgrade.

Step 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.

Step 2 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.

Step 3 Create a shared home directory

⚠ If you are upgrading an **existing Confluence Data Center** instance (Confluence 5.6 or later), you can skip this step, as you already have a Shared Home directory.

To set up your Shared Home directory:

1. Create a directory that is accessible to all cluster nodes via the same path. This will be your **shared home** directory.
2. Edit `confluence.cfg.xml` in the home directory on the first node and add a new property called `confluence.cluster.home` with the path of the shared home directory as the value. Example:

```
<property
name="confluence.cluster.home">/mnt/confluence-shared-home</property>
```

3. Move all the files/directories from the local home directory the first node to the new shared home directory **except** for the following:
   - config
   - confluence.cfg.xml
   - index
   - temp
   - bundled-plugins
   - plugin-cache-*
   - plugins-cache
   - plugins-osgi-cache
   - plugins-temp

   Remove the moved files/directories from the local home directories on all other nodes.

Step 4 Upgrade the first node

To upgrade the first node:

1. Extract (unzip) the files to a directory (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.properties` file to point to the **existing local home** directory on that node.
3. Copy the jdbc driver jar file from your existing Confluence installation directory to `confluence/WEB-INF/lib` in your new installation directory.
   The jdbc driver will be located in either the `<Install-Directory>/common/lib` or `<Installation-Directory>/confluence/WEB-INF/lib` directories.
4. Copy any other immediately required customisations 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)
5. Start Confluence, and and confirm that you can log in and view pages before continuing to the next

step.

You should now reapply any additional customisations from the old version to the new version, before upgrading the remaining nodes.

Step 5 Upgrade the other nodes

To upgrade the next node in your cluster:

1. Stop Confluence on the first node.
2. Copy the installation directory and home directory from the first node to the next node.
   If the path to the local home directory is different on this node, edit the `confluence-init.propert ies` to point to the correct location.
3. Start Confluence, and and confirm that you can log in and view pages before continuing with the next node.

Repeat this process for each remaining node.

Step 6 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).

The Cluster monitoring console (

> **General 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 means you can add and remove nodes from your Data Center cluster at any time.

Adding a node

To add a node:

1. Copy the installation directory and local home directory from the stopped node to your new node.
2. 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.

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.nod e.name` system property to set the node identifier on each node so that it's easier for your users and administrators to read.

See Configuring System Properties for more information on how to set the system property.

Change Node Discovery from Multicast to TCP/IP

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're switching from a multicast configuration to TCP/IP (unicast), however, you'll need to edit the `confluence.cfg.xml` file 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 the `confluence.cfg.xml` file in each local home directory)

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 the confluence.cfg.xml file:

```
<property name="confluence.cluster.address">[multicast IP]</property>
<property name="confluence.cluster.join.type">multicast</property>
```

If both lines exist in the file, change them to the lines below; where the `confluence.cluster.address` property exists, but there's no reference to the `confluence.cluster.join.type` property, 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.

You can now restart your cluster nodes.

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'`	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. This key is only used by confluence if `confluence.cluster.join.type` is set to `tcp_ip`

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 Confluence installer, you do not need to install Java manually.

Check the Supported Platforms page to find out which Java versions are supported for 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.

˅ For Windows: (click to expand)

Installing the JDK on Windows

Before you start, go to **Control Panel** > **Programs and Features** to check whether a JDK is already installed.

To install the JDK on Windows:

1. Download the appropriate JDK version.
   Check the Supported Platforms page to find out which JDK / JRE versions are supported for your version of Confluence.
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 `JAVA_HOME` 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.

˅ For Linux: (click to expand)

Installing the JDK on Linux

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/JDK7` or `/usr/lib/jvm/java-7`, 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.
- If you have installed an unsupported JDK and you want to use SSL then you need to install the Sun JSSE package.

To install the JDK on Linux:

1. Download the appropriate JDK version.
   Check the Supported Platforms page to find out which JDK / JRE versions are supported for your version of Confluence.
2. Run the Java installer. Detailed installation instructions are provided on Oracle's website.
3. Open a shell console and type `echo $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

After you've installed Java in Windows, you must set the `JAVA_HOME` environment variable to point to the Java installation directory.

> This information is only relevant if you're installing Confluence manually on a Windows server. If you're using the installer we'll do this part for you.

If you installed the Java Development Kit (JDK) you'll be setting the `JAVA_H OME` environment variable. If you installed the Java Runtime Environment (JRE) you will follow the same steps but set the `JRE_HOME` environment variable instead.
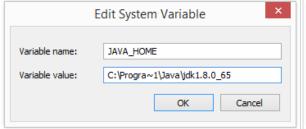
Related pages
- Starting Tomcat as a Windows Service
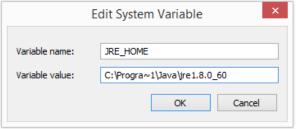- Installing Confluence in Linux

Set the JAVA_HOME Variable

To set the JAVA_HOME variable:

1. Find out where Java is installed. If you didn't change the path during installation, it will be something like this:
   `C:\Program Files\Java\jdk1.8.0_65`
2. In Windows 7 right click **My Computer** and select **Properties** > **Advanced**.
   In Windows 8 go to **Control Panel** > **System** > **Advanced System Settings**.
3. Click the **Environment Variables** button.
4. Under **System Variables**, click **New.**
5. In the **Variable Name** field, enter:
   - `JAVA_HOME` if you installed the JDK (Java Development Kit)
     or
   - `JRE_HOME` if you installed the JRE (Java Runtime Environment)
6. In the **Variable Value** field, enter your JDK or JRE installation path.

   If the path contains spaces, use the shortened path name, for example `C:\Progra~1\Java\jdk1.8 .0_65`)

Edit System Variable	✕
Variable name:	JAVA_HOME
Variable value:	C:\Progra~1\Java\jdk1.8.0_65
	OK Cancel

Edit System Variable	✕
Variable name:	JRE_HOME
Variable value:	C:\Progra~1\Java\jre1.8.0_60
	OK Cancel

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 any 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 even after reopening the command window, restart Windows.

Creating a Dedicated User Account on the Operating System to Run Confluence

ⓘ This step is optional if you are evaluating Confluence, but should be mandatory for Confluence installations used in production. If you have used the Confluence installer on Linux, this user will be created automatically.

A dedicated user should be created to run Confluence, because Confluence runs as the user it is invoked under and therefore can potentially be abused. For example:

- If your operating system is *nix-based (for example, Linux or Solaris), type the following in a console:
  `$ sudo /usr/sbin/useradd --create-home --comment "Account for running Confluence" --shell /bin/bash confluence`

- If your operating system is Windows:
  1. Create the dedicated user account by either:
     - Typing the following at the Windows command line:
       ```
       > net user confluence mypassword /add /comment:"Account for running
       Confluence"
       ```
       (This creates a user account with user name 'confluence' and password 'mypassword'. You should choose your own password.)
     - Opening the Windows 'Computer Management' console to add your 'confluence' user with its own password.
  2. *(Optional)* Use the Windows 'Computer Management' console to remove the 'confluence' user's membership of all unnecessary Windows groups, such as the default 'Users' group.
     ✅ If Windows is operating under Microsoft Active Directory, ask your Active Directory administrator to create your 'confluence' account (with no prior privileges).

Ensure that the following directories can be read and written to by this dedicated user account (e.g. 'confluence'):

- The Confluence Installation Directory, particularly these sub-directories:
  - `logs`
  - `temp`
  - `work`
- Your Confluence Home directory.

See also Best Practices for Configuring Confluence Security.

Confluence Setup Guide

Before running the Confluence Setup Wizard, as described below, you should have already completed installing Confluence.

When you access Confluence in your web browser for the first time, you will see the **Confluence Setup Wizard**. This is a series of screens which will prompt you to supply some default values for your Confluence site. It will also offer some more advanced options for setting up data connections and restoring data from a previous installation.

On this page:
- 1. Start the setup wizard
- 2. Choose your installation type and add-ons
- 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

1. Start the setup wizard

1. Start Confluence (if it is not already running)
   For Windows, go to **Start** > **Programs** > **Conf luence** > **Start Confluence Server**.
   Or, run the start-up script found in the `bin` fol der of your installation directory:
   - `start-confluence.bat` for Windows.
   - `start-confluence.sh` for Linux-based 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.

2. Choose your installation type and add-ons

In this step, you'll choose whether you want a trial or a production installation.

- **Trial installation**
  Set up Confluence with the embedded H2 database, default settings and sample content to get you started.

You'll need to migrate to a supported external database before using Confluence as a production system. This option is recommended if you're just taking Confluence for a test drive.

- **Production installation**
  Set up Confluence with your own external database. This option is recommended for setting up Confluence in a production environment.

If you've purchased a license bundle that includes Questions or Team Calendars for Confluence (or if you're keen to try these add-ons) you can get the setup wizard to install these add-ons automatically - choose the add-ons you want to automatically install.

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 to my.atlassian.com, or to purchase a new commercial license go to my.atlassian.com/purc hase.

If you selected a **Trial installation** in the previous step, Confluence will generate your license and then create the embedded database. 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 the supported platforms list to confirm that your chosen database and version is supported.
- See database configuration for information on setting up your database, including UTF-8 character encoding requirements.
- If you are using Confluence as a production system you **must** use an external database.
- The embedded H2 database option is available for evaluating or demonstrating Confluence, but should not be used for production use. If you choose this option, you'll need to migrate to an external database later on.

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 and reliability. A guide on how to migrate your data can be found here.

Embedded Database

---

5. Production installation: external database

> **Before you Start**
> - **Character encoding:**
>   - We strongly recommend that character encoding is consistent across your database, application server and web application, and that you use **UTF-8** encoding.
>   - Before setting up your database, please read about configuring character encoding.
> - **Database name:** When creating a new external database, give it the name '**confluence**'.

Choose how you want Confluence to connect to your database either via a direct JDBC connection or via a server-managed datasource connection.

Screenshot: Connection options

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. See Dat abase JDBC Drivers for details.
- **Database URL** – The JDBC URL for the database you will be connecting to. This will depend on the JDBC driver, and will be found in the documentation for your database.
- **User Name** 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, see Database 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:co mp/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.

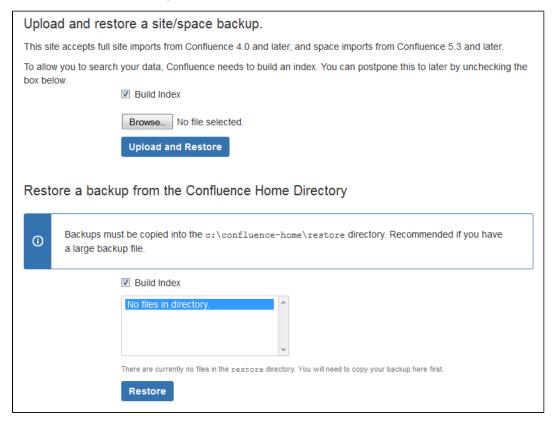**6. Production installation: load content**

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 a manual backup file 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>/restore` directory. Your backup 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 Desk.

- If you do not have a JIRA application installed, or if you would prefer to set up external user management later, choose **Manage 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**

Connect to JIRA

Fill in the details of the JIRA server you wish to retrieve user and group information from. You will need a valid administrator account on that JIRA server. Learn more about configuring JIRA integration.

JIRA Server Location

JIRA Base URL* `http://myjirasite:8080/`

For example: http://jira.mycompany.com or http://mycompany.com/jira

JIRA Administrator Login

Confluence will need to log in to JIRA with administrative privileges to retrieve user information.

Username* `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 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.

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 **Start** to 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 Desk 3.0 or later.

On this page:

- Connecting to a JIRA application in the Setup Wizard
- Troubleshooting
- Note about trusted applications authentication

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

Fill in the details of the JIRA server you wish to retrieve user and group information from. You will need a valid administrator account on that JIRA server. Learn more about configuring JIRA integration.

JIRA Server Location

JIRA Base URL* | http://myjirasite:8080/

For example: http://jira.mycompany.com or http://mycompany.com/jira

JIRA Administrator Login

Confluence will need to log in to JIRA with administrative privileges to retrieve user information.

Username* | admin

This username must have system administrator rights on your JIRA server.

Password* | ••••••

Advanced Options

Confluence Base* URL | http://myconfluencesite:8090

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 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 your `config.xml` file. 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.IllegalStateException: 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.	Follow the steps below to remove the partial configuration and 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 have multiple servers running on the same host...

    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 `<baseUrl>/rest/applinks/1.0/manifest` (where `<baseurl>`is the base URL of your Confluence site). The application ID

will be listed in the `<ID>` 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. See Connecting to Crowd or JIRA for User Management.

Note about trusted applications authentication

When you connect to a JIRA application in the setup wizard, the setup procedure will configure **Trusted Applications authentication** for your application. Please be aware of the following security implications:

- Trusted applications are a **potential security risk**. When you configure Trusted Applications authentication, you are allowing one application to access another as any user. This allows all of the built-in security measures to be bypassed. Do not configure a trusted application unless you know that all code in the application you are trusting will behave itself at all times, and you are sure that the application will maintain the security of its private key.
- Only use Trusted Applications authentication if both your servers have the same set of users and the servers fully trust each other.

Upgrading Confluence

This document describes the procedure for upgrading to the latest version of Confluence on Windows or Linux.

Before you start

- Check your Confluence **licence is valid**.
  To check go to

  ⚙ ▾

  > **General Configuration** > **Support Tools** and make sure the license support period has not expired. If your support period has expired renew your licence and reapply it before proceeding with the upgrade.

- Read the Release Notes and Upgrade Notes for both the version you are upgrading to, and any versions you are skipping.

- Check that your **Java version, application server, operating system, database and browsers** are supported.
  See Supported Platforms and End of Support Announcements for Confluence to confirm latest requirements.

- Check the latest **database setup guide** for your database and ensure that the database is configured correctly. There may be new configuration requirements. See Database Configuration.

- Check the **compatibility of any add-ons**.
  Some add-ons may not yet be compatible with the latest version of Confluence. See Checking add-on compatibility with application updates to find out how to check this in the Universal Plugin Manager.

- Check for any **known issues** that might affect your instance.
  See the Confluence Knowledge Base for known issues for the version you are upgrading to and Database Troubleshooting for known issues related to supported external databases.

- Make a note of any **modifications** to your Confluence instance (for example layouts or a custom theme). Any customisation you wish to maintain will need to be reapplied after upgrading.

Step 1 Determine your upgrade path and method

Find the upgrade path that works for your current version of Confluence and the version you plan to upgrade to.

The following table will help you to determine the most efficient upgrade path from your current version to the latest versions of Confluence. To use the table find your current installed version of Confluence in the left column and follow the suggested path.

Your Version	Recommended upgrade path to Confluence 5
2.7 or earlier	Upgrade to **2.7.4** then upgrade to **3.5.17**, 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** then upgrade to the latest version of Confluence 5.
4.0 to 4.3	Upgrade directly to the latest version of Confluence 5.
5.0 to 5.9	Upgrade directly to the latest version of Confluence 5.

There are several factors that will determine the upgrade method you should use. If you:

- Are upgrading from a version **earlier than Confluence 3.5** you will need to upgrade manually. You can contact Support for help determining an appropriate upgrade path.
- Are moving to a **different operating system or file location** you should upgrade manually.
- Currently use the **EAR/WAR distribution** to deploy Confluence into an existing application server, you can only upgrade manually.
- Are moving to a **different database** you should upgrade using the installer first and then follow the procedure outlined in migrating to a different database.
- Have a **clustered** instance of Confluence you should follow the procedure outlined in Upgrading Confluence Data Center.
- Are using the **embedded (trial) database** you should migrate to a different database **before** upgrading. This database is supplied for evaluation purposes only and is not recommended for production environments. See embedded HSQLDB database for more information.

otherwise you should follow the instructions below and use the Windows or Linux installer to upgrade Confluence. XML backups **should not** be used to upgrade Confluence.

Step 2 Upgrade Confluence in a test environment

We strongly recommend you recreate your production instance and test the upgrade in this cloned environment.
1. Create a snapshot of your current production Confluence environment on a test server - see Moving Confluence Between Servers for how to do this.
2. Follow the steps below to perform the upgrade on your cloned environment.
3. Test all your unsupported add-ons (plugins) and any customisation (for example custom themes and layouts) with the new version before proceeding with the upgrade in your production environment.

Step 3 Back up

Before you begin the Confluence upgrade you must back up:
- **your external database**
  You must perform a manual backup of your external database and confirm that the backup was created properly. If you are unfamiliar with the backup-restore facilities of your database, you can simply restore the backup to a different system to ensure the backup worked before proceeding.
- **your Confluence Home directory**
  The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. The location of the Home directory is stored in a configuration file called `confluence-init.properties`, which is located inside the `confluence/WEB-INF/classes` directory in your Confluence Installation directory.
  if you store attachments outside the Confluence Home directory, you should also backup your attachments directory.
- **the Confluence installation directory**
  This is where the Confluence application files and libraries were unpacked (unzipped) when Confluence was originally installed.

The installation wizard will back up your Confluence directories as part of the installation process, but you should also back these directions up manually before starting the upgrade.

Step 4 Upgrade Confluence in your production environment

1. Download the appropriate Windows or Linux installer from the Confluence Download Center.
2. Start the installer:
   - Windows Users: run the .exe file.
     If prompted to allow the upgrade wizard to make changes to your computer, choose '**Yes**'. If you do not, the installation wizard will have restricted access to your operating system and any subsequent installation options will be limited.
   - Linux users: open a Linux console and change directory (`cd`) to the '.bin' file directory and execute the '.bin' file.
     If the '.bin' file is not executable after downloading it, make it executable, for example `chmod a+x atlassian-confluence-5.6.1-x64.bin` (specify the exact filename of the installer you downloaded).
3. The installation wizard will guide you through the upgrade process. Some things to note:
   a. When prompted choose **Upgrade an existing Confluence installation** (for Linux users this is option 3).
   b. Verify that the **Existing Confluence installation directory** suggested by the wizard is correct. This is especially important if you have multiple Confluence installations running on the same machine.
   c. At the 'Back up Confluence directories' step, ensure '**Back up Confluence home**' is selected. This will create a .zip backup of the Confluence home and installation directories. **This is strongly recommended**.
   d. The installation wizard will notify you of customisations in the Confluence Installation directory. Make a note of these before proceeding as you will need to manually reapply these customisations after the upgrade is complete.
   e. If you have not already done so, the wizard will prompt you to backup your external database and check plugin compatibility. If your database does not support online backups you will need to stop the installation wizard at this point.
4. The wizard will shut down your Confluence instance and proceed with the upgrade. Once complete, it will restart Confluence and you can then launch Confluence in your browser to confirm the upgrade was successful.

During the upgrade the wizard will migrate following from your existing Confluence installation:

- TCP port values in your `server.xml` file.
- Custom values in your `confluence-init.properties` (confluence.home property) and `setenv.sh`/ `setenv.bat` files (`CATALINA_OPTS` or `JAVA_OPTS` parameters)

If you are using an Oracle or MySQL database, you'll need to copy the jdbc driver jar file from your existing Confluence installation directory to `confluence/WEB-INF/lib` in your new installation directory.

> Other configurations or customisations (including any other modifications in the `server.xml` file) are **n ot migrated** during the upgrade and need to be reapplied manually. See below for more information.

Additional steps when customisations are present

The installation wizard's ability to notify you about customisations will depend on how your existing Confluence instance was installed:

- If your current Confluence instance was installed using the installer, the wizard will check the entire Confluence Installation directory.
- If your current Confluence instance was installed manually it will only check the `confluence` subdirector y of the Confluence Installation directory. The installation wizard will **not** notify you of modifications in any other directory, for example modifications to start-up scripts under the `bin` directory or modifications to the `server.xml` file (such as an SSL configuration).

If customisations are present you will need to perform the following steps after the upgrade is complete:

1. Stop the upgraded Confluence instance.
2. Reapply the customisations to the relevant files in the upgraded Confluence Installation directory.
3. Restart the upgraded Confluence instance.

We **strongly recommend** you test your customisations in a test instance prior to upgrading your production

instance as changes may have been made to Confluence that make your customisations unsuable.

Troubleshooting

> **Did something go wrong?**
>
> If you need to retry the upgrade, **you must restore your pre-upgrade backups first.** Do not attempt to run an upgrade again, or start the older version of Confluence again after an upgrade has failed.

Some common issues encountered while upgrading:

- **Cannot proceed with upgrade because license has expired**
  If your licence has expired and was not renewed and reapplied before upgrading you will receive errors during the upgrade process. See upgrading beyond current license period for information on how to resolve this problem.
- **Unable to proceed with upgrade because of a conflict with anti virus**
  Some anti-virus or other Internet security tools may interfere with the Confluence upgrade process and prevent the process from completing successfully, particularly if you run Confluence as a Windows service. If you experience or anticipate experiencing such an issue with your anti-virus / Internet security tool, disable this tool first before proceeding with the Confluence upgrade.
- **Database does not support online backups**
  The upgrade wizard will prompt you to backup your database using your database's backup utilities. If your database does not support online backups, stop the upgrade process, shut down Confluence, perform your database backup and then run the installer again to continue with the upgrade.
- **Upgrade is taking a very long time**
  If you have a very large database (i.e. database backups take a very long time to complete), setting the `confluence.upgrade.recovery.file.enabled` system property to false will speed up the upgrade process. It should be used only when there is a process to back up database and verify the backup before performing an upgrade.

You can also refer to the Upgrade Troubleshooting guide in the Confluence Knowledge Base, or check for answers from the community at Atlassian Answers.

Upgrading Beyond Current Licensed Period

This page explains the recovery process should you mistakenly try to upgrade your Confluence installation to a version beyond your current license entitlement.

Related pages:

- Upgrading Confluence
- Working with Confluence Logs
- Confluence Administrator's Guide

License warnings

During an upgrade an obvious indication that your license has expired can be found in your log file. You will see a 'WARN' level entry similar to this:
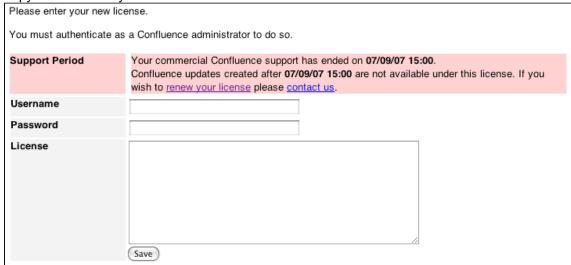
```
[confluence.upgrade.impl.DefaultUpgradeManager] isUpgradeAllowed
Your license is now outside of its support period.
You need to renew the license before you can upgrade to this version
of Confluence.
```

When you try to access the Confluence site in your browser, you will see the following warning screen:

You cannot access Confluence at present. Look at the table below to identify the reasons

Time	Level	Type	Description	Exception
2008-02-04 10:51:04	null	(EventType: upgrade)	Cannot proceed with upgrade. Your current license does not entitle you to upgrade to this version of Confluence. Please check that the support period of your license has not expired or that you have the correct partner license. If you wish to renew your license, please contact sales@atlassian.com. If you have a new license, please enter it on this page and **restart**.	fatal

Updating the Confluence license

1. Contact Atlassian Sales to arrange for a new license to be issued, as instructed on the warning screen illustrated above.
2. Once you have received a suitable license, supply the license key to Confluence:
   - Click link given on the license warning screen, illustrated above.
   - You will first be asked to log in as a Confluence administrator.
   - Then you will be presented with a simplified license administration screen. Enter the credentials of a Confluence system administrator.
   - Copy the license key into the **License** field and choose **Save**.

Please enter your new license.

You must authenticate as a Confluence administrator to do so.

Support Period	Your commercial Confluence support has ended on **07/09/07 15:00**. Confluence updates created after **07/09/07 15:00** are not available under this license. If you wish to renew your license please contact us.
Username	
Password	
License	

[Save]

3. Restart Confluence to continue the upgrade.

Confluence Post-Upgrade Checks

This article provides a list of items for Confluence Administrators to check after a Confluence upgrade to ensure that it has completed successfully. This list is not exhaustive, but it does cover common upgrade mistakes.

Before You Begin

After you have completed an upgrade, you should see the following message in the `atlassian-confluence.log` file:

```
2010-03-08 08:03:58,899 INFO [main]
[atlassian.confluence.upgrade.AbstractUpgradeManager] upgradeFinished Upgrade
completed successfully
```

If you do not see the line in your log similar to the one above, this means that your upgrade has not completed successfully. Please check our Upgrade Troubleshooting documentation to check for a suitable recommendation or fix. If there are no errors logged or if none of the errors are referenced in the the Troubleshooting Upgrades documentation, please contact Atlassian Support.

Upgrade Checklist

Below is a recommended list of items to check after completing an upgrade.

1. Layout and Menu

Visit the Confluence dashboard and check that it is accessible and displays as expected. Test the different Internet browsers that you have in use in your environment. In addition, confirm that the layout appears as expected and that the menus are clickable and functioning.

2. Search

Try searching for content, for example pages, attachments or user names. Check that the expected results are returned.

3. Permissions

Confirm that you can visit a page that has view restrictions, but you have permission to view. Confirm that you can edit a page that has edit restrictions but you have permission to edit. Make sure that the permissions of child pages are functioning as well. Involve as many space administrators as possible to confirm they are working. Confirm that anonymous or forbidden users cannot access or modify restricted pages.

4. Attachments

Confirm that attachments are accessible and searchable.

5. Plugins

Outdated third-party plugins can cause upgrade failure. Quite often, they will just be incompatible and simply do not work anymore. If you discover that your plugin is no longer working, please check for the latest version for your plugin in the The Atlassian Marketplace or 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 customised caches as described here that there is enough memory on the server for these caches should they reach maximum capacity.

Watching the migration logs during the upgrade

To monitor the progress of a site migration you should watch the output in the application log.

Typical logging progress will be shown by multiple log entries at the INFO level of the following format:

```
WikiToXhtmlMigrationThread-n - Migrated 2500 of 158432 pages, this batch
migrated 500/500 without error
```

There may be a wide array of messages logged from each individual page but any errors are also collected for display in a single migration report once all content has been processed. Here is a typical example of such a report:

```
Wiki to XHTML Exception Report:
Summary:
  0 settings values failed.
  0 PageTemplates failed.
  2 ContentEntityObjects failed.
Content Exceptions:
  1) Type: page, Id: 332, Title: Release Notes 1.0b3, Space: DOC -
Confluence 4.0 Beta. Cause:
com.atlassian.confluence.content.render.xhtml.migration.exceptions.Unkno
wnMacroMigrationException: The macro link is unknown.. Message: The
macro link is unknown.
  2) Type: comment, Id: 6919, Title: null, Global Scope. Cause:
com.atlassian.confluence.content.render.xhtml.migration.exceptions.Unkno
wnMacroMigrationException: The macro mymacro is unknown.. Message: The
macro mymacro is unknown.
```

Each entry in the report will identify the content that caused migration exceptions as well as displaying the exceptions themselves.

In almost all cases any content reported as errored will have been migrated to the new XHTML-based storage format, but will actually consist of wiki markup content wrapped within an XML 'unmigrated-wiki-markup' macro. This content will still be viewable in Confluence and editable within the new Confluence Editor.

However, in some cases a batch of content may actually have completely failed to 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.Migra
tionException: java.util.concurrent.ExecutionException:
org.springframework.transaction.UnexpectedRollbackException: Transaction
rolled back because it has been marked as rollback-only
```

Confluence provides no further report about this scenario and will also allow Confluence to restart as normal without retrying a migration. If a user tries to view any such unmigrated content they will see an exception similar to this:

```
java.lang.UnsupportedOperationException: The body of this
ContentEntityObject ('Page Title') was 'WIKI' but was expected to be
'XHTML'
```
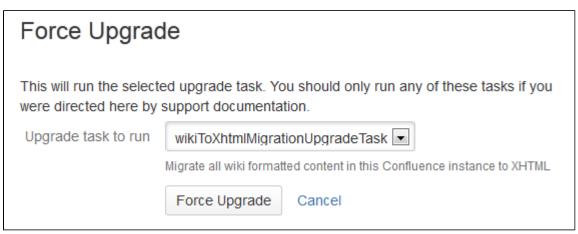
The solution is to ensure you manually re-run the site migration after the restart.

Re-running the migration – for content that completely failed the migration

A Confluence Administrator can restart the site migration if there was any content that failed migration (see previous section). Only the content that is still formatted in wiki markup will be migrated, so typically a re-migration will take less time than the original migration.

To manually re-run migration:

1. Open this URL in your browser: `<Confluence Address>/admin/force-upgrade.action`
2. Select **wikiToXhtmlMigrationUpgradeTask** in the **Upgrade task to run** dropdown list.
3. Choose **Force Upgrade**.

Re-attempting the migration – for content in 'unmigrated-wiki-markup' macro

The previous section was about dealing with the exceptional circumstance where certain content was left completely unmigrated. The most common migration problem is that the content was migrated but remains formatted as wiki markup on the page, within the body of an 'unmigrated-wiki-markup' macro. Any content which is referenced in the migration report will be found in this state. This content is still viewable and editable but since it is wiki markup it cannot be edited using the full feature set of the rich text editor.

The most common reason for content to be in this state is that the page contains an unknown macro, or a macro that is not compatible with Confluence 4.x.

There are two possible fixes for this situation:

1. Install a version of the macro that is compatible with Confluence 4.x. See Plugin Development Upgrade FAQ for 4.0 .
2. Edit the page and remove the problematic macro.

Regardless of the solution you choose, you can then force a re-migration of all the content (including content in templates) that was left wrapped in an 'unmigrated-wiki-markup' macro. This feature is found at `<Confluence Address>/admin/unmigratedcontent.action`

Update content with incompatible macros

Confluence has detected that there are **0 pages** with macros that are not yet Confluence 4+ compatible. To ensure backwards compatibility, these macros are still being rendered as wiki markup when editing your pages.

If you have recently updated plugins, you should update your content to ensure that any macros that are not Confluence 4 compatible become compatible. You may have to run the update several times as you update incompatible macros.

Update Check

> ⓘ **Update not required**
>
> You have not installed any new plugins since your last content upgrade. You do not need to run this upgrade unless you have been advised to by Atlassian Support staff.
>
> **Note:** Once an upgrade has commenced you will not be able to pause or undo the upgrade. An update can severely affect the performance of your instance, we recommend you conduct this update during a quiet time. Users editing a page as it is updated may receive notice of a conflicting edit.
>
> Update Content

Notes

We refer to the Confluence storage format as 'XHTML-based'. To be correct, we should call it XML, because the Confluence storage format does not comply with the XHTML definition. In particular, Confluence includes custom elements for macros and more. We're using the term 'XHTML-based' to indicate that there is a large proportion of HTML in the storage format.

Migration of Templates from Wiki Markup to XHTML-Based Storage Format

If you are **upgrading to Confluence 4.3 or later from an older version** (from Confluence 4.2.x or earlier) then as part of the upgrade an automatic migration of your page templates will take place. This is a non-destructive process. Your existing content is not overwritten. Instead, the migration process will create a new version of each space template and each global template on your Confluence site. The new version will use the new XHTML-based storage format, so that you can edit the template in the Confluence rich text editor.

Note: Nevertheless, you must be sure to perform a backup of your database and home directory prior to starting the new version of Confluence, as we recommend for any Confluence upgrade.

Watching the migration logs during the upgrade

To monitor the progress of a site migration you should watch the output in the application log.

A typical logging progress will be shown by multiple log entries at the INFO level of the following format:

```
WikiToXhtmlMigrationThread-n - Migrated 22 of 29 PageTemplates.
```

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.Unkno
wnMacroMigrationException: The macro link is unknown.. Message: The
macro link is unknown.
 2) Type: comment, Id: 6919, Title: null, Global Scope. Cause:
com.atlassian.confluence.content.render.xhtml.migration.exceptions.Unkno
wnMacroMigrationException: The macro mymacro is unknown.. Message: The
macro mymacro is unknown.
```

Each entry in the report will identify the content that caused migration exceptions as well as displaying the exceptions themselves.

In almost all cases any content reported as errored will have been migrated to the new XHTML-based storage format, but will actually consist of wiki markup content wrapped within an XML 'unmigrated-wiki-markup' macro. This content will still be viewable in Confluence and editable within the 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.Migra
tionException: java.util.concurrent.ExecutionException:
org.springframework.transaction.UnexpectedRollbackException: Transaction
rolled back because it has been marked as rollback-only
```

Confluence provides no further report about this scenario and will also allow Confluence to restart as normal without retrying a migration. If a user tries to view or edit an unmigrated template, the wiki template editor will be used.

The solution is to manually re-run the site migration after the restart, as described below.
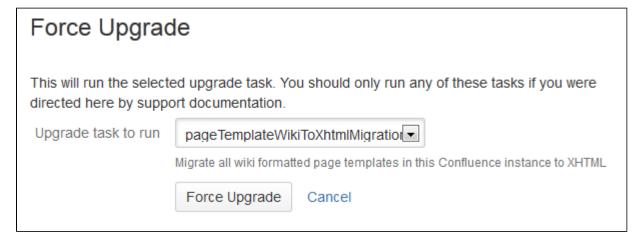
Re-running the migration

A Confluence administrator can restart the template migration if any templates have failed the migration (see previous section). Only the templates that are still formatted in wiki markup will be migrated again. Typically, a re-migration will take less time than the original migration.

To manually re-run the migration:

1. Open this URL in your browser: `<Confluence Address>/admin/force-upgrade.action`
2. Select **pageTemplateWikiToXhtmlMigrationUpgradeTask** in the **Upgrade task to run** dropdown list.
3. Choose **Force Upgrade**.

Screenshot: The 'Force Upgrade' screen in the Confluence administration console

Notes

We refer to the Confluence storage format as 'XHTML-based'. To be correct, we should call it XML, because the Confluence storage format does not comply with the XHTML definition. In particular, Confluence includes custom elements for macros and more. We're using the term 'XHTML-based' to indicate that there is a large proportion of HTML in the storage format.

Upgrading Confluence Manually

This document describes the procedure for upgrading to the latest version of Confluence on Windows or Linux manually (not using the upgrade wizard). See Upgrading Confluence to upgrade using the installer and upgrade wizard.

Before you start

- Check your Confluence **licence is valid**.
  To check go to

  ⚙ ▾

   > **General Configuration** > **Support Tools** and make sure the license support period has not expired. If your support period has expired renew your licence and reapply it before proceeding with the upgrade.

- Read the Release Notes and Upgrade Notes for both the version you are upgrading to, and any versions you are skipping.

- Check that your **Java version, application server, operating system, database and browsers** are supported.
  See Supported Platforms and End of Support Announcements for Confluence to confirm latest requirements.

- Check the latest **database setup guide** for your database and ensure that the database is configured correctly. There may be new configuration requirements. See Database Configuration.

- Check the **compatibility of any add-ons**.
  Some add-ons may not yet be compatible with the latest version of Confluence. See Checking add-on compatibility with application updates to find out how to check this in the Universal Plugin Manager.

- Check for any **known issues** that might affect your instance.
  See the Confluence Knowledge Base for known issues for the version you are upgrading to and Databas e Troubleshooting for known issues related to supported external databases.

- Make a note of any **modifications** to your Confluence instance (for example layouts or a custom theme). Any customisation you wish to maintain will need to be reapplied after upgrading.

Step 1 Determine your upgrade path and method

Find the upgrade path that works for your current version of Confluence and the version you plan to upgrade to.

The following table will help you to determine the most efficient upgrade path from your current version to the latest versions of Confluence. To use the table find your current installed version of Confluence in the left column and follow the suggested path.

Your Version	Recommended upgrade path to Confluence 5
2.7 or earlier	Upgrade to **2.7.4** then upgrade to **3.5.17**, 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** then upgrade to the latest version of Confluence 5.
4.0 to 4.3	Upgrade directly to the latest version of Confluence 5.
5.0 to 5.9	Upgrade directly to the latest version of Confluence 5.

There are several factors that will determine the upgrade method you should use. If you:

- Are moving to a **different database** you should upgrade first and then follow the procedure outlined in Mi grating to Another Database.
- Are running Confluence in a **cluster** you should follow the procedure outlined in Upgrading Confluence Data Center.
- Are using the **embedded (trial) database** you should migrate to a different database **before** upgrading. This database is supplied for evaluation purposes only and is not recommended for production environments. See Embedded H2 Database for more information.

Step 2 Upgrade Confluence in a test environment

We strongly recommend you recreate your production instance and test the upgrade in this cloned environment.
1. Create a snapshot of your current production Confluence environment on a test server - see Moving Confluence Between Servers for how to do this.
2. Follow the steps below to perform the upgrade on your cloned environment.
3. Test all your unsupported add-ons (plugins) and any customisation (for example custom themes and layouts) with the new version before proceeding with the upgrade in your production environment.

Step 3 Back up

Before you begin the Confluence upgrade you must back up:
- **your external database**
  You must perform a manual backup of your external database and confirm that the backup was created properly. If you are unfamiliar with the backup-restore facilities of your database, you can simply restore the backup to a different system to ensure the backup worked before proceeding.
- **your Confluence Home directory**
  The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. The location of the Home directory is stored in a configuration file called `confluence-init.properties`, which is located inside the `confluence/WEB-INF/classes` directory in your Confluence Installation directory.
  if you store attachments outside the Confluence Home directory, you should also backup your attachments directory.
- **the Confluence installation directory**
  This is where the Confluence application files and libraries were unpacked (unzipped) when Confluence

was originally installed.

Step 4 Upgrade Confluence in your production environment

1. Download the appropriate archive file from the Confluence Download Center.
2. Shut down Confluence.
3. Extract (unzip) the files to a directory (this will be 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 7Zip or Winzip.
4. Update the following line in the `<Installation-Directory>\confluence\WEB-INF\classes\confluence-init.properties` file to point to your **existing** Confluence home directory.

   ```
   # confluence.home=c:/confluence/data
   ```

   For example if your existing Confluence home directory is `c:\confluence\myhome` the file will read:

   ```
   confluence.home=c:/confluence/myhome
   ```

5. Copy the JDBC driver jar file from your existing Confluence installation directory to `confluence/WEB-INF/lib` in your new installation directory.
   The JDBC driver will be located in either the `<Install-Directory>/common/lib` or `<Installation-Directory>/confluence/WEB-INF/lib` directories.
6. There are some additional steps you make need to take if:
   - you are running Confluence as a **Windows Service**
     - ⌄ Click here to expand...
       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 use a **JIRA** application or **LDAP** for user management
     - ⌄ Click here to expand...
       If you use JIRA or LDAP for user management, copy the following files from your existing installation directory to your new installation directory:

       - `<Installation-Directory>/confluence/WEB-INF/classes/osuser.xml`
       - `<Installation-Directory>/confluence/WEB-INF/classes/atlassian-user.xml`

   - you use **Crowd** for user management
     - ⌄ Click here to expand...
       If you are using Crowd for user management, copy the following files from your existing installation directory to your new installation directory:

       - `<Installation-Directory>/confluence/WEB-INF/classes/osuser.xml`
       - `<Installation-Directory>/confluence/WEB-INF/classes/atlassian-user.xml` (if you are upgrading from Confluence 2.2 or later).
       - `<Installation-Directory>/confluence/WEB-INF/classes/crowd.properties`

   - you are upgrading from **Confluence 2.5** or earlier

> ⌄ Click here to expand...
>
> If you are upgrading from Confluence 2.5.5 or earlier and are copying the existing `atlassian-user.xml` file from your previous instance, make sure the hibernate cache parameter in this file is enabled, to avoid performance issues.
>
> ```
> <hibernate name="Hibernate Repository"
> key="hibernateRepository" description="Hibernate
> Repository" cache="true" />
> ```
>
> If you use Crowd for your user management, you do not need to do this.

- You are running Confluence on a different port (not the default 8090)
  - ⌄ Click here to expand...
    If you are not running Confluence on port 8090 update `<Installation-Directory>\conf\server.xml` file to include your ports

7. Start your new Confluence. You should not see the setup wizard.

Additional steps if further customisations are present

If you have customised Confluence (such as an SSL configuration in the `server.xml` file, or Cᴀᴛᴀʟɪɴᴀ_OPTS or JAVA_OPTS parameters in your `confluence-init.properties` file), you'll need to perform the following steps after the upgrade is complete:

1. Stop the upgraded Confluence instance.
2. Reapply the customisations to the relevant files in the newly upgraded Confluence Installation directory.
3. Restart the upgraded Confluence instance.

We **strongly recommend** you test your customisations in a test instance prior to upgrading your production instance as changes may have been made to Confluence that make your customisations unsuable.

Troubleshooting

> **Did something go wrong?**
>
> If you need to retry the upgrade, **you must restore your pre-upgrade backups first.** Do not attempt to run an upgrade again, or start the older version of Confluence again after an upgrade has failed.

You can also refer to the Upgrade Troubleshooting guide in the Confluence Knowledge Base, or check for answers from the community at Atlassian Answers.

Supported Platforms

This page describes the supported platforms for Confluence. Please review them before installing Confluence. The information on this page applies to **Confluence 5.9**.

Further information:

- End of support for various platforms and browsers when used with Confluence: End of Support Announcements for Confluence.
- More information about these supported platforms and hardware requirements: System Requirements.
- Go to
  ⚙ ▾
  > **General Configuration** > **Support Tools** to check your instance health. It looks at things like your licence validity, Tomcat version, basic database setup and more.

Related pages:

- Confluence Installation Guide
- Confluence Setup Guide

- Server Hardware Requirements Guide
- Supported Platforms FAQ

Key : ✅ = Supported. ❌ = Not Supported

Java version	
Oracle JRE / JDK	1.8 There's a [known issue](#) with `Java 1.8.0_25` and `1.8.0_31`, and [another known issue](#) with `Java 1.8.0_45`. We don't recommend running Confluence on these versions.
Operating systems for Confluence server installation	
Microsoft Windows (including 64-bit)[1]	([Microsoft Supported Versions](#) only)
Linux / Solaris [1, 2, 10]	
Apple Mac OS X	Not supported as server. Supported as client platform.
Application servers	
Apache Tomcat	8.0.x [12]
Databases	
PostgreSQL	9.2, 9.3
MySQL [3][11]	5.5, 5.6 (5.6.16 and later only - [more info](#)) MariaDB and Percona Server are not supported. There's currently a known issue with `MySQL 5.5`: **CONF-37829** - Improve select query that used in Confluence for copying a page **RESOLVED**
Oracle	12c
Microsoft SQL Server	2012, 2014
H2 [4]	(this embedded database is provided for evaluating Confluence only)
Web browsers – desktop	
Microsoft Internet Explorer (Windows) [5, 6, 13]	10, 11
Mozilla Firefox (all platforms)	Latest stable version supported
Google Chrome (Windows and Mac) [7]	Latest stable version supported

Safari (Mac)	✅ Latest stable version supported
Web browsers – mobile	
Mobile Safari (iOS) [(8)]	Latest stable version supported
Android (Android) [(9)]	The default browser on Android 4.0.3 (Ice Cream Sandwich)
Chrome (Android and iOS) [(8, 9)]	Latest stable version supported

1. Confluence is a pure Java application and should run on this platform provided the JRE or JDK requirement is satisfied.
2. While some of our customers run Confluence on SPARC-based hardware, Atlassian only officially supports Confluence running on x86 hardware and 64-bit derivatives of x86 hardware.
3. Ensure that you configure your Confluence MySQL database to use the InnoDB storage engine as the MyISAM storage engine could lead to data corruption.
4. Confluence includes an embedded H2 database. This database is fine for evaluation purposes, but for production environments, we only support running Confluence with one of the supported external database listed on this page.
5. Drag and drop in Confluence works as expected with Internet Explorer 10 in 'desktop' mode. The 'modern' mode of Internet Explorer 10 does not support drag-and-drop.
6. Confluence is tested with these versions of Internet Explorer in standards-compliant rendering mode, not compatibility mode. Enabling compatibility mode may cause problems because it emulates older, unsupported rendering modes.
7. Chrome does not have WEBDAV / plugin support so features such as Edit in Word for attachments will not work. Please refer to CONF-23322 for more information.
8. Confluence does not support editing in Mobile Safari on iOS devices (such as iPhone and iPad). Please refer to CONF-19523 for information on the progress of this issue.
9. Confluence does not support editing on Android devices.
10. NFS mounts are not supported on Linux operating systems due to Lucene requirements.
11. We don't currently support MariaDB or Percona Server. See **CONF-36471** NEW and **CONF-29060** OPEN for more information.
12. We only support the Tomcat version that is bundled with your specific Confluence version.
13. You need to have Javascript enabled to install Confluence. Setting your Internet Explorer security level as 'High' will disable Javascript and prevent you from installing Confluence. See this knowledge base article for more information.

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 summarises 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
View File macros	With the release of Confluence 5.9 (announcement)
Status Updates	With the release of Confluence 5.9 (announcement)
Documentation theme	With the release of Confluence 6.0 (announcement)
Oracle 11.1 and 11.2	With the release of Confluence 5.9 (announcement)
Microsoft SQL 2008	With the release of Confluence 5.9 (announcement)

Internet Explorer 9	With the release of Confluence 5.9 (announcement)

⌄ See the full list of announcements...

Most recent announcements first.

- Deprecated macros for Confluence (13 November 2015)
- Discontinued features for Confluence (10 July 2015)
- Deprecated databases for Confluence (19 May 2015)
- Deprecated Tomcat platform 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)
- Deprecated Tomcat platform for Confluence (22 April 2014)
- Deprecated Databases for Confluence (2 December 2013)
- Deprecated Web Browsers for Confluence (24 September 2013)
- Deprecated Databases for Confluence (13 August 2013)
- Deprecated Tomcat platform for Confluence (29 August 2012)
- Deprecated Java platform for Confluence (6 August 2012)
- Deprecated Databases for Confluence (1 May 2012)
- Deprecated Databases for Confluence (13 March 2012)
- Deprecated Operating Systems for Confluence (21 July 2011)
- Deprecated Databases for Confluence (7 January 2011)
- Deprecated Web Browsers for Confluence (7 January 2011)
- Deprecated Databases for Confluence (12 October 2010)
- Deprecated Web Browsers for Confluence (12 October 2010)
- Deprecated Databases for Confluence (6 July 2010)
- Deprecated Web Browsers for Confluence (6 July 2010)
- Deprecated Databases for Confluence (24 March 2010)
- Deprecated Application Servers for Confluence (27 January 2010)
- Deprecated Java Platforms for Confluence (27 January 2010)
- Deprecated Web Browsers for Confluence (14 December 2009)

Deprecated macros for Confluence (13 November 2015)

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 Confluence 5.9, 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

CONF-39829 - Plans to remove the view file macros
OPEN

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 HipChat gives 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 **CONF-38253** - Plans to remove status updates **OPEN** .

Documentation theme

In order to better focus our development efforts on a single theme, we plan to remove the Documentation theme in **Confluence 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 favourite 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 step-by-step guide which 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 **CONF-38256** - Plans to remove the documentation theme **OPEN** .

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-announcemen t at atlassian dot com`.

End of Life Announcement for Database Support

Database	Support End Date
Microsoft SQL 2008	After Confluence 5.8.x
Oracle 11.1	
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.

Deprecated Tomcat platform for Confluence (1 May 2015)

This section announces the end of Atlassian support for Tomcat 7.0.x for Confluence. As previously 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 email `eol-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 affecting Tomcat 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 that Atlassian 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 Platforms page 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 supports Internet Explorer 9
- Confluence 5.8.x and earlier versions will continue to work on Internet 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 our Security Bugfix Policy. If at any time you become aware of a vulnerability we've missed, please report it as described in How to report a security issue.

If you have more questions or concerns regarding this announcement, please contact us at `eol-announcemen t 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-announcemen t 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.

Deprecated Tomcat platform for Confluence (22 April 2014)

This section announces the end of Atlassian support for Tomcat 6.0.x for Confluence.

End of support means that Atlassian will not fix bugs related to the specified version of Tomcat, past the support end date for your version of Confluence. The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Tomcat 6.0.x Support

Platform	Support End Date
Tomcat 6.0.x	When Confluence 5.6 is released, due in mid 2014

Tomcat 6.0.x notes:

- Confluence 5.5 is the last major version that will support Tomcat 6.0.x. The Confluence 5.5.x bug-fix releases will also continue to support Tomcat 6.0.x.
- Confluence 5.5.x and previously-released versions will continue to work with Tomcat 6.0.x. However, we will not fix bugs affecting Tomcat 6.0.x after the end-of-life date for your version of Confluence.
- Confluence 5.6 will not be tested with Tomcat 6.0.x.

Deprecated Databases for Confluence (2 December 2013)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcemen t 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 utilise modern browser technologies.

End of support means that Atlassian will not perform any maintenance on Confluence related to IE8 after the final release of Confluence 5.5.x, except for security related issues. In order to minimise the impact on you and the way your company uses Confluence, we have provided this announcement as early as possible, and hope that the subsequent 6 month period will give you adequate time to prepare for this change without disruption.

Atlassian will continue to support Internet Explorer 9 (IE9) and Internet Explorer 10 (IE10) as well as the latest versions of Chrome, Firefox and Safari. For further information, please refer to the Supported Platforms page. If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

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-announcemen t 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 affecting Tomcat 5.5.x after the end-of-life date for your version of Confluence.
- Confluence 5.0 will not be tested with Tomcat 5.5.x.

Deprecated Java platform for Confluence (6 August 2012)

This section announces the end of Atlassian support for Java 6 for Confluence. Please note that Oracle has announced the end of public updates for Java 6: Java SE 6 End of Public Updates Notice.

End of support means that Atlassian will not fix bugs related to the specified version of Java, past the support end date for your version of Confluence. The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement 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-announcemen t 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-announcemen t 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-announcemen t 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-announcemen t 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-announcemen t 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-announcemen t 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-announcemen t 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-announcemen t 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 standardise on Tomcat, because it is the most widely used application server in our user population. It is fast, robust, secure, well-documented, easy to operate, open source, and has a huge community driving improvements. It is the de facto industry standard, with several companies available that specialise in providing enterprise grade support contracts for it, ranging from customisations to 24/7 support.

Deprecated Java Platforms for Confluence (27 January 2010)

This section announces the end of Atlassian support for certain Java Platforms for Confluence.

We will **stop supporting the following Java Platforms**:

- From Confluence 3.3, due Q3 2010, support for Java Platform 5 (JDK/JRE 1.5) will end.

We are ending support for Java Platform 5, in line with the Java SE Support Roadmap (i.e. "End of Service Life" for Java Platform 5 dated October 30, 2009). We are committed to helping our customers understand this decision and assist them in updating to Java Platform 6, our supported Java Platform.

The details are below. Please refer to the Supported Platforms for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email `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.

Supported Platforms FAQ

Q: How does Atlassian choose which JRE versions, application servers and databases to support?

For application servers and databases, we try to pick a good cross-section of open source options and popular commercial platforms. We then choose which JRE versions to support based on the recommended environments for these servers.

Q: What is a supported platform?

A supported platform is one that:

- Confluence is regularly tested on during the development cycle
- One that is available within Atlassian for support technicians and developers to reproduce problems
- Bugs raised against it will be given a high priority

Supporting a platform means we know how to get Confluence running in that environment and can troubleshoot Confluence issues within it. It does not mean we have any particular expertise beyond that. As such, we may not be able to provide assistance with customising or tuning that application server or database. (Atlassian support is not a substitute for a good database administrator.)

Q: Can I get assistance with running Confluence on a platform that is *not* supported?

If you are running Confluence on an unsupported platform, then we can not guarantee providing any support for it. Furthermore, we will recommend that you switch to a platform which is supported.

Q: If you write your application to standards like J2EE, JDBC and SQL, doesn't that mean it should run on any compliant server?

Confluence is a complicated application and we commonly encounter interesting edge-cases where different servers have interpreted the specifications differently. Then again, each server has its own different collection of bugs.

Q: How can I get Atlassian to support Confluence on a new platform?

Supporting a new platform involves a significant investment of time by Atlassian, both up-front costs to set up new testing environments and fix any issues we might encounter and the ongoing costs involved in maintaining the application against this new environment in the future. As such, supporting a new platform is not something we will do unless we know there is significant demand for it.

Please be aware that your interest alone will not be enough for us to add support for your application server or database. We would need to see a significant number of votes on the issue raised in our public JIRA site or a significant level of interest in our forums, before considering supporting that platform.

Q: My organisation has standardised on an operating environment that Confluence does not support. What can I do?

In this situation, you have the following two options:

1. Run Confluence in the unsupported environment, with the caveats mentioned above.
2. Make an exception to your standardised operating environment and set up Confluence based on its supported platforms.

Migrating Confluence Between Servers

This page describes how to move Confluence between physical servers. It is distinct from other functions. It does not cover database migration, application server migration, or upgrading. Atlassian suggests doing each of these steps separately. See also:

- Upgrading Confluence
- Migrating to Another Database
- Switching to Apache Tomcat

On this page:

- How to Create a Test or Development Site
- Transferring Confluence To Another Server Using The Same Operating System
- Transferring Confluence To Another Server Using a Different Operating System
- Ensuring no contact with production systems
- Migrating from HTTPS to HTTP
- Notes

How to Create a Test or Development Site

Administrators may need to move a Confluence site from one server to another for upgrades or downtime. This page tells you how to copy a Confluence site from one server to another. For example, you may want to transfer your current production snapshot to a test server as permitted in the licence agreement.

- **Avoid upgrades while migrating.** If you are planning to switch databases, application servers or Confluence versions, firstly perform the application transfer in isolation, and test that it was successful before making other changes.
- Development licenses are available for any Commercial or Academic license. Create one or c ontact Atlassian for help.

Transferring Confluence To Another Server Using The Same Operating System

If the operating systems on both servers are the same, then the home and install folders can be copied straight into an identical external database and user management setup.

1. On the original server, create zips of the Confluence install and home directories. Copy the zips to the new server.
2. On the new server, unzip the install and home directories. Windows users should avoid unzipping with the Windows built-in extractor, instead use Winzip or the free 7Zip.

   If you are changing the location of the home directory, open the Confluence install\confluence\WEB-INF\classes directory and edit confluence-init.properties by changing the line starting with 'confluence.home='.

3. If you are using the EAR/WAR distribution, modify the location of your war file if need be. If using Tomcat, this is likely in /Conf/Catalina/localhost. You'll want to make sure the docbase attribute is pointing to the right location.
4. This next step is dependent on your database:
    - For users of the *internal database*, the database content is stored inside the home directory. You should switch to an external database after the transfer is successful. The internal database is for evaluation only and is not recommended for use in Production systems.
    - For *external databases stored on another server*: change the user account or datasource permissions so that the new server has the same network access permissions as the original. Then confirm from the new server that the hostname can be resolved and is listening for database connections on the expected port.
    - For *external databases hosted locally (ie. localhost)*: on the original server, create a manual database backup using a native db dump backup tool. Copy the database backup to the new server.
5. On the new server, install or upgrade the database version to match the original server.
6. Import the database backup.
7. Add a database user account with the same username and password as the original.
8. Provide the database user with the full access to the imported database.
9. Use a database administration tool to confirm that the user can login from the localhost.
10. This step depends on your database connection:
    a. If you use JDBC (the default option) to connect to the database, to modify any database connection information, go to the Confluence home directory and edit confluence.cfg.xml. The connection URL is set under hibernate.connection.url. **Ensure it does not point to your production database server.**
    b. If you use a data source, follow the instructions for your database type and ensure the data source points to the new database: PostgreSQL, MySQL, SQL Server or Oracle.
11. If you are using internal user management, skip this step. For users who have JIRA or LDAP integration, provide the new server with network or local access to the same hosts as the original. If this is a true test site, set up a test of your JIRA site or LDAP server so as not to disrupt production systems and change the server.xml or atlassian-user.xml files (Confluence 3.4 and below), or modify the directory settings in Confluence Admin > User Directories (Confluence 3.5 and above) to point to the appropriate test servers. Note that it might be acceptable to use a production connection here, as users won't be logging on to the test system in high volume.
12. If appropriate, make sure no emails are sent out from the test system.
13. If you have previously installed Confluence using the guided installer *and* plan on starting Confluence using the `startup` or `start-confluence` scripts in the Confluence `install/bin/` directory, check `setenv.sh` (Unix/Linux) or `setenv.bat` (Windows) in the same directory. If there is a `JRE_HO ME` set, ensure that the path to the JRE is up to date in regards to the new environment.
14. Start Confluence.
15. Go to Administration > License Details and add your development license key. You can generate one at http://my.atlassian.com. There are more details in How to get a Confluence Developer license.
16. If you configured Confluence as a Windows service, repeat those instructions.
17. Add your development license key.
18. Some customers have experienced problems with Confluence's search functions after performing a migration, or that the content of their {`recently-updated`} macro is not being updated correctly. Errors in the `atlassian-confluence.log` file corroborate such problems. Hence, to avoid these issues, it is strongly recommended that you perform a rebuild of your content indices after performing a migration.

Transferring Confluence To Another Server Using a Different Operating System

Migrating from Windows to Linux

You will need to replace the backslash with forward slash in the following lines in `confluence.cfg.xml`:

```
<property
name="attachments.dir">${confluenceHome}/attachments</property>
    <property
name="lucene.index.dir">${confluenceHome}/index</property>
    <property
name="webwork.multipart.saveDir">${confluenceHome}/temp</property>
```

Using database tools (preferred option)

If you are using the Production backup strategy, follow these steps:

1. Download the proper distribution (**the same one you have from your original site**) from the Download Archive.
2. Copy your Confluence *home* (not install) directory from your original server (even if it was a different OS).
3. If you are changing the location of the home directory, open the Confluence install\confluence\WEB-INF\classes directory and edit confluence-init.properties by changing the line starting with 'confluence.home='.
4. For external databases stored locally, on the original server, create a manual database backup using a native db dump backup tool.
5. Copy the database backup to the new server.
6. On the new server, install or upgrade the database version to match the original server.
7. Import the database backup.
8. Add a database user account with the same username and password as the original.
9. Provide the user with the full access to the imported database.
10. Use a database administration tool to confirm that the user can login from the localhost.
11. To modify any database connection information, go to the Confluence home directory and edit confluence.cfg.xml. The connection URL is set under hibernate.connection.url. **Ensure it does not point to your production database server.**
12. If you are using internal user management, skip this step. For users who have JIRA or LDAP integration, provide the new server with network or local access to the same hosts as the original.
13. Copy server.xml, atlassian-user.xml, osuser.xml, any patches, and any other customized files velocity or properties files. If this is a true test site, set up a test of your JIRA site or LDAP server so as not to disrupt production systems and change the server.xml or atlassian-user.xml files to point to the appropriate test servers. Note that it might be acceptable to use a production connection here, as users won't be logging on to the test system in high volume.
14. If appropriate, make sure no emails are sent out from the test system.
15. If you have previously installed Confluence using the guided installer *and* plan on starting Confluence using the `startup` or `start-confluence` scripts in the Confluence`install/bin/` directory, check `setenv.sh` (Unix/Linux) or `setenv.bat` (Windows) in the same directory. If there is a `JRE_HOME` set, ensure that the path to the JRE is up to date in regards to the new environment.
16. Start Confluence.
17. Go to **Administration > License Details** and add your development license key. You can generate one at http://my.atlassian.com. There are more details in How to get a Confluence Developer license.
18. If you configured Confluence as a Windows service, repeat those instructions.
19. Add your development license key.
20. Some customers have experienced problems with Confluence's search functions after performing a migration, or that the content of their {`recently-updated`} macro is not being updated correctly. Errors in the `atlassian-confluence.log` file corroborate such problems. Hence, to avoid these issues, it is strongly recommended that you perform a rebuild of your content indices after performing a migration.

Using XML data backups (only for small to medium sized installations)

Note: The XML export built into Confluence is not suited for the backup or migration of large data sets. There are a number of third party tools that may be able to assist you with the data migration. If you would like help in selecting the right tool, or help with the migration itself, we can put you in touch with one of the Atlassian Experts.

If you're not yet using the Production backup strategy, you can migrate Confluence to a different server machine by creating an XML data backup as usual, and then importing that to Confluence on the new server.

1. Create an XML data backup from Confluence as follows:
   a. Choose the **cog icon**

      , then choose **General Configuration** under Confluence Administration
   b. Select **Backup & Restore**.
   c. Check the **Backup Attachments** option and click **Backup**.
2. Identify the version of Confluence that you are currently using. This is displayed at the bottom of each Confluence page.
3. Download Confluence to the new server. Get the version of Confluence that you identified above, but for the operating system of the new server. You may be using either the latest Confluence version, or an older version.
4. Install Confluence on the new server.
5. Go to **Administration** > **License Details** and add your development license key. You can generate a license at http://my.atlassian.com. You can find more details in How to get a Confluence Developer license.
6. Restore your XML data backup from **Administration** > **Backup and Restore**.
7. If appropriate, make sure that no email contact can be made with the test system.
8. Re-install any third party add-ons (including Atlassian add-ons such as Team Calendars or Questions for Confluence), and apply relevant licenses.
9. Some customers have experienced problems with Confluence's search functions after performing a migration, or that the content of their {recently-updated} macro is not being updated correctly. Errors in the atlassian-confluence.log file corroborate such problems. Hence, to avoid these issues, it is strongly recommended that you rebuild your content indices after performing a migration.

Ensuring no contact with production systems

To ensure no contact with external systems, you will need to disable both inbound and outbound mail services.

1. Disable global outbound mail by running the following database query:

   ```
   SELECT * FROM BANDANA WHERE BANDANAKEY =
   'atlassian.confluence.smtp.mail.accounts';
   ```

2. Disable space-level mail archiving by running the following database query:

   ```
   SELECT * FROM BANDANA WHERE BANDANAKEY =
   'atlassian.confluence.space.mailaccounts';
   ```

Change 'SELECT *' to '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 site without any mails being sent or retrieved. Think carefully about other plugins which may access production systems (SQL macro, JIRA macro, etc.). If these write content, or create unwanted load on external systems, they should be disabled promptly after starting the test site.

Migrating from HTTPS to HTTP

You may want to migrate from a server secured by SSL to one which is not secured by SSL. For example, this may be useful if you are copying a Confluence site from a production to a test site.

To migrate from HTTPS to HTTP, undo the HTTPS-specific settings that are described on this page: Adding SSL for Secure Logins and Page Security.

Notes

- If you wish to **merge** two Confluence sites, you can consider using the remote import plugin. This plugin is currently not supported. The supported method would be to export a space and then import each space one by one. The two Confluence sites must be running the same version of Confluence.

From Confluence Evaluation through to Production Installation

So, you want to try Confluence on an evaluation installation, then move to a production installation when you are ready? This page gives an overview of the steps to follow.

Assumptions:

- This page starts with telling you how to install an evaluation Confluence site. If you have already finished evaluating Confluence, you can safely skip steps 1 to 3.
- Your production installation will be an installed version of Confluence, **not** a Confluence 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.

On this page:

- Step 1. Set up your evaluation Confluence site
- Step 2. Add users and content to your evaluation site
- Step 3. Look for interesting add-ons as part of your evaluation
- Step 4. Set up your production Confluence site

Related pages:

- Supported Platforms
- Add and Invite Users
- Getting Started as Confluence Administrator
- Confluence Installation and Upgrade Guide

Step 1. Set up your evaluation Confluence site

If you have already set up an evaluation Confluence site, you can skip this step.

Below is a summary of the installation and setup procedure, focusing on the choice of database.

To install Confluence:

1. Download the 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).

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.
- Choose **Create** to add pages and blog posts.

To add users: Choose the **cog** icon

then choose **User Management**.

For more tips about getting started, see Confluence Getting Started.

Step 3. Look for interesting add-ons as part of your evaluation

If you have finished evaluating Confluence, you can skip this step.

Add-ons, also called plugins, provide additional features that you can install into your Confluence site. Some of them are provided free of charge. Many of the commercial add-ons are available free for an evaluation period.

You can browse and download add-ons on the Atlassian Marketplace. You can also find add-ons via the Confluence user interface, which interacts with the Atlassian Marketplace for you.

To find useful add-ons via the Confluence user interface:

1. Choose the **cog** icon

   then choose **Add-ons**.
2. Choose **Find new add-ons**.

Step 4. Set up your production Confluence site

When you are ready to move from an evaluation site to a production site, you need to migrate to a production-ready database. This involves installing a new Confluence site with a new database, and instructing Confluence to copy the data from your evaluation site to the new site. You will also need to check some important configuration settings, and define your backup strategy. The instructions below lead you through all the steps required.

Migrating your data to a production database:

1. Choose a database carefully, with a focus on reliability and backups. See our list of supported databases. If you are unsure which one to choose, we recommend PostgreSQL.
2. Install a new database and a new Confluence site, by following our guide to migrating to another database. The guide will lead you through the following steps:
   - Setting up your database server.
   - Adding a Confluence database (schema) to your database server.
   - Installing a new, production-ready Confluence site.
   - Copying your Confluence data from your evaluation site to your new production site.

Setting important configuration options on your production site:

- Set the base URL. See Configuring the Server Base URL.
- Make sure you have configured an email server. See Configuring a Server for Outgoing Mail.
- Decide on proxy setup and other settings that determine where Confluence fits into your network. See Web Server Configuration.
- Consider setting up a secure connection via SSL. See Running Confluence Over SSL or HTTPS.
- Read our guidelines on security. See Best Practices for Configuring Confluence Security.
- Decide whether you will manage your users in Confluence or connect to an external LDAP directory. See Configuring User Directories.
- Decide whether you want to allow public (anonymous) access to your site. See Setting Up Public Access.
- Set up your permission scheme. See Permissions and Restrictions.

- Connect Confluence to JIRA aplications such as JIRA Software or JIRA Service Desk 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.

Migrate from Confluence Cloud to Server

This page is for people who are currently using Confluence Cloud, and wish to move to Confluence Server (a Confluence installation hosted on your own servers).

> For migrating from Confluence Server to Confluence Cloud, see Importing wiki data.

On this page:
- Migration steps
- Support, limitations, and recommendations
    - Upgrading to a standard Confluence Server release
    - Third-party add-ons
    - Database considerations
    - Confluence license

To migrate from Confluence Cloud to Server you'll need to download and install a special **Confluence Cloud release** and then move your data from your Confluence Cloud site to your newly installed site.

Confluence Cloud is typically ahead of Confluence Server, which means you can't move your data from Confluence Cloud to a site installed from the standard Confluence Server download because backups taken from Confluence Cloud are only compatible with the current Cloud release (for example, 'Confluence 5.9.1-OD-2015.36').

Note: It is important that you export your data and download Confluence on the same day to ensure that your data and your Confluence installation use the same version.

Migration steps

To migrate from Confluence Cloud to Confluence Server:

1. Export the data from your Confluence Cloud site, using the Confluence backup manager.
    - For detailed instructions, see Exporting wiki data.
    - You now have an XML export of your Confluence data.
2. Download the current Confluence Cloud release.
    Go to the Confluence Cloud download page and get the release for your operating system:
    - For **Windows 64-bit:** Get '5.x.x-OD-x - Windows Installer (64 bit)'
    - For **Windows 32-bit:** Get '5.x.x-OD-x - Windows Installer (32 bit)'
    - For **Linux 64-bit:** Get '5.x.x-OD-x - Linux Installer (64 bit)'
    - For **Linux 32-bit:** Get '5.x.x-OD-x - Linux Installer (32 bit)'
    - To install Confluence **from an archive** on **UNIX** or **Mac OS X:** Get '5.x-OD-x - Standalone (TAR.GZ Archive)'
    - To install Confluence **from an archive** on **Windows:** Get '5.x-OD-x - Standalone (ZIP Archive)'.
3. Follow the conditional Confluence Installation Guide to install Confluence.
4. Import the data from your backup file (XML export) into your new Confluence installation - see Restoring a Site.
5. Log in to your new Confluence site, using the following credentials:
    - Username: `sysadmin`
    - Password: `sysadmin`
6. Change the password immediately after logging in.

Support, limitations, and recommendations

Please note the following about your Confluence site installed from a Cloud release.

> - 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 as per this guide.

- The User Management section might be missing due to

  > **CONF-35177** - User and Group Links Missing from Admin Console After Migrating From OnDemand to Confluence 5.6.x `OPEN`

  Follow the workaround notes on the issue to enable the feature again.
- If your Confluence Cloud instance 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 as per

  > **APL-1144** - Allow relocation of application links even if the target application is still accessible. `OPEN`

  You can address that by either editing the XML prior to import, or by bulk editing those references in Confluence database as per this guide.

Upgrading to a standard Confluence Server release

> You must upgrade to a standard Confluence Server release (**major version**) as soon as one is available. For example, if you installed 5.9.0-OD-43, you must upgrade to 5.9.0 when it's released. Cloud releases *are* supported, but only until the next Confluence Server release is available.

Third-party add-ons

Because Confluence Cloud is typically ahead of Confluence Server, some third-party add-ons may not be compatible with the Cloud release. You may have to wait until you upgrade to the next full Confluence Server release before installing incompatible add-ons. Any add-ons that you were already using in Confluence Cloud will be compatible.

Database considerations

If you are uncertain about which database to choose for your new Confluence installation, we recommend PostgreSQL - see Database Setup for PostgreSQL. The Confluence Cloud site runs on PostgreSQL, so there should be no compatibility issues.

If you choose another supported database, Contact Support if you encounter any compatibility issues with other supported databases.

Confluence license

You will need a new license to migrate to Confluence Server. Your existing Cloud license cannot be used. You can get a new license at https://my.atlassian.com.

Confluence Release Notes

Confluence release notes provide information on the features and improvements in each release. This page includes release notes for major releases and minor (bugfix) releases. If you are upgrading from an earlier version of Confluence, you will find essential information in the **upgrade notes** associated with the relevant release notes.

Release notes are maintained in the Confluence Latest space. All links below will take you to the Confluence Latest where you can browse all release notes and upgrade notes.

On this page:

- Confluence 5 release notes
- Confluence 4 release notes
- Confluence 3 release notes
- Confluence 2 release notes
- Confluence 1 release notes

Confluence 5 release notes

Confluence 5.9

- Confluence 5.9.2 Release Notes
- **Confluence 5.9.1 Release Notes**
- (Confluence 5.9.0 was an internal release)

Confluence 5.8

- Confluence 5.8.18 Release Notes
- Confluence 5.8.17 Release Notes
- Confluence 5.8.16 Release Notes
- Confluence 5.8.15 Release Notes
- Confluence 5.8.14 Release Notes
- Confluence 5.8.13 Release Notes
- (Confluence 5.8.11 and 5.8.12 were internal releases)
- Confluence 5.8.10 Release Notes
- Confluence 5.8.9 Release Notes
- Confluence 5.8.8 Release Notes
- (Confluence 5.8.7 was an internal release)
- Confluence 5.8.6 Release Notes
- Confluence 5.8.5 Release Notes
- Confluence 5.8.4 Release Notes
- (Confluence 5.8.3 was an internal release)
- **Confluence 5.8.2 Release Notes**
- (Confluence 5.8.0 and 5.8.1 were internal releases)

Confluence 5.7

- Confluence 5.7.5 Release Notes
- Confluence 5.7.4 Release Notes
- Confluence 5.7.3 Release Notes
- (Confluence 5.7.2 was an internal release)
- Confluence 5.7.1 Release Notes
- **Confluence 5.7 Release Notes**

Confluence 5.6

- Confluence 5.6.6 Release Notes
- Confluence 5.6.5 Release Notes
- Confluence 5.6.4 Release Notes
- Confluence 5.6.3 Release Notes
- (Confluence 5.6.2 was an internal release)
- Confluence 5.6.1 Release Notes
- **Confluence 5.6 Release Notes**

Confluence 5.5

- Confluence 5.5.7 Release Notes
- Confluence 5.5.6 Release Notes
- (Confluence 5.5.5 was an internal release)
- (Confluence 5.5.4 was withdrawn)
- Confluence 5.5.3 Release Notes
- Confluence 5.5.2 Release Notes
- Confluence 5.5.1 Release Notes
- **Confluence 5.5 Release Notes**

Confluence 5.4

- Confluence 5.4.4 Release Notes
- Confluence 5.4.3 Release Notes
- Confluence 5.4.2 Release Notes
- Confluence 5.4.1 Release Notes
- **Confluence 5.4 Release Notes**

Confluence 5.3

- Confluence 5.3.4 Release Notes
- Confluence 5.3.1 Release Notes
- **Confluence 5.3 Release Notes**

Confluence 5.2

- Confluence 5.2.5 Release Notes
- Confluence 5.2.4 Release Notes
- **Confluence 5.2.3 Release Notes**
- (Confluence 5.2.0 - 5.2.2 were internal releases)

Confluence 5.1

- Confluence 5.1.5 Release Notes
- Confluence 5.1.4 Release Notes
- Confluence 5.1.3 Release Notes
- Confluence 5.1.2 Release Notes
- Confluence 5.1.1 Release Notes
- **Confluence 5.1 Release Notes**

Confluence 5.0

- Confluence 5.0.3 Release Notes
- Confluence 5.0.2 Release Notes
- Confluence 5.0.1 Release Notes
- **Confluence 5.0 Release Notes**

Confluence 4 release notes

Confluence 4.3

- Confluence 4.3.7 Release Notes
- Confluence 4.3.6 Release Notes
- Confluence 4.3.5 Release Notes
- (Confluence 4.3.4 was an internal release)
- Confluence 4.3.3 Release Notes

Confluence 4.1

- Confluence 4.1.9 Release Notes
- (Confluence 4.1.8 was an internal release)
- Confluence 4.1.7 Release Notes
- Confluence 4.1.6 Release Notes
- Confluence 4.1.5 Release Notes

- Confluence 4.3.2 Release Notes
- Confluence 4.3.1 Release Notes
- **Confluence 4.3 Release Notes**

Confluence 4.2

- Confluence 4.2.13 Release Notes
- Confluence 4.2.12 Release Notes
- Confluence 4.2.11 Release Notes
- (Confluence 4.2.9 and 4.2.10 were internal releases)
- Confluence 4.2.8 Release Notes
- Confluence 4.2.7 Release Notes
- Confluence 4.2.6 Release Notes
- Confluence 4.2.5 Release Notes
- Confluence 4.2.4 Release Notes
- Confluence 4.2.3 Release Notes
- Confluence 4.2.2 Release Notes
- Confluence 4.2.1 Release Notes
- **Confluence 4.2 Release Notes**

Confluence 3 release notes

Confluence 3.5

- Confluence 3.5.17 Release Notes
- Confluence 3.5.16 Release Notes
- Confluence 3.5.13 Release Notes
- (Confluence 3.5.12 was an internal release)
- Confluence 3.5.11 Release Notes
- (Confluence 3.5.10 was an internal release)
- Confluence 3.5.9 Release Notes
- (Confluence 3.5.8 was an internal release)
- Confluence 3.5.7 Release Notes
- Confluence 3.5.6 Release Notes
- Confluence 3.5.5 Release Notes
- Confluence 3.5.4 Release Notes
- Confluence 3.5.3 Release Notes
- Confluence 3.5.2 Release Notes
- Confluence 3.5.1 Release Notes
- **Confluence 3.5 Release Notes**

Confluence 3.4

- Confluence 3.4.9 Release Notes
- Confluence 3.4.8 Release Notes
- Confluence 3.4.7 Release Notes
- Confluence 3.4.6 Release Notes
- Confluence 3.4.5 Release Notes
- (Confluence 3.4.4 was an internal release)
- Confluence 3.4.3 Release Notes
- Confluence 3.4.2 Release Notes
- Confluence 3.4.1 Release Notes
- **Confluence 3.4 Release Notes**

Confluence 2 release notes

Confluence 2.10

- Confluence 2.10.4 Release Notes
- Confluence 2.10.3 Release Notes
- Confluence 2.10.2 Release Notes

- Confluence 4.1.4 Release Notes
- Confluence 4.1.3 Release Notes
- Confluence 4.1.2 Release Notes
- (Confluence 4.1.1 was an internal release)
- **Confluence 4.1 Release Notes**

Confluence 4.0

- **Confluence 4.0 Release Notes**

Confluence 3.3

- Confluence 3.3.3 Release Notes
- (Confluence 3.3.2 was an internal release)
- Confluence 3.3.1 Release Notes
- **Confluence 3.3 Release Notes**

Confluence 3.2

- Confluence 3.2.1 Release Notes
- **Confluence 3.2 Release Notes**

Confluence 3.1

- Confluence 3.1.2 Release Notes
- Confluence 3.1.1 Release Notes
- **Confluence 3.1 Release Notes**

Confluence 3.0

- Confluence 3.0.2 Release Notes
- Confluence 3.0.1 Release Notes
- **Confluence 3.0 Release Notes**

Confluence 2.4

- Release Notes 2.4.5
- Release Notes 2.4.4
- Release Notes 2.4.3

- Confluence 2.10.1 Release Notes
- **Confluence 2.10 Release Notes**

Confluence 2.9

- Confluence 2.9.3 Release Notes
- Confluence 2.9.2 Release Notes
- Confluence 2.9.1 Release Notes
- **Confluence 2.9 Release Notes**

Confluence 2.8

- Confluence 2.8.3 Release Notes
- Confluence 2.8.2 Release Notes
- Confluence 2.8.1 Release Notes
- Confluence 2.8 Release Notes
- **Confluence 2.8 Beta Release Notes**

Confluence 2.7

- Confluence 2.7.4 Release Notes
- Confluence 2.7.3 Release Notes
- Confluence 2.7.2 Release Notes
- Confluence 2.7.1 Release Notes
- **Confluence 2.7 Release Notes**

Confluence 2.6

- Confluence 2.6.3 Release Notes
- Confluence 2.6.2 Release Notes
- Confluence 2.6.1 Release Notes
- **Confluence 2.6 Release Notes**

Confluence 2.5

- Release Notes 2.5.8
- Release Notes 2.5.7
- Release Notes 2.5.6
- Release Notes 2.5.5
- Release Notes 2.5.4
- Release Notes 2.5.3
- Release Notes 2.5.2
- Release Notes 2.5.1
- **Release Notes 2.5**

Confluence 1 release notes

Confluence 1.4

- Release Notes 1.4.4
- Release Notes 1.4.3
- Release Notes 1.4.2
- Release Notes 1.4.1
- **Release Notes 1.4**

Confluence 1.3

- Release Notes 1.3.6
- Release Notes 1.3.5
- Release Notes 1.3.4

- Release Notes 2.4.2

Confluence 2.3

- Release Notes 2.3.3
- Release Notes 2.3.2
- Release Notes 2.3.1
- **Release Notes 2.3**

Confluence 2.2

- Release Notes 2.2.10
- Release Notes 2.2.9
- Release Notes 2.2.8
- Release Notes 2.2.7
- Release Notes 2.2.6a
- Release Notes 2.2.5
- Release Notes 2.2.4
- Release Notes 2.2.3
- Release Notes 2.2.2
- Release Notes 2.2.1
- **Release Notes 2.2**

Confluence 2.1

- Release Notes 2.1.5
- Release Notes 2.1.4
- Release Notes 2.1.3
- Release Notes 2.1.2
- Release Notes 2.1.1
- **Release Notes 2.1**

Confluence 2.0

- Release Notes 2.0.3
- Release Notes 2.0.2
- Release Notes 2.0.1
- **Release Notes 2.0**

Confluence 1.2

- Release Notes 1.2.3
- Release Notes 1.2.2
- Release Notes 1.2.1
- **Release Notes 1.2**

Confluence 1.1

- Release Notes 1.1.2
- Release Notes 1.1.1
- **Release Notes 1.1**

- [Release Notes 1.3.2](#)
- [Release Notes 1.3.1](#)
- **[Release Notes 1.3](#)**

Confluence Release Summary

This page shows the highlights of the major Confluence releases.

Current Release

For information about the latest release, please go to the [Confluence Release Notes](#).

Confluence 5.9 - November 2015

- It's the dashboard, but not as you know it - redesigned dashboard
- A better way to restrict pages - a new, easier-to-use page restrictions dialog
- One-click create for blank pages
- CQL for search - advanced searching with CQL (Confluence Query Language) filters
- Say goodbye to unknown attachment fails - copy and paste images with confidence
- Notify watchers remembers your preference for each page
- More control over page excerpts - simple or rich excerpts
- Improvements to on-boarding, inline comments, Data Center migration, and lots more

Confluence 5.8 - June 2015

- Close the collaboration loop - present your files in full screen
- Massive improvements for big tables - auto row numbering and copy and paste for columns
- Restore deleted files from the trash
- More granular options for Page Properties Report and Content by Label macros with CQL
- View roadmaps by week or month
- Plus improvements to Data Center indexing, inline comments, notify watchers, onboarding and much more.

Confluence 5.7 - January 2015

- Inline comments on pages and pinned comments on files to simplify your feedback loop
- A whole new files experience - upload, preview, share and more.
- HipChat Cloud integration to supercharge your team collaboration.
- New macro for planning visual roadmaps.
- Data Center improvements including cluster monitoring and the return of workbox.
- New profile picture and space logo uploader
- Content search REST API (CQL)
- Improvements to the editor.

Confluence 5.6 - September 2014

- Introducing Confluence Data Center
- Keep track of your JIRA issues visually
- Restrict blog posts
- Collaborate faster with quicker comments
- Recently viewed pages on the go
- Tasks sorting improvements
- Insert date lozenges with ease
- Sidebar improvements
- PDF Export improvements
- Support for Office 2013

Confluence 5.5 - 30 April 2014

- Tasks that fit the way you work
- Move blog posts
- Disable space blueprints
- Add a custom favicon
- Improved JIRA integration
- Extended Page Properties and Page Properties Report macros
- Confluence Health Check
- New REST API

Confluence 5.4 - 3 December 2013

- Confluence and JIRA: better together for agile teams
- Create JIRA issues from Confluence
- Give feedback fast with Quote in Comment
- Communicate your releases with new JIRA Reports
- Reflect and improve on your sprints with Retrospectives
- Complete JIRA traceability
- Manage your docs with the Documentation space blueprint
- Macro and blueprint improvements including a new JIRA chart macro, and improvements to the JIRA Issues macro, Page Properties Report macro and Product Requirements blueprint

Confluence 5.3 - 1 October 2013

- Space Blueprints including Knowledge Base space and Team space
- Shared Links blueprint
- Ability to change usernames
- Display page tree in the sidebar
- Request access to a restricted page
- Improved macros including Livesearch, Include Page, Labels List and Code Block

Confluence 5.2 - 13 August 2013

- Faster and cleaner search
- Introducing the Decisions blueprint
- Blueprints keep getting better
- Add multiple layout sections to your pages
- Distraction free editing
- A single, smarter way of adding JIRA issues
- Collaborate with your team on the go
- Speedier interactions

Confluence 5.1 – 27 March 2013

- Introducing blueprints
- Meeting Notes blueprint
- File List blueprint
- Product Requirements blueprint
- New template features
- HTML5 viewers for PDF and PowerPoint files
- Page Properties Report macro
- Improved macros

Confluence 5.0 – 26 February 2013

- A visual refresh
- Updated global navigation
- Content creation made simple

- New sidebar for content discovery
- Editor improvements
- Quick access to recently viewed pages
- Redesigned space administration and space tools
- Improved theming and branding
- Improved user and group management for large sites

Confluence 4.3 – 4 September 2012

- Workbox notifications
- Personal tasks
- Tasks on pages
- Confluence mobile
- Table sorting and highlighting
- Draggable images and macros
- Rich text templates
- Space archiving
- Improved user invitations and signup options
- Default space permissions

Confluence 4.2 – 10 April 2012

- Page layouts
- Likes
- Quick comments
- Popular content on the dashboard
- Recommended content by email
- Labels on attachments
- Signup invitations via URL
- Easy upgrade, try and buy for plugins

Confluence 4.1 – 13 December 2011

- Autoconvert for Pasted Links
- Image effects
- Quick find and replace
- Follow Your Network On the Dashboard
- Space attachments macro
- Global PDF stylesheets
- Use any character in page titles
- New translation feature

Confluence 4.0 – 19 September 2011

- Brand New Editor
- Simplified Editing Experience
- New Macros
- Faster Editing Experience
- Introducing @mentions
- Improved Page Comparison Functionality
- Email Notification Improvements
- New Confluence Installer and Guided Upgrades
- New Editor Plugin Points for Developers

Confluence 3.5 – 16 March 2011

- Easy, Powerful Connections to Active Directory, LDAP and Crowd
- Improved JIRA Integration
- Drag-and-Drop for HTML5 Browsers
- Autowatch and Improved Notification Settings
- Sharing Pages and Blog Posts
- Enhanced Code Macro

- More Administrative Improvements
- "What's New" Feature Tour
- Categories, a New Way of Organising Spaces
- Embedding Audio and Video with the Multimedia Macro
- Infrastructure Changes

Confluence 3.4 – 12 October 2010

- New Keyboard Shortcuts, Mac-Friendly Too
- Keyboard Shortcut Dialog
- User Macros in Macro Browser and Autocomplete
- New Plugin Manager
- Improved Performance
- Infrastructure Changes

Confluence 3.3 – 7 July 2010

- Confluence Page Gadget
- Autocomplete for Inserting Macros
- Property Panels for Links
- Property Panels for Images
- Manage Watchers
- Email Notifications for Network Activity and Blogs
- Blog Improvements
- Context-Sensitive Help Links
- Security Features
- Infrastructure Changes

Confluence 3.2 – 24 March 2010

- Autocomplete for Inserting Links
- Autocomplete for Embedding Images and Documents
- A Link Browser that's Smarter, Smoother, Faster
- New Documentation Theme
- New Easy Reader Theme
- Template Bundles
- Reordering while Moving a Page
- New Keyboard Shortcuts and Editor Hints
- User Interface Enhancements

Confluence 3.1 – 8 December 2009

- Introducing Gadgets
- Drag-and-Drop
- Office 2007 Support
- New 'Move Page' Feature
- Enhanced Image Browser
- Draft Comparisons
- Page Restrictions Dialog Box
- Other Editor Enhancements
- New Web Browser Versions Supported

Confluence 3.0 – 1 June 2009

- Introducing the Macro Browser
- Enhanced User Profiles
- Introducing Your Network
- New User Status
- New Hover Profile Feature
- Customisable Enhanced PDF Exports

- Improved Rich Text Editor
- Performance Improvements
- Engine Room and Developer Community
- Administration Improvements

Confluence 2.10 – 3 December 2008

- Introducing the Widget Connector
- Improved Office Connector Now Bundled
- Introducing Quick Navigation
- 'Did You Mean', OpenSearch and More
- Custom Stylesheets for Confluence Spaces
- Updated JIRA Issues Macro with Custom Fields and Dynamic Display
- Enhanced User and Group Management
- Upgraded Rich Text Editor
- Universal Wiki Converter now with SharePoint Import and More
- Improved Activity Macros
- Plugin Framework 2

Confluence 2.9 – 7 August 2008

- Streamlined Search
- Auto Save
- Charts
- Page Tree
- Gallery
- New Tutorial
- More in the Menus
- Alphabetical Page Ordering
- Better Spam Prevention
- Plugin Repository
- Engine Room and Developers' Community

Confluence 2.8 – 10 April 2008

- Dynamic menus and simplified screen design
- Page ordering
- Collapsible comments
- Multiple-label filter
- Confluence installer
- Task list
- Performance enhancements
- Administration, management and monitoring

Confluence 2.7 – 12 December 2007

- JIRA Issues and Portlet macros use new trusted authentication
- Two-tier administrator permissions
- Inserting images and attaching files during page creation
- Sorting of images in Gallery macro
- Simplified and improved logging
- Performance, maintainability and administration

Confluence 2.6 – 27 September 2007

- Fresh look for the Default theme
- Personalised comments and Dashboard
- Space description on Dashboard
- Labels on templates
- Default content for space home pages

- Social Bookmarking plugin now bundled with Confluence
- Back-dating and renaming news items

Confluence 2.5 – 29 April 2007

- Introducing flexible page restrictions
- Dynamic task list JRE incompatibilities
- contentbylabel macro supports AND condition

Confluence 2.4 – 14 March 2007

- Editable comments
- Page mailing

Confluence 2.3 – 5 January 2007

- Confluence Massive – cluster support
- People directory
- Activity plugin – usage statistics
- Blogging RPC plugin – manage news in Confluence using blogger-compatible desktop clients
- WebDAV client support via WebDAV plugin – create, edit, move pages, attachments, etc via WebDAV

Confluence 2.2 – 27 April 2006

- Personal spaces
- Localisation/internationalisation – drop-in language packs (similar to JIRA)
- CAPTCHA support – spam protection
- Improved searching
- Improved LDAP performance
- Confluence ships with Tomcat 5.5

Confluence 2.1 – 20 December 2005

- Autosave
- Concurrent edit warnings
- LDAP integration with Atlassian User/POLIS

Confluence 2.0 – 17 November 2005

- Rich Text Editing – WYSIWYG editor
- Labels
- Dashboard tabs – All, My, Team, New
- RSS builder
- Export pages as Word documents
- Copy pages

Confluence 1.4 – 23 May 2005

- New user interface
- Enhanced editing – doing more in the edit interface
- Page permissions
- New plugin types
- Configurable themes
- Completely rewritten Wiki to HTML conversion engine

Confluence 1.3 – 30 November 2004

- Mail archiving
- Themes

- Trash can
- More granular space permissions

Confluence 1.2 – 23 August 2004

- Page list views – alphabetical, directory view and search view of all pages in a space
- Image thumbnails and thumbnail galleries
- Threaded comments
- Enhanced Search - indexing attachment comments and file names and contextual searching
- New permissions interface

Upgrade Notes Overview

Typically, each major release of Confluence comes with upgrade notes, which are specific recommendations for upgrading from the previous major version. If you plan to upgrade and skip a few Confluence versions, you must read the upgrade notes for all major versions between your current version and the version to which you are upgrading, to make sure you do not miss something important.

Please read our general information about upgrading Confluence.

> **For example:**
> When upgrading from Confluence 5.3 to Confluence 5.5, read the upgrade notes for Confluence 5.4, as well as those for Confluence 5.5.

Also, we strongly recommend that you read the upgrade notes for any minor releases in-between, since they contain important information that will affect your Confluence upgrade.

Below is a list of upgrade notes for previous major releases of Confluence, as well as the upgrade notes for important minor releases:

- Confluence 5.9 Upgrade Notes
- Confluence 5.8 Upgrade Notes
- Confluence 5.7 Upgrade Notes
- Confluence 5.6 Upgrade Notes
- Confluence 5.5 Upgrade Notes
- Confluence 5.4 Upgrade Notes
- Confluence 5.3 Upgrade Notes
- Confluence 5.2 Upgrade Notes
- Confluence 5.1 Upgrade Notes
- Confluence 5.0 Upgrade Notes
- Confluence 4.3 Upgrade Notes
- Confluence 4.2 Upgrade Notes
- Confluence 4.1 Upgrade Notes
- Confluence 4.0 Upgrade Notes
- Confluence 3.5 Upgrade Notes
- Confluence 3.4 Upgrade Notes
- Confluence 3.3 Upgrade Notes
- Confluence 3.2 Upgrade Notes
- Confluence 3.1 Upgrade Notes
- Confluence 3.0.1 Upgrade Notes
- Confluence 3.0 Upgrade Notes
- Confluence 2.10 Upgrade Notes
- Confluence 2.9 Upgrade Notes
- Confluence 2.8 Upgrade Notes
- Confluence 2.7 Upgrade Notes
- Confluence 2.6 Upgrade Notes

You will find the upgrade notes attached to the release notes for the relevant version.

Useful plugins

Before installing an add-on (also called a plugin) into your Confluence site, please check the add-on's information page to see whether it is supported by Atlassian, by another vendor, or not at all. See our guidelines on add-on support.

Getting Help and Support

Need an answer to your question? Looking for a real person to solve a problem? We're here to help!

Online help

Check the documentation for information about using and administering Confluence.

Support tools in Confluence

Administrators can access a range of Support tools within Confluence. Go to

⚙ ▾

> **General Configuration** > **Support Tools** to check the health of your instance, scan your logs, or send a support zip to our Support team.

Online community

If you can't find what you need in the documentation, try asking in our question and answers forum, Atlassian Answers.

Features and bugs

We love to hear your requests for new features and improvements! You make a feature suggestion in the Confluence (CONF) project in our JIRA issue tracker.

Alas, bugs do happen. If you find one, we'll do our best to fix it. Please raise a bug report in our JIRA issue tracker.

Features and bugs

We love to hear your requests for new features and improvements! You can add a request in the Confluence (CONF) project in our JIRA issue tracker.

Alas, bugs do happen. If you find one, we will do our best to fix it. Please raise a bug report in the issue tracker.

Support team

For one-on-one help, create a support request at Atlassian Support. A support engineer will follow up with you quickly.

If you don't have an account yet, you can create one - just follow the prompts. When creating your support request, please provide as much detail as possible. This will help us resolve your issue faster. See Troubleshooting Problems and Requesting Technical Support.

Confluence Resources

Resources for evaluators

- Free trial
- Feature tour

Resources for administrators

- Confluence knowledge base
- Guide to installing an Atlassian integrated suite
- The big list of Atlassian gadgets

Resources for developers

- Atlassian Developers site
- Developer topics on Atlassian Answers

Downloadable documentation

- Confluence documentation in PDF format

Add-ons and plugins

- Documentation for the Confluence SharePoint Connector
- Atlassian Marketplace

Support

- Atlassian Support
- Support policies

Training

- Atlassian training

Answers

- Confluence at Atlassian Answers

Mailing lists

- Visit http://my.atlassian.com to sign up for mailing lists relating to Atlassian products, such as technical alerts, product announcements and developer updates.

Feature requests and bug reports

- Issue tracker for Confluence

Confluence SharePoint Connector

> **Latest version of the Confluence SharePoint Connector**
> Confluence SharePoint Connector 1.9.5 has now been released. See the SharePoint Connector 1.9.5 Release Notes.

With the Confluence SharePoint Connector you can combine Confluence's free-form, easy to edit wiki with the document management and workflow strengths of SharePoint.

- Display SharePoint document libraries, calendars, links, discussions and more on your Confluence wiki pages. Edit SharePoint's Office documents directly from Confluence and save them back to SharePoint.
- Embed Confluence pages and Confluence page trees into a SharePoint page. Click through from SharePoint to Confluence.
- Enjoy automatic login (single sign-on) between Confluence and SharePoint.
- Search Confluence and SharePoint content together, retrieving a unified set of results

Please refer to the SharePoint Connector documentation for more information.

Support Policies

Welcome to the support policies index page. Here, you'll find information about how Atlassian Support can help you and how to get in touch with our helpful support engineers. Please choose the relevant page below to find out more.

- Bug Fixing Policy
- New Features Policy
- Security Bugfix Policy

To request support from Atlassian, please raise a support issue in our online support system. To do this, visit support.atlassian.com, log in (creating an account if need be) and create an issue under Confluence. Our friendly support engineers will get right back to you with an answer.

Bug Fixing Policy

Summary

- Atlassian Support will help with workarounds and bug reporting.
- Critical bugs will generally be fixed in the next maintenance release.
- Non critical bugs will be scheduled according to a variety of considerations.

Raising a Bug Report

Atlassian Support is eager and happy to help verify bugs — we take pride in it! Please open a support request in our support system providing as much information as possible about how to replicate the problem you are experiencing. We will replicate the bug to verify, then lodge the report for you. We'll also try to construct workarounds if they're possible.

Customers and plugin developers are also welcome to open bug reports on our issue tracking systems directly. Use the appropriate project on http://jira.atlassian.com to report bugs for Atlassian products.

When raising a new bug, you should rate the priority of a bug according to our JIRA usage guidelines. Customers should watch a filed bug in order to receive e-mail notification when a "Fix Version" is scheduled for release.

How Atlassian Approaches Bug Fixing

Maintenance (bug fix) releases come out more frequently than major releases and attempt to target the most critical bugs affecting our customers. The notation for a maintenance release is the final number in the version (ie the 1 in 3.0.1).

If a bug is critical (production application down or major malfunction causing business revenue loss or high numbers of staff unable to perform their normal functions) then it will be fixed in the next maintenance release provided that:

- The fix is technically feasible (i.e. it doesn't require a major architectural change).
- It does not impact the quality or integrity of a product.

For non-critical bugs, the developer assigned to fixing bugs prioritises the non-critical bug according to these factors:

- How many of our supported configurations are affected by the problem.
- Whether there is an effective workaround or patch.
- How difficult the issue is to fix.
- Whether many bugs in one area can be fixed at one time.

The developers responsible for bug fixing also monitor comments on existing bugs and new bugs submitted in JIRA, so you can provide feedback in this way. We give high priority consideration to security issues.

When considering the priority of a non-critical bug we try to determine a 'value' score for a bug which takes into account the severity of the bug from the customer's perspective, how prevalent the bug is and whether roadmap features may render the bug obsolete. We combine this with a complexity score (i.e. how difficult the bug is). These two dimensions are used when developers self serve from the bug pile.

Further reading

See Atlassian Support Offerings for more support-related information.

New Features Policy

Summary

- We encourage and display customer comments and votes openly in our issue tracking system, http://jira.atlassian.com.
- We do not publish roadmaps.
- Product Managers review our most popular voted issues on a regular basis.
- We schedule features based on a variety of factors.
- Our Atlassian Bug Fixing Policy is distinct from this process.
- Atlassian provides consistent updates on the top 20 issues.

How to Track what Features are Being Implemented

When a new feature or improvement is scheduled, the 'fix-for' version will be indicated in the JIRA issue. This happens for the upcoming release only. We maintain roadmaps for more distant releases internally, but because

these roadmaps are often pre-empted by changing customer demands, we do not publish them.

How Atlassian Chooses What to Implement

In every major release we *aim* to implement highly requested features, but it is not the only determining factor. Other factors include:

- **Customer contact**: We get the chance to meet customers and hear their successes and challenges at Atlassian Summit, Atlassian Unite, developer conferences, and road shows.
- **Customer interviews**: All product managers at Atlassian do customer interviews. Our interviews are not simply to capture a list of features, but to understand our customers' goals and plans.
- **Community forums**: There are large volumes of posts on answers, of votes and comments on jira.atlassian.com, and of conversations on community forums like groups on LinkedIn.
- **Customer Support**: Our support team provides clear insights into the issues that are challenging for customers, and which are generating the most calls to support
- **Atlassian Experts**: Our Experts provide insights into real-world customer deployments, especially for customers at scale.
- **Evaluator Feedback**: When someone new tries our products, we want to know what they liked and disliked and often reach out to them for more detail.
- **In product feedback**: The JIRA Issue Collectors that we embed our products for evaluators and our Early Access Program give us a constant pulse on how users are experiencing our product.
- **Usage data:** Are customers using the features we have developed?
- **Product strategy:** Our long-term strategic vision for the product.
- Please read our post on Atlassian Answers for a more detailed explanation.

How to Contribute to Feature Development

Influencing Atlassian's release cycle
We encourage our customers to vote on issues that have been raised in our public JIRA instance, http://jira.atlassian.com. Please find out if your request already exists - if it does, vote for it. If you do not find it you may wish to create a new one.

Extending Atlassian Products
Atlassian products have powerful and flexible extension APIs. If you would like to see a particular feature implemented, it may be possible to develop the feature as a plugin. Documentation regarding the plugin APIs is available. Advice on extending either product may be available on the user mailing-lists, or at Atlassian Answers.

If you require significant customisations, you may wish to get in touch with our partners. They specialise in extending Atlassian products and can do this work for you. If you are interested, please contact us.

Further reading

See Atlassian Support Offerings for more support-related information.

Security Bugfix Policy

See Security @ Atlassian for more information on our security bugfix policy.

Upgrading from Confluence 3.5 to 5.1 - Simplified Guide

> In accordance with Atlassian's standard support policy, Confluence 3.5 will reach End of Life on **6 September 2013**.

If your organisation is still using Confluence 3.5, we strongly recommend you to upgrade to Confluence 5.1.4.

This is a **simplified guide** to upgrade from Confluence 3.5 to 5.1.4. A **more detailed guide** is available here.

What's new in Confluence 5.1?

The following blog posts outline the main features and changes in each version:

- **4.0:** Faster, richer, more reliable editor
- **4.3:** Tasks, mobile support, and in-app notifications
- **5.0:** All-new design
- **5.1:** Powerful page templates for meeting notes, requirements, and decisions

Wiki Markup

One of the the major differences between Confluence 3.5 and newer versions is the migration of wiki markup to our new editor. We understand that you may still require wiki markup from time to time and we feel that the following resources will make things easier:

- Confluence 4 Editor - What's Changed for Users of the Old Rich Text Editor
  - The editor auto converts wiki markup to the new formatting.
  - The editor has the ability to **Insert** > **Wiki Markup**.
- Bob Swift's Wiki Markup Addon allows you to preserve wiki markup if necessary. This plugin is currently offered free of charge.
- **API Changes** - Confluence now stores information in an XHTML format, which means that scripted behavior against the Confluence API will break. We have provided the convertWikiToStorageFormat method to ease the transition.

Preparing to upgrade

1. Check if your currently installed Plugins are compatible with the chosen target Confluence version

Verify whether a compatible version of the plugin is available in Confluence 5.1.4. We've made this easy for you: This information is available via the Confluence Upgrade Check in the plugin administration section of Confluence. See the documentation: Checking add-on compatibility with application updates. You can also check the respective home pages for these plugins on the Atlassian Marketplace.

2. Test the upgrade in a staging environment before upgrading your production instance

This will allow you to anticipate any unexpected behavior during the upgrade. Testing will also allow you to plan the upgrade of the production instance, as it will give you a better understanding of how long the upgrade will take and how much time you will need to re-install Plugins and re-apply any customisation.

3. Check the Upgrade Notes for the 3.5 to 5.1 Upgrade

Make sure that you are aware of the changes listed in the Simplified Upgrade Notes - Upgrading from 3.5 to 5.1 page. These **Upgrade Notes** are a summary of important changes in the platforms supported by Confluence and information to be considered during the upgrade.

4. Prepare your team

Confluence 5 introduced significant changes to the user experience. To help you manage the change in your organization, we have provided the Planning for Confluence 5 guide to help prepare your team.

Upgrading Confluence

This is a **simplified version** of the Confluence Upgrade process. A more detailed guide is available here.

The following instructions will help you to run the upgrade in **two steps**, on Windows or Linux environments:

1. Upgrade from **3.5** to **5.0.3**.
2. Upgrade from **5.0.3** to **5.1.4**.

This means that you will run the following steps twice. First, to upgrade from **3.5** to **5.0.3** and then a second time to upgrade from **5.0.3** to **5.1.4**.

Upgrading Confluence on Windows

1. Download the **Confluence Windows Installer** (.exe) file (for the new version of Confluence) from the Confluence Download Center.
2. Run the .exe file to start the upgrade wizard.
   If a Windows 7 (or Vista) 'User Account Control' dialog box requests if you want to allow the upgrade wizard to make changes to your computer, specify '**Yes**'. If you do not, the installation wizard will have restricted access to your operating system and any subsequent installation options will be limited.
3. At the 'Upgrading Confluence?' step, choose the '**Upgrade an existing Confluence installation**' option.
4. In the 'Existing Confluence installation directory' field, specify the Confluence Installation Directory of your Confluence installation to be upgraded.
   The upgrade wizard will attempt to find an existing Confluence installation and use its location to pre-populate this field. However, always verify this location, particularly if you have multiple Confluence installations running on the same machine.
5. During subsequent steps of the upgrade wizard, you will be prompted to specify or do the following options:
   a. At the 'Back up Confluence directories' step, ensure the '**Back up Confluence home**' option is selected. This creates 'zip' archive file backups of your existing Confluence Installation and Confluence Home Directories in their respective parent directory locations.
      Please Note: Choosing this option is strongly recommended!
   b. At this point, the upgrade wizard notes any customisations in your existing installation directory which it cannot automatically migrate to your upgraded Confluence installation. If you are notified by the installer about any files containing such customisations, please make a note of the locations of these files as you will need to manually migrate their customisations (which are not mentioned in the overview above) to your upgraded Confluence installation. One relatively common customisation that the upgrade wizard cannot automatically migrate is an SSL configuration defined in the `conf/server.xml` file of the Confluence Installation Directory.
      Please Note: when upgrading from the version that was not installed by the installer the customisations can only be detected in the `confluence` subdirectory of your existing Confluence Installation Directory. Modifications to files in directories other than `conflu ence` will not be detected when you upgrade, for example, modifications to `conf/server.xml`. However the next time you upgrade (e.g. to version 4.1.1) the upgrade feature will cover modifications across the whole Confluence Installation Directory.
   c. At the 'Upgrade Check List' step, perform the following steps if you have not already done so:
      - Back up your external database.
      - If your Confluence site includes third-party plugins (add-ons), check that they will be compatible with your upgraded Confluence version. See the documentation: Checking

add-on compatibility with application updates.

    d. Choose '**Next**', your existing Confluence installation will be shut down if it is still running. The upgrade wizard will then:

        i. Back up your existing Confluence installation.
        ii. Delete the contents of the existing Confluence Installation Directory.
        iii. Install the new version of Confluence to the existing Confluence Installation Directory.
        iv. Starts your new (upgraded) Confluence installation.
        ⚠️ If you noted any files that contain customisations which must be migrated manually to your upgraded Confluence installation (above), then:

            1. Stop the upgraded Confluence installation.
            2. Migrate the customisations from these files into the upgraded installation directory.
            3. Restart the upgraded Confluence installation.

6. At the last step of the upgrade wizard, select the option to launch the upgraded Confluence installation in a browser so you can check the upgrade.

Congratulations, you have completed upgrading your Confluence installation on Windows!

> Finished running the steps to upgrade from **3.5** to **5.0.3**? Remember to run them again, upgrading from **5.0.3** to **5.1.4**.

Upgrading Confluence on Linux

1. Download the appropriate **Confluence 'Linux 64-bit / 32-bit Installer'** (.bin) file that suits your operating system (for the new version of Confluence) from the Confluence Download Center.
2. Open a Linux console and change directory (`cd`) to the '.bin' file's directory.
   ⚠️ If the '.bin' file is not executable after downloading it, make it executable, for example:
   `chmod a+x atlassian-confluence-X.Y.bin`
   (where X.Y represents your version of Confluence)
3. Execute the '.bin' file to start the upgrade wizard.
4. When prompted to choose between creating a new Confluence installation or upgrading an existing installation, choose the '**Upgrade an existing Confluence installation**' option.
5. Specify the Confluence Installation Directory of your Confluence installation to be upgraded.
   The upgrade wizard will attempt to find an existing Confluence installation and will provide its location as a choice. However, always verify this location, particularly if you have multiple Confluence installations running on the same machine.
6. During subsequent steps of the upgrade wizard, you will be prompted to specify or do the following options:
   a. Choose the option to back up Confluence's directories. This creates 'zip' archive file backups of your existing Confluence Installation and Confluence Home directories in their respective parent directory locations.
      Please Note: Choosing this option is strongly recommended!
   b. At this point, the upgrade wizard notes any customisations in your existing Confluence Installation Directory which it cannot automatically migrate to your upgraded Confluence installation. If you are notified of any files containing such customisations, please make a note of the locations of these files as you will need to manually migrate their customisations (which are not mentioned in the overview above) to your upgraded Confluence installation. One relatively common customisation that the upgrade wizard cannot automatically migrate is an SSL configuration defined in the `conf/server.xml` file of the Confluence Installation Directory.
      Please Note: when upgrading from the version that was not installed by the installer the customisations can only be detected in the `confluence` subdirectory of your existing Confluence Installation Directory. Modifications to files in directories other than `confluence` will not be detected when you upgrade, for example, modifications to `conf/server.xml`. However the next time you upgrade (e.g. to version 4.1.1) the upgrade feature will cover modifications across the whole Confluence Installation Directory.
   c. At the 'Upgrade Check List' step, perform the following steps if you have not already done so:
      - Back up your external database.
      - If your Confluence site includes third-party plugins (add-ons), check that they will be compatible with your upgraded Confluence version. See the documentation: Checking add-on compatibility with application updates.

d. Upon proceeding, your existing Confluence installation will be shut down if it is still running. The upgrade wizard will then:
  i. Back up your existing Confluence installation.
  ii. Delete the contents of the existing Confluence installation directory.
  iii. Install the new version of Confluence to the existing Confluence installation directory.
  iv. Starts your new (upgraded) Confluence installation.
    ⚠ If you noted any files that contain customisations which must be migrated manually to your upgraded Confluence installation (above), then:
      1. Stop the upgraded Confluence installation.
      2. Migrate the customisations from these files into the upgraded Confluence Installation Directory.
      3. Restart the upgraded Confluence installation.
7. The last step of the upgrade wizard provides you with a link to launch the upgraded Confluence installation in a browser, so you can check the upgrade.

Congratulations, you have completed upgrading your Confluence installation on Linux!

> Finished running the steps to upgrade from **3.5** to **5.0.3**? Remember to run them again, upgrading from **5.0.3** to **5.1.4**.

Possible issues during the upgrade

Known issues upgrading to Confluence 5

CONF-26596 - Using color macro inside headings in Wiki Markup results in double headings
`RESOLVED`

CONF-23287 - Confluence 4.0 upgrade does not take into account modifications in server.xml
`RESOLVED`

CONF-28495 - ConvertToRelativeLinksUpgradeTasks leaves the version comment "Corrected links that should have been relative instead of absolute." `RESOLVED`

CONF-28818 - Upgrading from 3.5.x with accentedpage title, anchors break content
`OPEN`

Getting help

- **Check for known issues**. Please check the known issues for the relevant release on this page of the Knowledge Base and follow the instructions to solve the problem.
- **Check for answers from the community.** Other users may have encountered the same issue. You can check for answers from the community at Atlassian Answers.
- **Did you encounter a problem during the Confluence upgrade?** Please refer to the guide to trouble shooting upgrades in the Confluence Knowledge Base.
- **If you encounter a problem and cannot solve it**, please create a support ticket and one of our support engineers will help you.

Upgrading from 3.5 to 5.1 - Upgrade Notes

This page is a summary of the Upgrade Notes released for 4.0 and newer versions. If you wish, you can still see the complete notes for each version below:

- Confluence 4.0 Upgrade Notes
- Confluence 4.1 Upgrade Notes
- Confluence 4.2 Upgrade Notes
- Confluence 4.3 Upgrade Notes
- Confluence 5.0 Upgrade Notes
- Confluence 5.1 Upgrade Notes

Preparing your team for Confluence 5

Confluence 5.0 introduces a number of **significant changes to the user experience**: A new way of creating content, a redesigned header, a new sidebar, and other changes to the look and feel of your site. People in your organisation will need to be aware of the coming changes, so that they can plan and prepare for them. We have written a guide to help you: Planning for Confluence 5.

On this page:

- Preparing your team for Confluence 5
- Integration and Supported Platforms
  - Team Calendars 3.1 or later required
  - End of support for Java 6
  - End of support for Tomcat 5.5
  - Complete List of Supported Platforms
- Upgrade Process and Effects
  - Content updates after upgrade
  - User macros migration
  - Email notifications
  - Plain text emails no longer available
- Checking for known issues and troubleshooting the Confluence upgrade

In addition, if you are **upgrading from Confluence 3.5 or earlier** please note that the change to the Confluence editing experience is significant. See the guide to Planning for Confluence 4.

Integration and Supported Platforms

Team Calendars 3.1 or later required

If you are using the Team Calendars add-on for Confluence, you will need to upgrade to Team Calendars 3.1 or later. Earlier versions of this add-on will not work with the new design in Confluence 5.0.

End of support for Java 6

As previously announced, from this release onwards we no longer offer support for Java 6 (JRE and JDK 1.6). Please see End of Support Announcements for Confluence.

End of support for Tomcat 5.5

As previously announced, from this release onwards we no longer offer support for Tomcat 5.5.x. Please see End of Support Announcements for Confluence.

Complete List of Supported Platforms

See the complete list of supported platforms for Confluence 5 at Supported Platforms.

Upgrade Process and Effects

Content updates after upgrade

You will notice once you upgrade to Confluence 5.0, every page will have been edited during the migration process. This manifests itself in the appearance of changes to every page's content by various different users. When migrating pages, Confluence will show the same username that last contributed to your content as the author of the change. This is the upgrade task of migrating your wiki page storage format to XHTML from Wiki Markup. Do not be alarmed by this. This will only happen once for all current versions of content. This will not trigger email notifications, but will update your RSS feeds. Where possible, a change comment of "Upgraded to Confluence 4" will be added to pages and blog posts.

User macros migration

After upgrading, user macros will not be able to be inserted until they have defined macro parameters. See U ser Macro Template Syntax. Defining macro parameters will mean that they can once again insert these user macros.

Email notifications

With the update of the email templates, the textual content of the notifications has also changed. These wording changes may invalidate any email filters set up to label or organise your email notifications.

Plain text emails no longer available

Emails are now only available in HTML format, which allows for the display of images and formatted content, such as changes made and the full content of the updated page or blog post.

Checking for known issues and troubleshooting the Confluence upgrade

After you have completed the steps required to upgrade your Confluence installation, check all the items on the **Confluence post-upgrade checklist** to ensure that everything works as expected. If something is not working correctly, please check for known Confluence issues and try troubleshooting your upgrade as described below:

- **Check for known issues**. Sometimes we find out about a problem with the latest version of Confluence after we have released the software. In such cases we publish information about the known issues in the Confluence Knowledge Base.
- **Check for answers from the community.** Other users may have encountered the same issue. You can check for answers from the community at Atlassian Answers.
- **Did you encounter a problem during the Confluence upgrade?** Please refer to the guide to trouble shooting upgrades in the Confluence Knowledge Base.

Troubleshooting Problems and Requesting Technical Support

This document tells you how to troubleshoot problems in Confluence and how to obtain technical support from Atlassian.

Troubleshooting a problem

If you have a problem with Confluence, follow the steps below.

To diagnose a problem and search for a solution:

1. If you are not a Confluence administrator, report your problem to the person in charge of your Confluence site and ask them to follow up on the issue.
2. Use the health check and log scanner (see be low) to check your Confluence logs for errors and attempt to match them against known issues in our knowledge base and bug tracker.
3. Check our knowledge base for solutions to known issues.
4. Check our issue tracker for known bugs.
5. If your problem may be related to an add-on, you can run Confluence in Safe Mode or temporarily disable any third party add-ons.
6. If you are having problems configuring a feature, please take a look at the appropriate guides:
   - Confluence Installation Guide
   - Confluence Setup Guide

On this page:

- Troubleshooting a problem
- Using the Health Check
- Scanning your Confluence logs to match known issues
- Raise a support request with an add-on vendor
- Raise a support request
  - Method 1: Create a support request from within Confluence
  - Method 2: Create a support request at our support site
- Creating a support zip file
- Logging a bug

Related pages:

- Requesting Support for External User Management
- Requesting Performance Support
- Confluence Knowledge Base

- Configuring Confluence
- Database Configuration

7. Check the following guides for troubleshooting specific problems:
    - Issues related to your database server: Database Troubleshooting.
    - Issues related to user management: R equesting Support for External User Management.
    - Issues related to performance: Reques ting Performance Support.
8. If the above tools and documentation don't solve your problem, you can create a support request and attach your support zip file. If you think you may have found a bug, you can create a bug report instead.

Using the Health Check

The Confluence Health Check provides a simple way to check the set up of your Confluence instance. The health check looks at things like your licence validity, basic database setup and more.

To access the health check go to

⚙ ▾

> **General Configuration** > **Support Tools**.

The health check will let you know if there are any problems.

Scanning your Confluence logs to match known issues

Confluence has an inbuilt log scanner that will check your Confluence logs for errors and try to match them against known issues in our knowledge base and bug tracker. You'll need Confluence Administrator or System Administrator permissions to do this.

To use the Confluence log scanner:

1. Go to

    ⚙ ▾

    > **General Configuration** > **Support Tools** > **Log Scanner**.
2. Choose **Scan** to scan the Confluence log file in its default location (change the location if necessary).
3. The log scanner will return a list of links, pointing to articles in our knowledge base and/or bug reports in our issue tracker.
    - The latest-reported problems are displayed first. By default only the most recent 10 matches are displayed. If you have more than 10 matches and want to display all results, click the 'Show All' link that appears on the top of the results page.
    - Choose a link to read the reported problem and possible solutions or workarounds.

Raise a support request with an add-on vendor

If you have an add-on related issue, please check whether the add-on is supported by Atlassian or a third party developer.

- Check support information for the add-on on the Atlassian Marketplace
- If the plugin is not supported by Atlassian, you'll need to contact the plugin developer directly.

You can read more about Atlassian support for add-ons.

Raise a support request

There are two ways to raise a support request with Atlassian - from within Confluence, or via our support

site.

If your problem concerns user management or performance, please take a look at the additional requirements in Requesting Support for External User Management or Requesting Performance Support bef ore creating a request. .

Method 1: Create a support request from within Confluence

This is the most convenient method, however if your zip file is too large, or there are security restrictions on your mail server, your request may not be correctly sent. You can also use this method to append system information to an existing support request.

You'll need Confluence Administrator or System Administrator permissions to do this. Your site will also need a mail server set up that allows you to send zip files.

To create a support request:

1. Go to
    ⚙ ▾
    > **General Configuration** > **Support Tools** > **Create Support Request**.
2. Provide as much information as possible, following these guidelines:
    *   **Contact Email** - this defaults to the email address of the logged-in user. It will be used to find your support account or, if no account is found, to create a new account. Confluence will send all further notifications and updates to this address.
    *   **Summary** - enter a short description of the problem.
    *   **Description** - enter as much information as possible, including any error messages that are appearing or any steps our support team can take to reproduce the problem.
    *   **Support Data to Attach** - choose the log and configuration information you would like to include. These will be attached to your support request as a zip file, and will help our support engineers resolve your problem. Refer to Working with Confluence Logs for information about the log files.
        ⓘ Log files can be very large (up to 20mb with the default log4j configuration, or more if you've customised the log settings) and your mail server may bounce the message if it is too large.
3. Choose **Send**.
    We recommend checking whether the email message has been successfully sent, and contacting your email administrator if you need special provisions for this email message.

Confluence will submit your request via email to Atlassian support. If you don't already have a support account, we'll automatically create one for you. You'll receive email updates about the request progress. You can also view the status of your support request and add any additional information at any time at Atlassian Support.

Screenshot: The support request form.

Support Request

To: * **Atlassian Support Team**

Contact Email: * rach@email.com.au

Summary: *

Priority: * Moderate loss of application functionality or performance resulting in multip ▼

Description: *

Support Data to Attach

Attaching logs and configuration information to your Support Request helps our engineers give you a faster and more complete response. Usernames and passwords are removed from your configuration files automatically (usernames will still appear in log files).

☑ Application Properties

Memory statistics, configuration options, and other key information. (highly recommended)

☑ Thread Dump

Generate and include a Thread Dump. (recommended)

☑ Confluence confluence.cfg.xml File

Your confluence.cfg.xml file, which contains database settings and other useful information. (recommended)

☑ Tomcat Configuration Files

Tomcat configuration files such as server.xml and web.xml.

☑ Authentication Configuration Files

Authentication configuration files such as atlassian-user.xml, osuser.xml and crowd.properties. Atlassian will handle your data with strict confidence and will not disclose it to any third parties. NOTE: Usernames, etc. referenced in your log files are NOT anonymised. (recommended)

☑ Cache Configuration Files

Cache configuration files, which are useful in diagnosing performance problems.

☑ Confluence Application Logs

Your Confluence application logs. (highly recommended)

☑ Tomcat Logs

Tomcat log files, e.g. catalina.out. (recommended)

☑ Limit File Sizes?

Limit the size of files included in your support zip to no more than 25Mb.

Send

Method 2: Create a support request at our support site

If your Confluence instance is not running, or does not have a mail server configured, you use the Atlassian Support site to create a request.

To create a support request:

1. If you don't have a free Atlassian support account, create one here.
2. Log in to https://support.atlassian.com.
3. Follow the prompts to choose your product and type of request.
4. Provide us as much information as possible about your problem, and your Confluence environment. (If Confluence is running, you can go to **System Information** in the administration console and copy that information into your request). You may also want to attach:

- A support zip. See Creating a support zip to find out how to generate it.
- Your log files (if your instance won't start up). See Working with Confluence Logs to find out how to access your log files.

You'll receive email updates about the request progress, and see the status of your request at any time in the Support site.

Creating a support zip file

We recommend that you attach a support zip file to every interaction with our support team. You'll need Confluence Administrator or System Administrator permissions to create a support zip.

To create a support zip:

1. Go to

   ⚙️ ▾

   > **General Configuration** > **Support Tools** > **Create Support Zip**.
2. Keep all checkboxes selected then choose **Create**.
3. Confluence will create the support zip file and display its location on the screen. Log in to the Confluence server to retrieve the file.

You can then attach the zip file to your support request (a new request or an existing one).

Screenshot: Creating a support zip file

Support Zip

This tool will allow you to create a zip file containing useful information about your instance. No information will be sent to Atlassian.

☑ Application Properties

    Memory statistics, configuration options, and other key information. (highly recommended)

☑ Thread Dump

    Generate and include a Thread Dump. (recommended)

☑ Confluence confluence.cfg.xml File

    Your confluence.cfg.xml file, which contains database settings and other useful information. (recommended)

☑ Tomcat Configuration Files

    Tomcat configuration files such as server.xml and web.xml.

☑ Authentication Configuration Files

    Authentication configuration files such as atlassian-user.xml, osuser.xml and crowd.properties. Atlassian will handle your data with strict

    confidence and will not disclose it to any third parties. NOTE: Usernames, etc. referenced in your log files are NOT anonymised. (recommended)

☑ Cache Configuration Files

    Cache configuration files, which are useful in diagnosing performance problems.

☑ Confluence Application Logs

    Your Confluence application logs. (highly recommended)

☑ Tomcat Logs

    Tomcat log files, e.g. catalina.out. (recommended)

☑ Limit File Sizes?

    Limit the size of files included in your support zip to no more than 25Mb.

[Create]

Logging a bug

If you have found a bug, the easiest way to report it is to create numbered instructions on how to reproduce the bug, create a support request, as described above, and then the Atlassian support team will confirm your bug and lodge a bug report.

Alternatively, you can check to see if anyone else has reported the bug, then log a bug report yourself.

To report a bug:

1. Go to the Confluence issue tracker.
2. Type keywords for your problem into the **Contains text** search box then click the Search icon

Confluence ▼	Type: All ▼	Status: All ▼	Assignee: All ▼	Contains text	More ▼	🔍

3. Read through the summaries of the bugs shown. If any summary describes your problem, you can watch the issue to be notified when something changes, or add more information as a comment.
4. If there's no existing bugs that match your problem, you can create a new bug report - including any information on how to duplicate the problem.
5. Sometimes it may be useful to include the result of the 500 error page, which you will find at this location:

```
https://<domain><host>:<port>/500page.jsp
```

Content Anonymiser for Data Backups

Atlassian may request a copy of the `entities.xml` file from a customer's exported zip file (backup file), in order to diagnose database corruption or to find a bug in Confluence.

If your data is confidential, you can run an anonymiser program over your `entities.xml` file to remove all your data and leaving only the structure of the export.

Usage

To run the Content Anonymiser on your backup file:

1. Download the anonymiser JAR (attached to this page).
2. Extract the `entities.xml` file from your zipped backup file to the same directory as the JAR.
3. Use the command prompt to go to the directory where all three files are located.
4. To create `cleaned.xml`, run the command

```
java -jar confluence-export-cleaner-1.1-jar-with-dependencies.jar entities.xml
cleaned.xml
```

5. Re-ZIP `cleaned.xml` with its `exportDescriptor.properties` to ensure we (eg. Atlassian Support) knows exactly which version of Confluence does the XML backup exported from.

How it works

The Content Anonymiser program replaces all the text content in file `entities.xml` with 'x' characters. For example, the word "Atlassian" will be transformed to "xxxxxxxxx". The resulting `cleaned.xml` file is expected to have the same size as the original file.

This release of the Content Anonymiser uses STX, a fast and efficient XML transformation technology. It should not require a lot of memory to run, even for a large backup.

Development

For Atlassian developers:

- Source code.
- Maven repository.

Enabling Detailed SQL Logging

Confluence uses the open source persistence framework Hibernate. This page tells you how to configure Confluence's logging to report individual SQL requests that are sent to the database by Hibernate. It is useful for

troubleshooting the following events:

- XML site backups that fail to import.
- Exceptions caused by an illegal database operation.

> **Enable SQL logging via the Administration Console**
> Since the 2.7 release, you can also enable SQL logging at runtime via the Administration Console —
> read the instructions. This runtime option does not allow you to enable logging of SQL parameter values.
> To log the parameters as well, continue with the steps below.

To enable detailed SQL logging in Confluence, you need to modify `log4j.properties`, located in `confluence/WEB-INF/classes`.

ℹ After you have enabled hibernate logging, please replicate the action that is causing the error in the first place. This is the best way to ensure that the Confluence log file contains relevant SQL logging.

> If you require support assistance with a database related problem, it is advisable to enable detailed SQL logging **before** sending us the log files. This will assist us in determining what SQL queries were running during the reported problem.

ℹ If the entries mentioned below are not defined in the `log4j.properties file`, you can manually add the entries to the file in the 'Hibernate Logging' section.

To Log SQL Queries with Parameters

1. Stop Confluence
2. Change the following lines in `log4j.properties` from `ERROR` to `DEBUG`:

   ```
   ## log hibernate prepared statements/SQL queries (equivalent to
   setting 'hibernate.show_sql' to 'true')
   ## DEBUG, TRACE or ALL is required to see query strings
   log4j.logger.net.sf.hibernate.SQL=DEBUG
   ```

3. Uncomment the following lines in `log4j.properties`:

   ```
   ## log hibernate prepared statement parameter values. Note: due to
   caching in net.sf.hibernate.type.NullableType, requires restart to
   take effect
   ## TRACE or ALL is required to see parameter values
   ## Note that log4j.appender.confluencelog.Threshold (or other
   appenders) also must be TRACE or ALL to see any trace messages in
   the logs
   log4j.logger.net.sf.hibernate.type=TRACE
   ```

4. Change the following lines in `log4j.properties` from `DEBUG` to `TRACE`:

   ```
   log4j.appender.confluencelog.Threshold=TRACE
   ```

5. Start Confluence.

RELATED TOPICS

Troubleshooting SQL exceptions
Working with Confluence Logs
Generating a Heap Dump

Sometimes you may see that Confluence is holding onto a chunk of memory over a period of time (for example, tenured space is increasing close to Xmx). In such a situation, it is useful to find out what is stacking up in the memory by analysing the heap dump.

On this page:

- Automatically Generating a Heap Dump when Confluence Hits OutOfMemory Error
- Manually Generating a Heap Dump when Confluence Stops Responding
- Submitting a Heap Dump to Atlassian Support

> **Tips when getting a heap dump**
> If you have a large Xmx size, please limit your Xmx size to 1024m. This will help to keep Confluence heap dump smaller while still containing sufficient information to analyse it.

Automatically Generating a Heap Dump when Confluence Hits OutOfMemory Error

Typically, we would like to analyse the heap dump produced when Confluence died from an OutOfMemory Error. For this, you can add additional JVM parameters like below:

```
-XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=<path to this heap dump file>
```

If you do not set the HeapDumpPath parameter, by default the heap dump will be saved in the folder where Tomcat is run from.

ℹ️ If you are setting these parameters in the Windows registry (when running Confluence on Windows as a service), be sure that each parameter is on its own line.

Manually Generating a Heap Dump when Confluence Stops Responding

It is also possible to get a heap dump manually using a JDK bundled tool called jmap, although we recommend that you use the automatic method above for best result.

For Linux/Solaris-based Operating Systems:
Please execute the following command on Linux OS:

```
$JAVA_HOME/bin/jmap -dump:format=b,file=heap.bin <pid>
```

For Windows:
Please find your Confluence process ID (see below) and then execute the command below on a Windows command line:

```
%JAVA_HOME%\bin\jmap -dump:format=b,file=heap.bin <pid>
```

To find out the process ID for your Java process in Windows, you can use Process Explorer from Microsoft. This is what it looks like:

Using Process Explorer to find your Tomcat process ID

Submitting a Heap Dump to Atlassian Support

Please zip the file and then send it to Atlassian Support.

Generating a Thread Dump

If Confluence is performing poorly, behaving unexpectedly or stops responding and you can generate a thread dump to help diagnose the cause of the problem. Including a thread dump can be very useful to the Atlassian Support team when investigating your problems.

The thread dump shows the current state of each thread in the JVM, including a stack trace, so it needs to be taken at the time the application is experiencing problems.

Information about what locks are being held and waited upon by a thread are not included in the simple thread dump that you can generate from within Confluence. If you need this additional information, you'll need to generate a thread dump e xternally.

On this page:

- Stack traces and security
- Generating a thread dump within Confluence
- Generating a thread dump externally

Stack traces and security

To help debug support cases and provide legendary support, Confluence provides stack traces through the web interface when an error occurs. These stack traces include information about what Confluence was doing at the time, and some information about your deployment server.

Only non-personal information is supplied such as operating system and version and Java version. With proper network security, this is not enough information to be considered dangerous. No usernames or passwords are included.

Generating a thread dump within Confluence

To generate a thread dump via Support Tools:

1. Go to

   > **General Configuration** > **Support Tools**.
2. Choose **Create Support Zip**
3. Make sure **Thread Dumps** is selected (you can deselect any other options that you don't want to include - if you're creating the support zip to provide to Atlassian Support, we recommend including as

much information as possible)
4. Click **Create**.

A zip file containing your thread dumps (and anything else you chose to include) will be saved to your home directory.

Generating a thread dump externally

If Confluence stops responding or you require information on locks being held and waited upon by threads, then use one of methods described in Generating a Thread Dump Externally.

Our support team may ask you to use this method if the simple thread dump generated within Confluence does not include enough information to diagnose the problem.

Generating a Thread Dump Externally

If Confluence stops responding and you can't create a thread dump within Confluence (via Support Tools), you can create thread dumps outside the application. External thread dumps are also useful if you require information on locks being held or waited upon by threads.

> **Take Multiple Thread Dumps**
> Typically you'll want to take several dumps about 10 seconds apart, in which case you can generate several dumps and output the stack traces to a single file.

On this page:

- Generating thread dumps on Linux
- Generating thread dumps on windows
    - Generating a thread dump using our scripts
    - Generating a thread dump from the console
    - Generating a thread dump using jstack
- Output
- Thread dump analysis tools

Generating thread dumps on Linux

To generate a thread dump on Linux (or Solaris or other Unixes):

1. Identify the **java** process that Confluence is running in.: This can be achieved by running a command similar to:

   ```
   ps -ef | grep java.
   ```

2. Use the process ID from the above to generate the thread dumps:

   ```
   kill -3 <pid>
   ```

   This won't kill your server (so long as you included the "-3" option with no space in between).

Generating thread dumps on windows

Generating a thread dump using our scripts

We now have scripts for generating a series of thread dumps externally on Windows. Check out this BitBucket Repository for more information.

Generating a thread dump from the console

If you are not running Confluence as a service, you can generate a thread dump directly in the console.

Click the console window and press <CTRL>+BREAK (or SHIFT+CTRL+PAUSE on some keyboards). The thread dump will print directly to the console.

Generating a thread dump using jstack

The JDK (Java Development Kit) includes jstack, which is used for generating thread dumps.

> **Note**: The JRE (Java Runtime Environment) that is bundled with the Confluence installer does not include jstack. You'll need to have the full JDK installed.

To generate a thread dump using jstack:

1. Identify the process. Launch the task manager by, pressing `Ctrl + Shift + Esc` and find the Process ID of the Java (Confluence) process. (If you can't see the PID column right click a column heading in Task Manager and choose PID).
2. Run jstack <pid> to Capture a Single Thread Dump. This command will take one thread dump of the process id <pid>, in this case the pid is 22668:

   ```
   adam@track:~$ jstack 22668 > threaddump.txt
   ```

   This will create a file called threaddump.txt to your current directory.

> Common issues with jstack:
>
> - You must run jstack as the same user that is running Confluence
> - If the jstack executable is not in your $PATH, then please look for it in your <JDK_HOME>/bin directory
> - If you receive `java.lang.NoClassDefFoundError: sun/tools/jstack/JStack` check that tools.jar is present in your JDK's lib directory. If it is not, download a full version of the JDK.
> - If you see the following message: 'Not enough storage is available to process this command,' see this article.
> - If you get the error "Not enough storage is available to process this command", download the 'psexec' utility from here, then run the following command using:
>   ```
>   psexec -s jstack <pid> >> threaddumps.txt
>   ```

Output

Thread dumps appear in the catalina.out file in the application directory's logs folder. Search for "thread dump" in the log file for the beginning of the dump.

Often Support may ask you to generate a sequence of thread dumps over a short period, so that they can compare what each dump contains and to look for any long running threads that could be the cause of the performance issue.

You can manually generate multiple thread dumps by executing the command repeatedly, but it is often easier to use a small script to automate the process. Here's an example that you can adapt to run on your server:

```
for i in `seq 1 10` ; do
  echo ${i}
  your/path/to/jstack `ps aux | grep java | grep confluence | grep -v grep | awk '{print $2}'` >> threaddump.log
  sleep 10
done
```

Thread dump analysis tools

- Thread Dump Analyzer TDA
- Samurai

Getting Java Crash Log File

If you discovered that Confluence died without manual intervention, there may be something goes wrong during a local Java session. Java will produce a log file that looks like the following: `hs_err_pid20929.log`.

The location of this log file is usually in the directory where Tomcat is run eg. `/bin` folder. For Windows Services, it should be located in where Windows Services are run, eg. `c:\Windows\System32` if you are on 32 bit.

Useful VM Option

If using Java 6 or above, it's possible to define the path to the `hs_err_pid` file.

Add the following JVM Parameter to your existing ones:

```
-XX:ErrorFile=./hs_err_pid<pid>.log
```

RELEVANT TOPICS

Generating a Heap Dump
Java Crashes

Profiling using the YourKit Plugin

There is a plugin for Confluence 2.2 and later which allows easy profiling using the YourKit profiler. No license is required to generate a memory or CPU snapshot, but you will need at least an evaluation license to analyse the results.

On this page:

- Configuring YourKit in your JVM
    - Windows Configuration
    - Linux/Mac OS X Configuration
- Performance Impact
- Installing the YourKit Plugin
- Plugin Source Code

Configuring YourKit in your JVM

Download YourKit 6.0 for your platform and follow the installation instructions to install it.
⚠ Note: YourKit version 7 is not compatible with the Confluence yourkit plugin.

The following instructions apply to Confluence and Tomcat installations with the Oracle (previously Sun) JDK. They should be easily applicable to other application servers and JVMs, however. The YourKit documentation covers this in more detail.

Windows Configuration

On Windows, add to your PATH environment variable the `bin/win32/` directory underneath the YourKit installation directory. For example, you might add "C:\Program Files\YourKit Java Profiler 6.0.12\bin\win32" to your PATH, via Control Panel, System, Advanced, Environment Variables.

To configure Confluence's JVM to use the YourKit agent, you need to add a parameter to CATALINA_OPTS in the `bin/setenv.bat` file in your Confluence application directory. This file controls the startup parameters for Tomcat, so you'll need to restart it after making the changes.

Add the 'agentlib' parameter to the end like this:

```
set CATALINA_OPTS=%CATALINA_OPTS% –Xms1024m –Xmx1024m
–XX:MaxPermSize=256m –XX:+UseG1GC –agentlib:yjpagent
```

Linux/Mac OS X Configuration

On Unix-based systems, include the installation directory in the library path environment variable, as shown below:

- For the Mac: `export DYLD_LIBRARY_PATH=$DYLD_LIBRARY_PATH:/path/to/yourKitAgent`
- For other Unix-based systems: `export`

```
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/path/to/yourKitAgent
```

In general, to configure the JDK, you add the `agentlib` parameter:

```
java -agentlib:yjpagent ...
```

You can add this to Tomcat's `bin/setenv.sh` like this:

```
CATALINA_OPTS="$CATALINA_OPTS -Djava.awt.headless=true"
-agentlib:yjpagent "
```

Performance Impact

Running YourKit can have detrimental effects on performance.

To minimize performance problems use the following modifications to the agentlib parameter:

```
-agentlib:yjpagent=disablecounts,disablealloc,disablej2ee
```

See also Profiling overhead: how to reduce or avoid in the YourKit documentation.

Installing the YourKit Plugin

Download the plugin and upload it into Confluence through the Administration, Plugins page.

A new menu option, "YourKit Profiling' will appear under the 'Administration' heading. Click it and you should see the options to take a memory or CPU snapshot.

This profiler dump will be saved to a local temp directory, and the path shown once it is complete. For the CPU snapshot, this will take at least 30 seconds. For the memory snapshot, 10-15 seconds.

Screenshot: You can take a memory or CPU snapshot

Why would I do this?

Analysing a profiler dump enables Atlassian Support (or you, if you are interested) to see exactly what is happening in your application: what classes are using the memory, what is using CPU and so on. This can help us debug tricky performance problems which would otherwise be impossible to analyse remotely.

Take a CPU snapshot if:

- Confluence is sometimes unresponsive
- Pages take a long time to load
- Confluence's CPU usage is peaking.

Take a memory snapshot if:

- Confluence's memory usage is higher than you expect
- You are getting OutOfMemoryError's in your logs.

ℹ️ If you run into situations where Confluence is unresponsive and you are not able to trigger a memory snapshot, please ensure that you start Confluence with the `onexit=memory` parameter in the JVM options (as

in the example below) and simply shut down Confluence. Before shutting down a memory snapshot will be created.

```
-agentlib:yjpagent=onexit=memory
```

Plugin Source Code

The source code for this Confluence plugin is available from Subversion and browseable in Fisheye. The JAR produced by 'mvn package' includes a copy of the YJP redistributable bundled in META-INF/lib/.

Confluence Data Center

This page provides a high level overview of options and considerations for large enterprises using Confluence. If you want to find out how to make sure Confluence can scale with your organisation this information is for you.

For a detailed overview of Confluence's clustering solution see Confluence Data Center Technical Overview. For help with installation, take a look at Installing Confluence Data Center.

On this page:

- Clustered Confluence
- Performance at scale
- Instant scalability
- Achieving high availability
- Failover

Clustered Confluence

Confluence Data Center is designed to support the unique, and complex requirements of enterprise organisations. It is a clustered solution that can provide performance at scale and high availability.
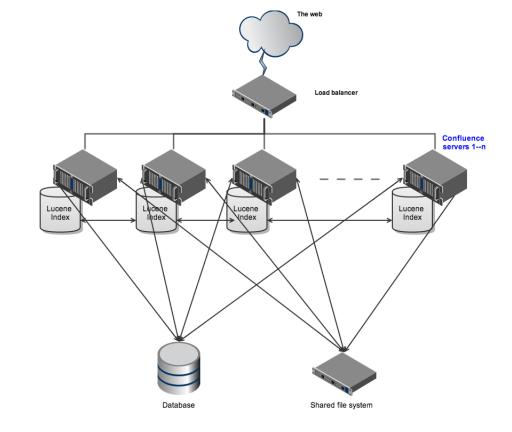
Confluence Data Center enables you to configure a cluster similar to the one pictured here with:

- Multiple server nodes
- A shared file system
- A database that all nodes read and write to.
- A load balancer to evenly direct requests to each node.

All nodes are active and process requests. A user will access the same node for all requests until their session times out, they log out, or a node is removed from the cluster.

Performance at scale

If your Confluence instance has a very heavy load (you have a lot of users accessing Confluence at the same time) a clustered installation will allow you to serve a greater number of concurrent requests than a

single server.

Confluence Data Center spreads the load evenly between cluster nodes enabling you to serve more requests before performance degrades. It's not faster than a single node and won't improve performance in instances with a small number of users.

Instant scalability

Licensing is based on users, not the number of nodes in your cluster. This means you can join additional nodes to your cluster at any time, making it very easy to adapt as your usage grows.

Achieving high availability

High availability, or HA, means that a system is available without interruption. HA is necessary if you have SLA's or other commitments such as 99.99% up time.

If a Confluence Data Center node fails, is shut down or disconnected from the network the rest of the cluster will continue operating - as long as one node remains. When a node fails the load balancer stops sending requests to that node and redirects requests to the remaining nodes. Users will be directed to another node seamlessly.

Some things you can do alongside Confluence Data Center to minimise single points of failure include ensuring that your:

- load balancer
- database
- shared file systems
- external user management

are all highly available. An example of a highly available Confluence installation is shown here.

With all of the above in place, the only time you should need to bring down your entire cluster is to upgrade Confluence itself. Confluence Data Center requires all nodes to be stopped to perform the upgrade.

Most other maintenance can be performed by bringing down individual nodes, as nodes can leave and rejoin the cluster without disruption.

Failover

It is important to note that Confluence Data Center nodes must be able to communicate between one another with minimal latency and therefore must be located in the same physical location.

Warm failover can be achieved by syncing your Confluence data to a data center in a different geographical location and switching over in the event of a failure. This solution is not automatic, and requires manual intervention, but can provide failover in the case of a major failure in your primary data center.

> **Ready to get started?**
>
> Contact us to speak with an Atlassian or get started with Data Center today.

Confluence Data Center Technical Overview

This page provides information on Confluence Data Center, which is a clustered solution that can provide performance at scale and high availability. This is essential if your Confluence instance is mission critical or has very high load.

How it works

The basics

On this page:

- How it works
- Infrastructure and hardware requirements
- Plugins and add-ons

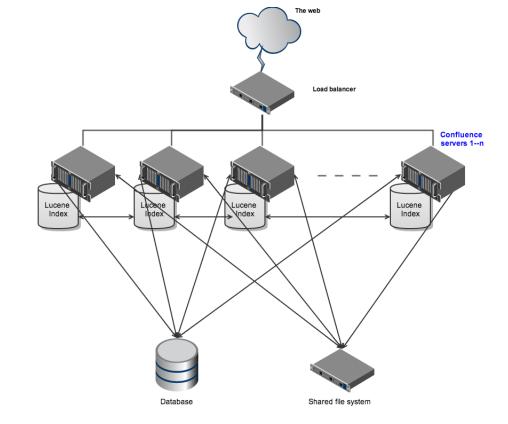Confluence Data Center enables you to configure a cluster similar to the one pictured here with:

- Multiple server nodes that store:
    - logs
    - caches
    - Lucene indexes
    - configuration files
    - plugins
- A shared file system that stores:
    - attachments
    - avatars / profile pictures
    - icons
    - export files
    - import files
    - plugins
- A database that all nodes read and write to.
- A load balancer to evenly direct requests to each node.

All nodes are active and process requests. A user will access the same node for all requests until their session times out, they log out or a node is removed from the cluster.

Licensing

Your Data Center license is based on the number of users in your cluster, rather than the number of nodes. You can monitor the available license seats in the License page.

If you wanted to automate this process (for example to send alerts when you are nearing full allocation) you

can use the REST API.

~ REST API...

The following GET requests require an authenticated user with system administrator permissions. The requests return JSON.

`<confluenceurl>/rest/license/1.0/license/userCount`	Number of active users
`<confluenceurl>/rest/license/1.0/license/remainingSeats`	Number of users you can add before reaching your license limit
`<confluenceurl>/rest/license/1.0/license/maxUsers`	Maximum number of users allowed by your license

Home directories

Confluence has a concept of a local home and shared home. Each 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 node in the cluster. Attachments, icons and avatars are stored in the shared home as are export and import files.

Add-ons can choose whether to store data in the local or shared home, depending on the needs of the add-on.

If you are currently storing attachments in your database you can continue to do so, but this is not available for new installations.

Caching

Confluence uses a distributed cache that is managed using Hazelcast. Data is evenly partitioned across all the nodes in a cluster, instead of being replicated on each node. This allows for better horizontal scalability, and requires less storage and processing power than a fully replicated cache.

Because of this caching solution, to minimise latency, your nodes should be located in the same physical location.

Indexes

A full copy of the Confluence indexes are stored on each node individually. A journal service keeps each index in synch. If you need to reindex Confluence for any reason, this is done on one node, and then picked up by the other nodes automatically.

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 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 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 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 startup process.

Cluster safety mechanism

A scheduled task, ClusterSafetyJob, runs every 30 seconds in Confluence. In a cluster, this job is run on one node only. The scheduled task operates on a safety number – a randomly generated number that is stored both in the database and in the distributed cache used across 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.

This mechanism also exists in standalone Confluence.

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 add-on.

Cluster node discovery

When configuring your cluster nodes you can either supply the IP 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 set-up of the first node.

Infrastructure and hardware requirements

The choice of hardware is up to you. Below are some areas to think about when planning your hardware and infrastructure requirements.

Servers

We recommend your servers have at least 4GB of physical RAM. A high number of concurrent users means that a lot of RAM will be consumed. You usually don't need to assign more than 4GB per JVM process, but can fine tune the settings as required.

You should also not run any 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're using multicast, you can't run Confluence Data Center in Amazon Web Services (AWS) environments as AWS doesn't currently support multicast traffic.

Cluster nodes

Your Data Center license does not restrict the number of nodes in your cluster. We have tested the performance and stability with up to 4 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

- run the same Confluence version
- 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 cluster nodes don't diverge, as it can result in a range of problems with your cluster.

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 instance has a connection pool of 20 connections and you expect to run a cluster with four nodes, your database server must allow at least 80 connections to the Confluence database. In practice, you may require more than the minimum for debugging or administrative purposes.

You should also ensure your intended database is listed in the current Supported 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.

Shared home directory and storage requirements

All 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.

Load balancers

We suggest using the load balancer you are most familiar with. The load balancer needs to support 'session affinity'.

Network adapters

Use separate network adapters for communication between servers. The Confluence cluster nodes should have a separate physical network (i.e. separate NICs) for inter-server communication. This is the best way 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 high availability

Confluence Data Center removes the application server as a single point of failure. You can further minimise single points of failure by ensuring your load balancer, database and shared file system are also highly available.

Plugins and add-ons

The process for installing add-ons in Confluence Data Center is the same as for a standalone instance of Confluence. You will not need to stop the cluster, or bring down any nodes to install or update an add-on.

The Atlassian Marketplace indicates add-ons that are compatible with Confluence Data Center.

Add-on licenses for Data Center are sold at the single server rate, but must match or exceed your Confluence Data Center license tier. For example, if you are looking to have 3,000 people using Confluence Data Center, then you would buy any add-ons at the 2-001-10,000 user tier.

If you have developed your own plugins for Confluence you should refer to our developer documentation on How do I ensure my add-on works properly in a cluster? to find out how you can confirm your plugin is cluster compatible.

> **Ready to get started?**
>
> Contact us to speak with an Atlassian or get started with Data Center today.

Restricted Functions in Confluence Data Center

There are some features that are disabled or limited in Confluence Data Center. 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. 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 **CONF-34680** - Make quick reload plugin available in Confluence Data Center **RESOLVED** for more info.

Application links authentication: • Basic access (http) • Trusted Applications	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	LIMITED	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.

Confluence Data Center Performance

This document describes the performance tests we conducted on 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 your own organization.
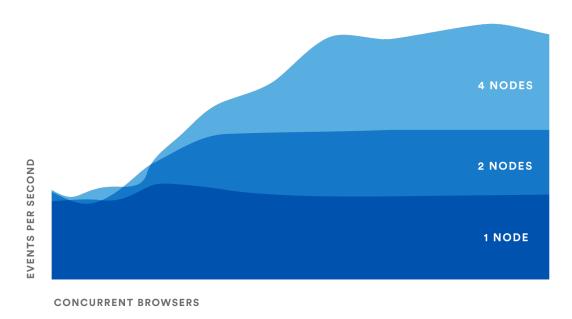
We started our performance tests by taking a fixed load profile (read/write ratio), then tested different cluster set ups against multiples of that load profile.

Below you will find more information on our testing methodology, hardware specs and detailed results.

Testing results summary

- Under a high load, clustered Confluence has improved performance overall.
- Adding more nodes increases throughput, handles higher load and decreases response times.
- You might observe a different trend/behavior based on your configuration and usage. For details, please see What we tested.

Confluence Data Center Performance

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 browsers to 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 show the ratio of actions performed 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.

* The tests did not include admin actions as these are assumed to be relatively infrequent.

Hardware

All performance tests were all run on the same controlled, isolated lab at Atlassian using the hardware listed below.

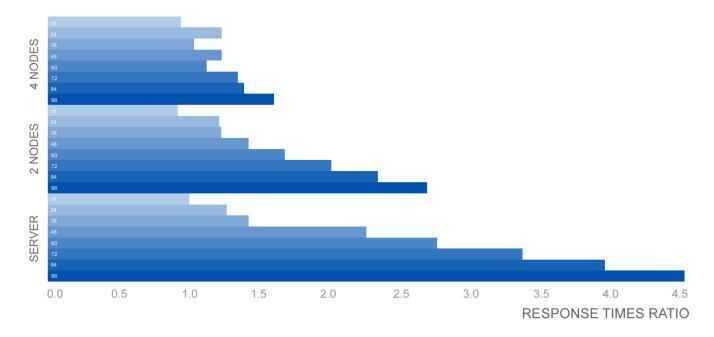| Hardware	Description	How many?
Rackform iServ R304.v3	CPU: 2 x Intel Xeon E5-2430L, 2.0GHz (6-Core, HT, 15MB Cache, 60W) 32nm RAM: 48GB (6 x 8GB DDR3-1600 ECC Registered 2R DIMMs) Operating at 1600 MT/s Max NIC: Dual Intel 82574L Gigabit Ethernet Controllers - Integrated Controller: 8 Ports 3Gb/s SAS, 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	20
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.
- 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 us to speak with an Atlassian or get started with Data Center today.

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 ensuring Confluence's availability in the event that your primary site becomes unavailable.

Confluence Data Center is the only Atlassian-supported high-availability solution for Confluence. However, if you don't choose Confluence Data Center, our Experts may be able to help implementing a high availabili ty solution for your environment. Contact our Experts team for more information.

This page demonstrates how you can use Confluence Data Center 5.9 or later in 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 & RCO[1]), 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 need Confluence Data Center 5.9 or later to implement the strategy described in this guide.

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 organisation.

> **Maintaining a runbook**
> The detailed steps will vary from organisation to organisation 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. Synchronisation 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 RPO[1]
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. Install Confluence Data Center 5.9 or 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 standb y 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: Show me... • 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 strategy that 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.
Starting the standby cluster	It's important to **initially start 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.
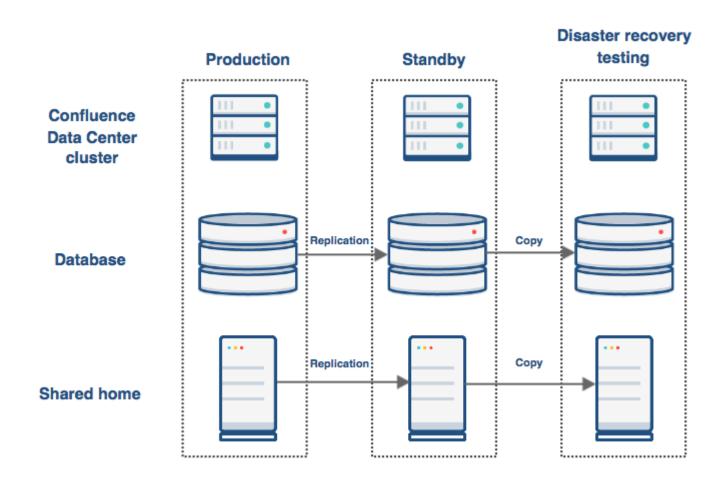
The key is to **keep the main data center as isolated as possible from the disaster recovery testing**
.

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	1. Temporarily pause all replication to the standby database 2. 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	1. Clone your standby installation separate from both the live and standby instances 2. Change the connection to the database in the `<confluencelocalhome>/conflu ence.cfg.xml` file to avoid any conflict 3. Change the location of the shared home directory in the `<confluencelocalhome >/confluence.cfg.xml` file to avoid any conflict 4. If using TCP/IP for cluster setup, change the IP addresses to that of your testing instances in `<confluencelocalhome>/confluence.cfg.xml`

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 latest snapshot and no replication
2. Ensure that the new shared home directory is ready, with the latest snapshot and no replication
3. Ensure you 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.home mapped, as it was in the standby instance, in the test server
5. Disable email (See `atlassian.mail.senddisabled` in Configuring System Properties)
6. Start Confluence

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 Confluence on 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:

- Synchronise your primary database with the state of the secondary
- Synchronise 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

Atlassian Answers

Our community and staff are active on Atlassian Answers. Feel free to contribute your best practices, questions and comments. Here are some of the answers relevant to this page:

- Disaster Recovery Configuration for Jira and Confluence

Troubleshooting

If you encounter problems after failing over to your standby instance, check these FAQs for guidance:

⌄ What should I do if my database isn't synchronized correctly?
  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 either 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.

⌄ What should I do if my search index is corrupt?
  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.

⌄ What should I do if attachments are missing?
  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.

⌄ What happens to my application links during failover?
  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.